Fachbereich Erziehungswissenschaften und Psychologie

der Freien Universität Berlin

Ecologically Rational Strategy Selection

Dissertation

zur Erlangung des akademischen Grades

Doktor der Philosophie

(Dr. Phil.)

vorgelegt von

Dipl. Psych.

Julian Nicolas Marewski

Datum der Disputation 30.01.2009

Erstgutachter: Prof. Dr. Gerd Gigerenzer

Zweitgutachter: Prof. Dr. Arthur M. Jacobs

Berlin, 2010

Nothing that we have discovered about memory requires us to revise our basic verdict about the complexity or simplicity of human cognition. We can still maintain that,

> Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity the environment in which we find ourselves …

provided that we include in what we call the human environment the cocoon of information, stored in books and in long-term memory, that we spin about ourselves.

(Herbert A. Simon, 1996, pp. 109–110.)

Acknowledgements

This dissertation is the result of research I carried out as a predoctoral research fellow at the Center for Adaptive Behavior and Cognition (ABC), located at the Max Planck Institute for Human Development.

Some people are special. Lael Schooler, my main advisor for all these years, is one such person. I want to thank Lael for being an incredible source of knowledge and intellectual stimulation, and for all his patience with my unpredictable requests for advice. Lael has been a tolerant and generous teacher, investing a great deal of effort and brainpower in our joint projects. I have learned an enormous amount about science from him, ranging from how to ask good questions to how to build models to try to answer them. (At least I hope I have improved in this regard.) Another special person is ABC director Gerd Gigerenzer. Over the years, Gerd has planted many ideas in my mind that have found their way into this dissertation, and I hope there will be many more to come. He is a tremendous source of inspiration. His door is always open, and he always takes the time to listen to and teach his students. I would like to thank him for all the feedback and input he has provided me during my time at ABC. I would also like to thank him for creating this very exciting research environment here at ABC and, particularly, for giving me the opportunity to work here.

Two colleagues deserve special recognition. I would like to thank Wolfgang Gaissmaier for being a great colleague in my years at ABC. Wolfgang, I hope that we will continue to work on many more projects together. Maybe one day we will also publish one of those *Science* papers we envisioned! I would like to thank Henrik Olsson, as well. Henrik, it has been great fun to write a paper with you. When will we write the next one?

There are many more special people. In particular, I would also like to thank those colleagues who commented on earlier drafts of the chapters of this dissertation, and on the papers that are based on them. Once more, I would like to mention here Wolfgang Gaissmaier, Gerd Gigerenzer, and Lael Schooler, who, together with Mirta Galesic, Daniel Goldstein, and Henrik Olsson, have been my coauthors on one or another of the papers on which this dissertation is based. Moreover, I am indebted to Uwe Czienskowski, Nadine Fleischhut, Bettina von Helversen, Konstantinos Katsikopoulos, Christiane Marewski, Dieter Marewski, Hansjörg Neth, Magnus Persson, Jose Quesada, Jeffery Stevens, and Jenny Volstorf for many constructive comments on earlier drafts of the chapters and papers of this dissertation, and to all my former and current colleagues at ABC for many stimulating conversations! I would also like to thank Christel Fraser, Julia Schooler, and Anita Todd for editing this dissertation and/or the papers that are based on it. I would like to thank Gregor

Caregnato for helping me run almost a thousand subjects, some of them through more than a thousand trials—thank you, Gregor! I would like to thank Christian Elsner for helping me out with any computer problems I may have had, and, importantly, for being around, contributing to making ABC a wonderfully friendly and fun workplace.

On a more personal note, I would like to dearly thank my family. How could I have gotten to the point of being able to submit a dissertation thesis without you? The most special person in this long list of special people is Begoña Mera del Cabo. Bego, I would like to apologize to you for all the suffering you went through, for instance, when I—once more—worked late. Thank you for giving me all your support throughout the past years, and for repeatedly reminding me that there is a life beyond science. Without you, Bego, this dissertation would not exist.

Julian Marewski, Berlin, November 2008

QUE LE SORPRENDE MÁS DE LA HUMANIDAD?

LOS HOMBRES PORQUE PIERDEN
LA SALUD PARA GANAR DINERO;
DESPUÉS PIERDEN EL DINERO
PARA RECUPERAR LA SALUD...

Y POR PENSAR ANSIOSAMENTE
EN EL FUTURO NO DISFRUTAN EL PRESENTE,
POR LO QUE NO DISFRUTAN,
NO VIVEN NI EL PRESENTE NI EL FUTURO

Y VIVEN COMO SI NO TUVIESEN
QUE MORIR NUNCA...
Y MUEREN COMO SI NUNCA HUBIERAN VIVIDO.
DALAI LAMA

Table of Contents

Chapter 1

Introduction: The Goal and Contents of This Dissertation

At the time of writing this dissertation, international financial markets are in turmoil. Large banks are going bankrupt almost daily. Today, September 30, 2008, the Dow Jones has crashed more than 700 points—the largest intraday course decline in its history. It is a difficult situation for financial decision makers—regardless of whether they are lay investors, trying to make small-scale profits here and there, or professionals employed by the finance industry. To safeguard their investments, they need to foresee uncertain future economic developments, such as which investments are likely to be the safest harbors and which companies are likely to crash next. In times of rapid waves of potentially devastating financial crashes, these informed bets must often be made quickly, with little time for extensive information search or computationally demanding calculations of likely future returns. Especially lay stock traders have to trust the contents of their memories, relying on incomplete, imperfect knowledge and facts that are quickly accessible, for example, from a news ticker.

Humans are not omniscient. They do not come equipped with the ability to run computationally demanding calculations quickly in the mind. Rather, we make decisions under the constraints of limited information processing capacity, knowledge, and time—be they about the likely performance of stocks; which movie to watch in the cinema, for example, when several are about to start; whom to court in a speed-dating session, or whether to admit to the hospital a patient who has registered at the emergency room reception. According to the *fast and frugal heuristics research program* (Gigerenzer, Todd, & the ABC Research Group, 1999), humans can nevertheless make such decisions successfully because they can rely on a large repertoire of simple decision strategies, or *heuristics*. These simple rules of thumb can perform well even under the constraints of limited knowledge, time, and information-processing capacity because they exploit the structure of information in the environment in which a decision maker acts and build on the ways evolved cognitive capacities work, such as the human memory system. Importantly, together, these simple rules of thumb form an *adaptive toolbox* of the cognitive system, where the tools are heuristics a decision maker uses to respond adaptively to different decision situations, each one appropriate for a given task.

However, even though it is an important assumption of the fast and frugal heuristic approach that decision makers respond to different decision situations by selecting the heuristic that is appropriate for the task, relatively little is known about *how* such a choice is made. The goal of my dissertation is to contribute to our understanding of the corresponding mechanisms of *heuristic choice*—or, to use a more general term, *strategy selection*. Specifically, my dissertation focuses on the selection of decision strategies for making inferences about unknown quantities and uncertain events in situations in which all available information must be retrieved from memory (i.e., *inferences from memory*; see Gigerenzer & Goldstein, 1996). In doing so, I investigate how the interplay between the human memory system, the environment in which decision makers act, and available decision strategies can lead to the emergence of adaptive mechanisms of heuristic selection.

My research thus brings together different theories, namely, about memory, decision environments, decision strategies, and strategy selection. While my work on decision strategies and the environment is grounded in the fast and frugal heuristics research program, parts of this dissertation will show how this framework can be combined with another ecological approach to psychology, that is, with John R. Anderson and colleagues' *ACT-R* (*adaptive control of thought–rational*) *cognitive architecture*, (e.g., Anderson et al., 2004). ACT-R provided me with an ecologically grounded and quantitative model of memory. A number of other ecological theories have also directly or indirectly influenced this dissertation work. For instance, James J. Gibson's (e.g., 1979) writings offered me a heuristic way of thinking about what questions one might ask about environmental structure and strategy selection. Importantly, it is the fit between human cognition and the environment that binds together the different approaches taken here and that is exemplified by what Gerd Gigerenzer and colleagues call *ecological rationality*, defined as the "adaptive behavior resulting from the fit between the mind's mechanisms and the structure of the environment in which it operates" (Todd & Gigerenzer, 2000, p. 728).

In this chapter, I will first give a short introduction to Gerd Gigerenzer and colleagues' fast and frugal heuristics approach, which is the central theory in the context of this dissertation. Second, I will briefly review Anderson and colleagues' ecological approach to human cognition, providing a quick glance at the second framework that will play a major role in the theory development reported below. Last, I will offer an overview of the contents of this dissertation.

Ecologically Rational Heuristics

In which stocks to invest, which movies to watch, whom to court, and what to eat—our days are filled with decisions, yet how do we make them? The answer to this question depends on one's view of human rationality because this, in turn, determines what kinds of models of cognitive processes one believes represent people's decision strategies. There are at least two major approaches.

Visions of Rationality

Unbounded rationality. The study of *unbounded rationality* asks the question, if people were omniscient and omnipotent, that is, if they could compute the future from what they know, how would they behave? The maximization of subjective expected utility is one suggestion (e.g., Edwards, 1954). When judging, for instance, in which stocks to invest, such models assume that decision makers will collect and evaluate all information, weight each piece of it according to some criterion, and then combine the pieces to reach the mathematically optimal solution to maximize the chance of attaining their goals (e.g., profit maximization). Typically, unbounded rationality models assume unlimited time to search for information, unlimited knowledge, and large computational power (i.e., information-processing capacity) to run complex calculations and compute mathematically optimal solutions. These models are common in economics, optimal foraging theory, and computer science.

Bounded rationality. According to the second approach, unbounded rationality models are implausible descriptions of how people make decisions. Our resources—knowledge, time, and computational power—are limited. Herbert Simon (1956, 1990), the father of this *bounded rationality* view, argued that people rely on simple strategies to deal with situations of sparse resources. One research program that is often associated with Simon's work is the *heuristics-and-biases framework* (e.g., Kahneman, Slovic, & Tversky, 1982; Tversky & Kahneman, 1974), which proposes that humans rely on rules of thumb, or heuristics, as cognitive shortcuts to make decisions.[1] Even though this program thus differs from the unbounded rationality view, it still takes unbounded rationality models—such as maximization of subjective expected utility models—as the normative yardstick against which to evaluate human decision making. According to the heuristics-and-biases tradition, decisions deviating from this normative yardstick can be explicated by assuming that people's heuristics are error prone and subject to systematic cognitive biases. Conversely, people's use of heuristics explains why decisions can be suboptimal, or irrational, when compared to the normative yardstick. In short, in this tradition, the term bounded rationality mainly refers to the idea that limitations in our cognitive abilities, in our knowledge, and in other reasoning resources produce

[1] Kahneman et al. (1982) credited Simon in the preface to the anthology although their major early papers, which appear in the anthology, do not cite Simon's work on bounded rationality. Thus, this connection was possibly made in hindsight (Lopes, 1992).

errors, biases, and judgmental fallacies (for a discussion of the "irrationality" rhetoric of the heuristics-and-biases tradition, see Lopes, 1991).

However, H. A. Simon (e.g., 1990) not only stressed the cognitive limitations of humans and proposed simple strategies that we may rely on but also emphasized how the strategies are adapted to our decision-making environment: "Human rational behavior … is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor" (p. 7). The fast and frugal heuristics research program (e.g., Gigerenzer et al., 1999) has taken up this emphasis. In this framework, the term bounded rationality conveys the idea that by exploiting the structure of information available in the environment, heuristics can lead to good decisions even in the face of limited knowledge, computational power, or time. This approach thus shares with the heuristics-and-biases program the idea that people rely on heuristics to make decisions, but dispenses with the normative yardsticks that are used in the heuristics-and-biases tradition to invoke cognitive deficits and irrational errors. Instead, the fast and frugal heuristics framework has developed an ecological view of rationality through which it tries to understand *how* and *when* heuristics result in adaptive decisions. In this view, heuristics are ecologically rational with respect to the environment, and being rational here means that a heuristic is successful with regard to some outside criterion, such as accuracy or speed of prediction.[2]

Four Questions About Heuristics: What Heuristics Are Used, When Are They Used, Where Should They Be Used, and How Can They Help?

Research in the fast and frugal heuristics program focuses on four interrelated questions (see Gigerenzer, Hoffrage, & Goldstein, 2008).[3] The first two questions are descriptive and concern the adaptive toolbox: What heuristics do organisms use to make decisions? When do people rely on a particular heuristic from the toolbox, that is, when and how are different decision strategies from the available repertoire selected? The third question is prescriptive and deals with ecological rationality: To what environmental structures is a given heuristic adapted—that is, in what situations does it perform well, for example, by allowing us to make accurate, fast, and effortless decisions? In contrast to these three theoretical questions, the fourth question focuses on practical applications: How can the

[2] Hammond (1996) called such outside criteria *correspondence criteria* as opposed to *coherence criteria*, which take the laws of logic or probability theory as a normative yardstick for rationality.

[3] Gigerenzer et al. (2008) list three questions rather than four, counting the above-mentioned first and second questions as one. While the two questions may not be answerable independently of each other, I do believe they have slightly different emphases: The first question is primarily concerned with the descriptive adequacy of a heuristic as a model of behavior. The second question, in turn, is concerned with identifying the conditions that determine *when* a heuristic is used, rather than *whether* it is used at all.

study of people's repertoire of heuristics and their fit to environmental structure aid decision making in the applied world?

Ecologically rational heuristics are studied in diverse domains, including more applied areas, such as first-line antibiotic prescriptions for children (Fischer et al., 2002), the improvement of coronary care unit allocations (Green & Mehr, 1997), and risk communication in medicine and law (Gigerenzer, 2002; Gigerenzer & Edwards, 2003; Hoffrage, Lindsey, Hertwig, & Gigerenzer, 2000). At the same time, the fast and frugal heuristics approach is discussed in several branches of science, including the law (e.g., Gigerenzer & Engel, 2006), philosophy (e.g., Bishop, 2006), and biology (e.g., Hutchinson & Gigerenzer, 2005). In basic research, this program has proposed a range of heuristics for different tasks—parental investment (Davis & Todd, 1999), mate search (Todd & Miller, 1999), estimation (Hertwig, Hoffrage, & Martignon, 1999), inferential judgments (e.g., Gigerenzer & Goldstein, 1996; Goldstein & Gigerenzer, 2002), categorization (Berretty, Todd, & Martignon, 1999), and choices between risky alternatives (Brandstätter, Gigerenzer, & Hertwig, 2006), to name a few. Moreover, it has produced a large amount of research investigating whether people, both young and old, rely on given heuristics (Bröder & Schiffer, 2003; Mata, Schooler, & Rieskamp, 2007; Pachur, Bröder, & Marewski, 2008; Pachur & Hertwig, 2006; Pohl, 2006; Rieskamp & Hoffrage, 1999, 2008; Rieskamp & Otto, 2006), under what environmental structures the heuristics perform well (e.g., Gigerenzer & Goldstein, 1996; Hogarth, & Karelaia, 2007; Katsikopoulos & Martignon, 2006b; Martignon & Hoffrage, 1999), and how accurate they are for predicting quantities and events in the real world, such as the outcomes of sports competitions (e.g., Pachur & Biele, 2007; Scheibehenne & Bröder, 2007; Serwe & Frings, 2006), the performance of stocks on the stock market (Borges, Goldstein, Ortmann, & Gigerenzer, 1999), the results of political elections (see Chapter 4), or professors' salaries (Czerlinski, Gigerenzer, & Goldstein, 1999; see also Brighton, 2006).

ACT-R: A Unified Theory of Cognition

Research in the fast and frugal heuristics framework places emphasis on specifying precise formal models of heuristics that can be submitted to vigorous testing. For instance, in a two-alternative choice situation, say, whether to read this dissertation or another one, the computer code or mathematical equations formally specifying a model of a heuristic decision strategy should predict both which alternative will be chosen and how different reasons to choose one alternative over the other will be processed in order to derive a decision. In fact, a number of research programs in judgment and decision making and other disciplines take a similar approach, precisely specifying formal models of behavior (e.g., Busemeyer, & Myung, 1992; Dougherty, Gettys, & Ogden, 1999; Hintzman, 1988; Payne, Bettman, & Johnson, 1993; Raaijmakers & Shiffrin, 1981; Ratcliff, Van Zandt, & McKoon, 1999; Rumelhart, McClelland, & the PDP Research Group, 1986). Yet, one dimension along which many research programs diverge is their generality; that is, one may be able to roughly classify formal models into at least two categories: those that focus on capturing the aspects of an isolated task (e.g., the Stroop task) or phenomenon (e.g., probability matching) and those that *in addition* are integrated into an overarching architecture that formally specifies the assumptions of a broader theory (e.g., about cognition in general). Arguments for developing more integrative systems are provided by A. Newell (1990):

> A single system (mind) produces all aspects of behavior. It is one mind that minds them all. Even if the mind has parts, modules, components, or whatever, they all mesh together to produce behavior. Any bit of behavior has causal tendrils that extend back through large parts of the total cognitive system before grounding in the environmental situation of some earlier times. If a theory covers only one part or component, it flirts with trouble from the start. It goes without saying that there are dissociations, independencies, impenetrabilities, and modularities. These all help to break the web of each bit of behavior being shaped by an unlimited set of antecedents. So they are important to understand and help to make that theory simple enough to use. But they don't remove the necessity of a theory that provides the total picture and explains the role of the parts and why they exist. (pp. 17-18)

In the past decades, a number of more overarching architectures have been developed (e.g., *EPIC*, Meyer & Kieras, 1997; *Soar*, A. Newell, 1990). Some of these theories are integrative enough to allow modeling the most diverse phenomena, ranging from simple syllable counting to driving behavior, within the same system. The ACT-R architecture (e.g., Anderson & Lebiere, 1998) is one such system that, as is important in the context of this dissertation, allows modeling the interplay between the contents of people's memories, the environment in which they act, and the decision strategies they employ.

The Development of the ACT-R Theory of Cognition

ACT-R is a broad quantitative theory of behavior that covers much of human cognition. Its core is constituted by a set of modules, each of which is devoted to processing a different kind of information. For instance, there is a goal module for keeping track of intentions, a declarative module for information retrieval from memory, a visual module for identifying objects in the visual area, and a manual module for executing motor commands, such as hand movements. These modules are coordinated through a production system. That is, the overall theory consists of *production rules* (i.e., if–then rules) whose conditions (i.e., the "if" parts of the rules) are matched against the contents of a module. If the conditions of a rule are met, then the rule fires and the actions specified in the "then" part are carried out. Specifically, with ACT-R researchers can derive predictions of at least three kinds of data: (1) overt behavior, such as the outcomes of decisions; (2) the temporal aspects of the behavior, such as time involved in making a decision; and (3) the associated patterns of activity in the brain, as measured with functional magnetic resonance imaging (fMRI) scanners. The temporal resultion of the system lies at the millisecond level.

Since the 1970s, ACT-R has been repeatedly modified in order to be able to account for new phenomena, a fact that is also reflected in small changes in the theory's name (e.g., ACT, ACT*, ACT-R). For instance, ACT has its origins in the *human associative memory* theory (*HAM*, Anderson & Bower, 1973). This theory did not deal with the many different types of knowledge a person can have about the world. In 1976, Anderson suggested distinguishing between *declarative knowledge* (knowing that), which HAM dealt with, and *procedural knowledge* (knowing how), which HAM ignored. Based on ideas from A. Newell (e.g., 1973a), production rules were used to implement such procedural knowledge. The result was a production system called *ACTE*, which was later replaced by ACT* (Anderson, 1983). Among other things, this new system incorporated assumptions about how production rules might be acquired. The development of ACT-R followed. ACT-R embodied the insight that the cognitive system can give rise to adaptive processing by being tuned into the statistical structure of the environment. Just as in the fast and frugal heuristics program, research in ACT-R is thus concerned with studying the interplay between human cognition and its surroundings. In the context of this dissertation, the ecological foundations of ACT-R's memory system are particularly important.

The Ecological Foundations of ACT-R's Theory of Memory

Much of the ecological component of ACT-R's memory system has its theoretical roots in what Anderson and colleagues called the *rational analysis* of memory (e.g., Anderson & Milson, 1989; Anderson & Schooler, 1991, 2000; Schooler & Anderson, 1997; see also Chapter 3). One key tenet of this analysis is that memories are retrieved as a function of how likely it is that they will be needed to achieve some processing goal. According to Anderson and colleagues' analysis, in doing so human memory capitalizes on a person's history of past encounters with objects (e.g., stock names), which, in turn, can be indicative of how likely objects are to reoccur in the environment and be needed in the future. In their view, human memory essentially makes a bet, namely, that as the recency and frequency with which a piece of information has been encountered increases, so too does the probability that this information will be needed to achieve a given processing goal in the future. Conversely, the more time that has passed since an object has been encountered, the less is the likelihood that memories of the object will need to be retrieved in the future and, ultimately, memories of such objects can be forgotten. This way, memory drops outdated, largely irrelevant information and gives a retrieval advantage to recently and frequently encountered, most likely more relevant information. To illustrate their point, Anderson and Schooler (1991) analyzed environments consisting of text and word utterances. For instance, they observed that the probability of a particular word utterance decreases as a function of the amount of time that passes since the word was last uttered. Similarly, the likelihood of recalling a memory of a given object drops as a function of the amount of time since the object was last encountered. In fact, in various environments, they found strong correspondences between regularities in the patterns of occurrence of information (e.g., a word's recency and frequency of occurrence) and the classic forgetting and learning functions (e.g., as described by Ebbinghaus, 1885/1964). Recently, Schooler and Hertwig (2005) tied the fast and frugal heuristics research program to the ACT-R theory of cognition (see also Nellen, 2003; see Gaissmaier, Schooler, & Mata, 2008, for a review): In a series of computer simulations, they implemented two simple inference heuristics in ACT-R. They then used the architecture to derive quantitative predictions as to when the forgetting of memories can help a person using these heuristics to make accurate decisions, illustrating how cognitive capacities, such as human memory, interplay with people's decision strategies. It is this line of research that my dissertation extends by showing how memory can lead a person to use different heuristics adaptively, aiding the selection of different decision strategies from the adaptive toolbox.

Organization of the Dissertation

Specifically, my dissertation tackles the first three questions of the fast and frugal heuristics framework, namely, (i) what heuristics people use, (ii) when they use them, and (iii) to which environmental structure the heuristics are adapted. As will become clear in the chapters, none of these questions can in fact be answered independently of the others, and, given the focus of this dissertation on memory-based decision making, all of them require a theory of memory.

The dissertation is organized as follows. Chapter 2 briefly discusses different ways to evaluate formal models. This way, I set the methodological preliminaries for the theoretical, simulation, and experimental work to follow in Chapters 3 and 4. Impatient readers are invited to skip this introductory discussion and turn directly to Chapters 3 and 4. The main foci of Chapter 3 are the second and third questions. This chapter outlines a way in which the interplay between the cognitive system and the environment can aid the selection of different heuristics from the adaptive toolbox. Importantly, in this chapter, I will not only focus on models of decision making but will also propose and test a formal model of memory that quantifies the ways in which the human memory system interacts with the environment. Feeding this memory model with environmental data, I will be able to make systematic quantitative predictions concerning memory retrieval, such as on the probabilities of retrieving certain memories about objects in our world (e.g., stock names) as well as about the associated retrieval time distributions. By submitting people's reliance on different heuristics to tests, the simulation studies and experiments reported in this third chapter also help to answer the first question, that is, to identify which heuristics people use. In Chapter 4, this question becomes even more important: Here, I will pit competing models of decision strategies against each other, systematically evaluating the descriptive adequacy of a particularly simple heuristic as a model of behavior. In addition, I will propose extensions of this heuristic and, by addressing the second and third questions once more, tackle the problem of how this heuristic might be selected from the adaptive toolbox as a function of the environment in which a decision maker acts. As the theoretical implications of the experiments and simulation studies reported in Chapters 3 and 4 will be discussed in detail in these chapters, in Chapter 5 I will limit myself to summarizing the main findings and will briefly return to the topic of model testing, previously addressed in Chapter 2.

Please note that the chapters of this dissertation were written to be understandable on their own, meaning that important concepts will be discussed in more than one place in the text. Therefore, there may be some redundancies between individual chapters, which, however, rather than being obstructive, may help the reader keep track of theses and findings when reading this dissertation.

Chapter 2

Methodological Preliminaries

Introduction

If theories are underspecified, then they can be used post hoc to "explain" lots of diverse observed data patterns. In the worst case, they become *one-word explanations*, labels that are broad in meaning and hence vague, and that therefore provide little or no specification of the underlying mechanisms or theoretical structure (see Gigerenzer, 1996, 1998). Consider the *representativeness heuristic* in the field of judgment and decision making (Kahneman & Tversky, 1972). A probability assessed by this decision strategy, say, whether a newly encountered animal is a dog, is derived from how representative this animal is of the target category—in this case, dogs. However, exactly *how* the category is represented or *how* representativeness is derived was not defined when the heuristic was proposed. This vagueness made it possible to apply the notion of representativeness to a wide range of phenomena, such as misperception of regression, the conjunction fallacy, and base-rate neglect. At the same time, this lack of specification made it difficult—if not impossible—to test (see e.g., Ayton & Fischer, 2004). In fact, after the definition of the heuristic was finally somewhat strengthened (see Kahneman & Frederick, 2002), a number of studies found that theories assuming different psychological processes outperform this heuristic in predicting people's behavior (e.g., Nilsson, Olsson, & Juslin, 2005).

One good way to make theories more precise is to cast them as formal models. Models can provide strong bridges between theories and empirical evidence, for instance, by enabling researchers to test competing quantitative predictions against each other.[4] In Chapters 3 and 4 of this dissertation, I will report several tests of formal models. In order to offer an introduction to these tests, I will next discuss different ways of evaluating models. As already mentioned above, readers who feel impatient about reading these introductory remarks, are invited to directly turn to Chapters 3 and 4.

[4] I use the term "predicting" (or "prediction") to refer to situations in which a model's free parameters are fixed such that they cannot adjust to the data on which the model is tested. In contrast, I use the term "fitting" to refer to situations in which a model's parameters are allowed to adapt to the test data. In doing so, I use the term "quantitative" prediction in a broad sense to refer to both categorical and numerical statements. For instance, the ACT-R model developed in Chapter 3 allows numerically predicting response time distributions. The heuristics tested in Chapters 3 and 4 make categorical predictions as to which one of two objects people choose in a decision task. In addition, they can make numerical predictions as to how accurate such decisions will be. Note that there are also ways to evaluate models based on qualitative predictions (see Pitt, Kim, Navarro, & Myung, 2006).

Overview of the Chapter and Introductory Definitions

Before going into detail, I would like to comment on the scope of this chapter. Many have written about the complications and merits of formal modeling (e.g., Forster, 2000; Hintzman, 1991; Jacobs & Grainger, 1994; Lewandowsky, 1993; Pitt, Myung, & Zhang, 2002; Roberts & Pashler, 2000; Sedlmeier & Renkewitz, 2007). Given space limitations, and given that the main focus of this dissertation is a different one, my intention here is to provide a very short overview of selected issues, going into only as much detail as might be useful for understanding the rationales of the subsequent model tests.

This chapter is organized as follows. First, I will briefly define the term model and comment on the goals of modeling. Second, I will explain the problem of overfitting, which is a major complication that can arise when testing models. Third, I will discuss different ways of selecting between competing models that can help avoid this problem.

What Is a Model?

In the broadest sense, a model is a simplified representation of the world that aims to explain observed data. Countless verbal, that is, informal explanations of psychological phenomena fit this definition. In a more narrow sense, a model is a formal instantiation of a theory that specifies the theory's predictions, for example, in mathematical equations or computer code.[5] This category also includes statistical tools, such as structural equation or regression models. Unless one believes that the mind works like a regression analysis or other statistical procedure, such tools are not typically meant to mirror the workings of psychological mechanisms, say, those determining *how* a person processes information (but see Gigerenzer, 1991, for examples of theories inspired by statistical tools). In this dissertation, I mainly discuss *algorithmic-level theories* (Marr, 1982), that is, formal instantiations of theories that are designed with the goal of reflecting psychological processes, although the subsequent points about model selection could also be made for *computational-level theories* that aim to explain the functional goals of behavior.

[5] Sometimes mathematical models are contrasted with computer models. For instance, according to Fum, Del Missier, and Stocco (2007), mathematical models can be used to describe a phenomenon but they do not reproduce it, whereas computer models can produce observable behavior. For the sake of simplicity, I do not distinguish between these different types of formal models in this chapter. Mathematical models can be implemented in computer code and some computer models can be expressed in terms of mathematical equations. Similarly, sometimes one can derive a theory's predictions through both mathematical analysis and computer simulation; however, for more complex theories one typically has to rely on computer simulations.

What Is the Scope of Modeling?

Formal modeling is not meant to be applied equally to all questions. As most research approaches, modeling should be seen as a tool tailored to specific problems that scientists should pull out of their methodological toolbox whenever it is most advantageous to do so. For instance, when investigating which of two treatments for depression is more powerful, it might be pointless to model the processes underlying each treatment's effectiveness. Instead, it may be better to examine differences between groups of patients receiving one treatment or the other, using meta-analysis as research tool. Modeling is especially suited for basic research about the cognitive system. Here, it has been used to investigate a large variety of phenomena, ranging from orthographic processing in visual word recognition (e.g., Grainger & Jacobs, 1996), to strategy selection (e.g., Rieskamp & Otto, 2006)—to name just two examples, as many more will follow in the chapters below.

How to Select Between Formal Models: A Short Overview

Consider two models that compete as explanations for a behavior in a task. How can one decide which model provides a better explanation for the data? This comparison of alternative models is called *model selection*. Model selection can have various technical meanings in different fields, but for my purposes it suffices to say that it is the task of choosing a model from a set of potential models, given available data.

A number of model selection criteria are available (see e.g., Jacobs and Grainger, 1994, for a detailed overview). These include falsifiability, that is, whether the models can be proven wrong, and the number of assumptions the models make. For instance, one could ask which of many competing models accounts for the data in the simplest way. In Chapter 4, I will rely on this criterion when pitting simpler heuristics against more complex ones. Other criteria address standards for psychological plausibility, such as whether the computations postulated by a model are tractable in the world beyond the laboratory (e.g., Gigerenzer et al., 2008). Moreover, one could also ask whether a model is consistent with overarching theories of cognition (see e.g., Dougherty, Franco-Watkins, & Thomas, 2008). Integrative architectures, such as ACT-R, can impose precise theoretical constraints on which models represent acceptable developments of a theory. As a matter of fact, when I developed the ACT-R memory model described in Chapter 3, I was careful not to violate the basic assumptions of the architecture. Possibly the most widely used model selection criterion is a model's *descriptive adequacy*—which is the yardstick for model selection I will focus on in the remainder of this section and which will also play a major role in all subsequent model tests. Often descriptive adequacy is evaluated in terms of goodness of *fit*, that is, when two or more models are compared, the model that provides the smallest deviation from existing data, measured, for instance, in terms of R^2, is favored over a model that results in a larger deviation from that data. Yet, there is a limitation of model selection procedures that are based exclusively on such measures of fit.

The Problem of Overfitting

To conclude that one model provides a better account of data than another based on R^2 or other standard goodness of fit indices might be reasonable if psychological measurements were noise free. However, noise-free data are practically impossible to obtain. Hence, researchers are confronted with the problem of disentangling the variation in data due to noise from the variation due to the psychological process of interest. Goodness of fit measures alone cannot make this distinction. As a result, a model can end up *overfitting* the data; that is, it can capture not only the variance due to the cognitive process of interest but also that from random error. Figure 2.1 illustrates a situation in which one model, call it Model A (thin line), overfits existing data by chasing after idiosyncrasies in that data. This model fits the existing data (filled circles) perfectly but does a poor job of predicting new data (pluses). Model B (thick line), while not fitting the existing data as well as Model A, captures the main tendencies in that data and ignores the idiosyncrasies. This makes it better equipped to predict new observations, as can be seen from the deviations between the model's predictions and the new data, which are indeed smaller than the deviations for Model A.

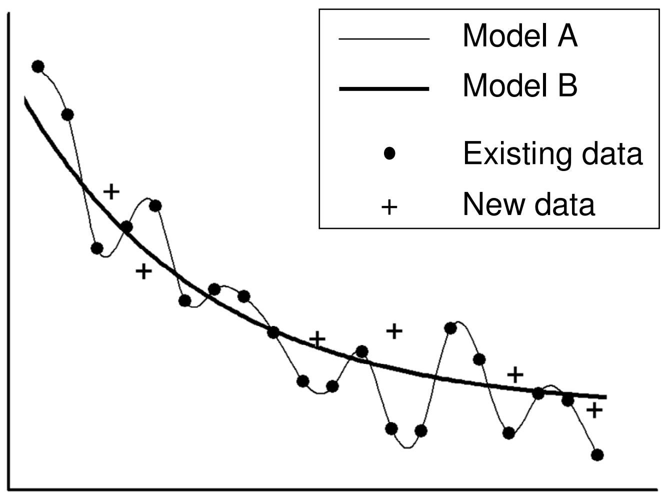

Figure 2.1. Illustration of how two models fit existing data (filled circles) and how they predict new data (pluses). Model A (thin line) overfits the data and is not as accurate in predicting new data as Model B (thick line; see Pitt et al., 2002).

The ability of a model to predict new data is called its *generalizability*, that is, the degree to which it is capable of predicting *all* potential samples generated by the same cognitive process, rather than fitting only a *particular* sample of existing data. The degree to which a model is susceptible to overfitting, in turn, is related to the model's *complexity* (Pitt et al., 2002), that is, a model's inherent flexibility that enables it to fit diverse patterns of data. Two factors contribute to a model's complexity: the number of free parameters it has, and how the parameters are combined in it—in other words, its *functional form*. The impact of many free parameters is illustrated in Figure 2.1, where the model that overfits the data (Model A) has more free parameters than the model that captures the main tendencies in the data (Model B). The impact of the functional form can be illustrated by comparing Stevens' (1957) and Fechner's (1860/1966) models of the relationship between physical dimensions (e.g., the intensity of light, called X here) and their psychological counterparts (e.g., brightness, called Y here). In both models, there are two free parameters, a and b, but they are combined differently (Stevens' model: $Y = a\,X^{b}$; Fechner's model: $Y = a \ln[X + b]$). Townsend (1975) noted that Stevens' model is more complex than Fechner's model. Since it assumes that a power function relates the psychological and physical dimensions, Stevens' model can fit data that have negative, positive, and zero curvature. Fechner's model, in turn, can only fit data with a negative curvature because it assumes a logarithmic relationship.

The relation between model complexity and generalizability can be summarized in the following way. Increased complexity makes a model more likely to end up overfitting the data while its generalizability to new data decreases. At the same time, a model's generalizability can also increase positively with the model's complexity—but only to the point at which the model is complex enough to capture systematic variations in the data. Beyond that point, additional complexity can result in decreases in generalizability, because then the model may also start to absorb random variations in the data (Pitt et al., 2002). In short, a good fit to existing data does not necessarily imply good generalizability to new data, which can make it difficult to tell which of two models provides a better explanation for data.

How to Select Between Models: A Rough Orientation

There are many different tools available for reducing the risk of selecting the wrong model. These model selection approaches can be roughly classified into *practical*, *simulation*, and *theoretical approaches*.[6]

[6] This classification is not meant as taxonomy. Rather, it is intended to aid the exposition.

Practical approaches. Practical approaches mirror the intuition that in a comparison of models, the one that can predict unseen data better than other models should be preferred. Typically, corresponding procedures estimate how well models can generalize to all possible samples generated by the same process by dividing up available data into a *calibration* (or training) set and a *validation* (or test) set. Each model's parameters are estimated on the calibration set. The resulting fixed parameters are then used to test the models on the validation set. The model that predicts the data best on the validation set (according to criteria such as mean squared error) is selected. Most common is to use some form of cross-validation (Browne, 2000; Stone, 1974, 1977). One general scheme is called *K-fold cross-validation*. Here, the data is partitioned into K subsets, and one of these K subsets is successively used as a calibration set and the remaining $K - 1$ subsets are used as the validation set. The overall prediction error is the average of the prediction error on the K validation sets. As special cases, this scheme includes *split-half cross-validation* ($K = 2$) and *leave-one-out cross-validation* (where K equals the number of observations). In Chapter 4, I will pit different models against each other using such practical approaches.

Another practical way of dealing with the problem of overfitting consists of dispensing with as many free parameters as possible—either by fixing them or by designing simple models with few or no free parameters. ACT-R, for instance, comes with an array of free parameters that can make ACT-R models flexible in fitting data. However, rather than estimating parameters each time a study is run, many researchers use the default values for these parameters (i.e., values set by the ACT-R system) or try to estimate them in separate studies (see A. Newell, 1973b), which is an approach I will take in the ACT-R modeling reported in Chapter 3. Other formal theories of cognition involve fewer free parameters. For example, several of the models that have been developed in the fast and frugal heuristics research program require no or only very few free parameters to be fitted, and in fact, in the model comparisons reported in Chapter 4, I will pit a model that does not have any free parameters against others that do have them.

Simulation approaches. When one compares models with free parameters it can be difficult to ascertain, a priori, what the models actually predict, as the predictions are dependent on the specific values of their free parameters. By simulating the predictions of competing models for a specific task, one can gain insight into the behavior of the models and use the results to design the task to maximize the possibility of discriminating between the models. It might be that the models predict the same outcome for most items (e.g., a judgment) over most of their parameter spaces, but that there are some items for which they make divergent predictions. These items could then be included in the design of the task in order to be better able to distinguish between the models.[7]

A more advanced form of the simulation approach is called *landscaping* (e.g., Navarro, Pitt, & Myung, 2004). Here the focus is on the problem of *model mimicry*, which refers to a model's ability to fit not only data generated by its own process, but also data generated by some other model. By letting a model generate many data sets, and then fitting this model as well as competing ones to that data, one can evaluate the separability of models and the informativeness of data in distinguishing between them. Moreover, this way the relative flexibility of models in fitting different data can be determined, which, in turn, can be informative for assessing the models' risk of overfitting. In Chapter 4, I will pit nested models against each other that—depending on their specific parameter values—will be able to completely mimic each other as well as result in different predictions.

Theoretical approaches. In most theoretical approaches to model selection, goodness of fit measures are combined with theoretically derived estimates of model complexity, resulting in an estimate of generalizability. Overall, such estimates can usually be expressed as generalizability = goodness of fit + complexity. Often, generalizability measures are based on the maximized log likelihood as goodness of fit index. The complexity measure, in turn, takes different forms for different generalizability measures (for an overview, see Pitt & Myung, 2000). Two widely used generalizability measures, or model selection criteria, are the *Akaike information criterion* (*AIC*; Akaike, 1973) and the *Bayesian information criterion* (*BIC*; sometimes called the *Schwarz information criterion*, see Schwarz, 1978). AIC represents the complexity of a model as the number of parameters. BIC is also sensitive to the number of parameters in the model, but in addition, it takes the log of the sample size into account. As a result, BIC favors simpler models to a greater extent than AIC does (see Forster, 2000). AIC and BIC are only sensitive to one dimension of complexity: the number of parameters. More advanced theoretical generalizability measures also take into account the functional form of a model's equation. Examples are *Bayesian model selection* (e.g., Myung & Pitt, 1997) and *minimum description length* (*MDL*; Pitt et al., 2002; see Grünwald, 2007, for a comprehensive treatment of MDL).

---

[7] However, sometimes it can be problematic to draw conclusions from tests on a set of selected items. For instance, a simple decision rule might be outperformed by another in predicting people's decisions on selected items. This, however, does not imply that the former decision rule is an invalid model of behavior. As will be demonstrated in Chapters 3 and 4 and discussed in Chapter 5, such a finding can also point to the mechanisms that underlie a person's decision to go with one decision rule rather than another; that is, such findings can point to mechanisms of strategy selection.

Choosing Between Model Selection Approaches

Choosing between different model selection approaches is not easy; they all have their pros and cons, and an approach that works in one situation might not in another. For instance, the more advanced theoretical approaches such as Bayesian model selection and MDL outperform AIC and BIC in model recovery simulations (i.e., one model is used to generate data and other models are then fitted to that data; see e.g., Myung, Balasubramanian, & Pitt, 2000). However, applying these more advanced procedures usually requires a high level of mathematical ability on the part of the researcher, making it difficult for many investigators to rely on them in practice. Sometimes the nature of a model rules out certain approaches. To illustrate, when comparing extremely complex ACT-R models, it might not be possible to derive the equations necessary for Bayesian model selection or MDL. Instead, one may often have to rely on practical approaches. Also the choice between simpler theoretical procedures is not straightforward. For example, mathematical and simulation results have fuelled a long-running debate in the model selection literature about using AIC or BIC. Many researchers side with BIC, because it identifies the correct model when the number of observations approaches infinity, and in addition, it has outperformed AIC in many simulation studies (e.g., Wagenmakers, 2003). These findings, however, have been challenged, in particular by Burnham and Anderson (e.g., 2002). They argued that most of the simulation results are not relevant to realistic model selection problems.[8]

In short, there are a range of model selection criteria that allow researchers to pit competing models against each other. While not all criteria are always applicable, in an ideal situation, all of the applicable criteria favor the same model. Sometimes, however, not all criteria will point unanimously to the same model, making it difficult to determine which model is the best.[9] Here, I cannot give general guidelines as to which method to choose. However, what I can say is that I personally tend to favor practical approaches such as cross-validation—mostly because they have a high face validity and intuitive appeal, are relatively easy to put into practice, and focus on prediction as their

[8] For instance, according to Burnham and Anderson (2002), in model recovery simulations the model that generated the data is known, and hence perfectly recoverable, favouring BIC. In real model selection problems, however, the data generating process is unknown and can at best only be approximated.

[9] There are many other complications that can arise when designing and testing models. Let me just mention two of them. First, if precision is the major virtue of modeling, it can also be a curse. Modelers need to decide how to bridge the gaps between informal verbal descriptions of theories and formal implementations, which can lead to unintended discrepancies between theories and their various formal counterparts, otherwise known as the *irrelevant specification problem* (see Lewandowsky, 1993). In Chapter 4, I somewhat face this problem, being forced to implement underspecified hypotheses in different formal models. Another problem that can arise in complex models is the *Bononi paradox*: When models become more complete and realistic, they become less understandable and more opaque (Dutton & Starbuck, 1971). To illustrate, if one has a model of how the brain works and constantly adds more intricate layers of simulated neurons to this model, then it might end up as no more understandable than the workings of an actual brain.

benchmark.

Conclusion

To conclude, I would like to highlight a point that I did not address in the previous sections: Often, there may exist a universe of different models, all of which are equally capable of reproducing and explaining the behavior—a dilemma that is also known as the *identification problem* (Anderson, 1976). As a result it appears unreasonable to ask which of many models is more "truthful"; rather, one needs to ask which model is better than another given a set of criteria, for example, the models' practical relevance, simplicity, or usability. As Box (1979) puts it—and I agree—"All models are false, but some are useful" (p. 202). Importantly, however, just as universes of functionally equivalent models may abound, there are an infinite number of vague theories for which nobody will ever be able to decide whether one is better than another. Thus, even though all models might be wrong, often there is no good alternative to building and testing them. In the following two chapters, let me do exactly this: design and test (wrong) formal models.

Chapter 3

How Memory Aids Strategy Selection

This chapter was submitted for publication to the journal Psychological Review. A revised and extended version of this chapter has received strong encouragement for resubmission to this journal. A copy the original dissertation chapter is available from Julian Marewski and can be requested at marewski[AT]mpib-berlin[DOT]mpg[DOT]de.

Marewski, J. N. & Schooler, L. J. (2010). Cognitive niches. An ecological model of emergent strategy selection.

Chapter 4

Models of Recognition-based Multi-alternative Inference

Abstract

The recognition heuristic is a simple mnemonic decision strategy for two-alternative inference. Initial experiments suggested that people employ this heuristic, basing inferences solely on recognition. Later, diverging results led researchers to conclude that recognition is integrated with other information. As this alternative hypothesis was never formalized as a testable model, I formulate corresponding models and pit them against the heuristic. Assuming that the heuristic is used by default, I also specify under what conditions people employ it and when they instead rely on other strategies. For example, people's reliance on the recognition heuristic depends on the retrieval of episodic knowledge about the source of recognition, semantic knowledge about cues in conflict with recognition, and the strength of the recognition signal. I re-formulate the heuristic for tasks with multiple alternatives, showing how it can generally aid decision making. Six experiments and 8 model comparisons suggest that the heuristic is a default strategy for multi-alternative inference.

Introduction

As of writing this chapter, eight Democratic and nine Republican candidates in the U.S. primaries have invested thousands of dollars to get their names into Americans' recognition memories. Even after years in politics, higher name recognition considerably increases their chances of being included in voters' consideration sets of candidates potentially worth a vote. At the same time, not only in the United States but all over the globe, people are wondering who will emerge from the primaries a winner, possibly becoming the next President.

In this chapter, I examine a simple cognitive strategy that can be relied on to make inferential decisions, such as forecasting which candidates voters are most likely to favor in an election. This strategy is known as the recognition heuristic (Gigerenzer & Goldstein, 1996; Goldstein & Gigerenzer, 2002). As does the fluency heuristic, which was the focus of my research in Chapter 3, the recognition heuristic operates on the accessibility of memories. It helps people infer which of two objects, one recognized and the other not, has the larger value on a given criterion. The heuristic reads as follows:

> If only one of two objects is recognized, infer the recognized one to have a larger value on the criterion.

Recognition information is not only useful for making social judgments about other people's minds, such as forecasting which candidates people will vote for. Relying on recognition is a far more general principle that can often lead us to make accurate inferences about uncertain future events or unknown quantities. Recall, as mentioned in Chapter 3, that our recognition of soccer teams and tennis players can be used to forecast their future success in competitions (e.g., Pachur & Biele, 2007; Scheibehenne & Bröder, 2007). Our recognition of universities and cities allows us to predict their quality and size, respectively (Hertwig & Todd, 2003; Reimer & Katsikopoulos, 2004), and our recognition of the names of billionaires reflects their fortunes (Hertwig et al., 2008). Yet the recognition heuristic, as formulated by Goldstein and Gigerenzer (2002) for making inferences about two objects, is of little help when evaluating three or more. The first goal of this chapter is to test whether the heuristic can be directly generalized to N objects—a generalization that may help explain how people form the consideration sets from which they make their final choice. This has been a key question in marketing, politics, and other decision-making domains.

Generalizing the Recognition Heuristic: Elimination by Recognition

In the marketing literature, many theories of choice assume a two-stage process: When evaluating multiple objects, for instance, when deciding which of eight candidates to vote for, or which of 20 cars to buy, first a smaller set of relevant alternatives is formed, and then a choice is made after more detailed examinations of the objects in this *consideration set* (e.g., Alba & Chattopadhyay, 1985; Hauser & Wernerfelt, 1990; Howard & Sheth, 1969). When recognition correlates strongly with the criteria on which objects are evaluated, the recognition heuristic generates "consideration sets" consisting of recognized objects:

> If there are N objects, then rank all n recognized objects higher on the criterion than the N–n unrecognized ones.

Once they are identified, in a second stage recognized objects can be ranked with heuristics that use objects' attributes as *cues*, say, knowledge about a candidate's party affiliation, or a car's carbon-dioxide output.

Consideration sets facilitate decisions by reducing the number of objects. To illustrate, a voter may want to forecast the final rank order of eight Democratic candidates in the primaries. But there are a total of 8! (40,320) possible rank orders. In contrast, if the recognition heuristic is used, and, say, three candidates are recognized and five unrecognized, then only 3! (6) ranks need to be considered. Unrecognized objects can be put aside (or ranked at random) because they are likely to score low on the criterion, and people typically know nothing about them in the first place. By ignoring the unheard-of, the recognition heuristic reduces complexity without necessarily harming accuracy.

Alternative Models of Inference: A Competition

In contrast to the recognition heuristic, many models of consideration-set generation posit that people evaluate objects by weighting and adding their values on a range of cues (e.g., Hauser & Wernerfelt, 1990; Roberts & Lattin, 1991). The assumption is that an object's low value on one cue can be *compensated for* by a high value on another cue.

However, there is evidence that people do not always make trade-offs (e.g., Einhorn, 1970; Fishburn, 1974; Hogarth, 1987; Payne et al., 1993; Yee, Dahan, Hauser, & Orlin, 2007). For instance, in a review of 45 process-tracing studies, Ford, Schmitt, Schechtman, Hults, and Doherty (1989) concluded that noncompensatory processes are the rule and compensatory processes are almost only observed in situations with few objects and cues. Surprisingly, by ignoring information rather than integrating it all, noncompensatory heuristics can yield more accurate judgments than compensatory ones (e.g., Gigerenzer & Goldstein, 1996) and at the same time simplify tasks (e.g., Einhorn, 1970; Simon, 1955).

The recognition heuristic is a noncompensatory model. Even when other cues are retrieved, when the heuristic is used, these cues are ignored.[29] Recently, findings that people systematically

[29] The recognition heuristic is a model of *probabilistic inference under uncertainty*, not a model of *deduction in situations of certainty*, which would be a *local mental model* (Gigerenzer et al., 1991). Consistent with this distinction, when Goldstein and Gigerenzer (2002) stressed that the recognition heuristic is a noncompensatory strategy, stating that "no other information can reverse the choice determined by recognition" (p. 82), they referred to knowledge about the values of objects on probability cues that correlate with objects' criterion values but do not directly reveal them. In fact, results reported by Oppenheimer (2003) and Pachur and Hertwig (2006) suggest that people do not use the recognition heuristic when they know that a recognized alternative has a very small criterion value, that is, when they can construct a local mental model to solve a task.

make inferences inconsistent with this heuristic have raised doubts about its adequacy as a model of behavior (Bröder & Eichler, 2006; Dougherty et al., 2008; B. R. Newell & Fernandez, 2006; B. R. Newell & Shanks, 2004; Oppenheimer, 2003; Pohl, 2006). For instance, Richter and Späth (2006) ran a series of studies and—observing that fewer decisions were consistent with the recognition heuristic when cues that contradicted recognition were available—concluded that there was no evidence of a noncompensatory use of recognition. According to them, there was clear evidence that recognition is integrated with knowledge. However, such conclusions may be premature since no alternative hypothesis—compensation by integration—was formally specified and formulated as testable model. (Moreover, reanalyses of Richter and Späth's data have shown that the large majority of participants' decisions were consistent with the prediction of the recognition heuristic even in the presence of contradictory cues; Pachur, Todd, Gigerenzer, Schooler, & Goldstein, in press.) Much the same can be said of all other previous work including my own where I have also not reported corresponding comparative model tests of the recognition heuristic (Pachur et al., 2008). In short, except for Pachur and Biele (2007),[30] no study has ever tested a single compensatory strategy, or any other model, for that matter, against the recognition heuristic, which is what is really needed to evaluate this model's ability to account for behavior. In this chapter, for the first time a competition between different recognition-based models is conducted. These models include various formalizations of the alternative hypotheses to the recognition heuristic that are discussed in the literature. All models are listed in Table 4.1 and are explained in the text below.

In carrying out a total of eight model comparisons, I will not only evaluate the recognition heuristic as a model of behavior but also formulate a theory that explains the findings reported by Richter and Späth (2006) and many others: As stressed by B. R. Newell and Fernandez (2006), such findings can be interpreted in two ways. One is that they challenge the recognition heuristic's plausibility; the other is that they point to the mechanisms that determine when people rely on the recognition heuristic and when they adopt other strategies.

[30] In their study, Pachur and Biele (2007) did not assess individual participants' knowledge about the objects. As a result, all alternative models made the same predictions for all participants. While I believe that Pachur and Biele took a laudable step in the right direction, their analyses differ substantially from the kind of comparative model tests I have in mind: I tailor all models to individual participants' information about the objects, assessing for each participant optimal weights for the information, and comparing the models' ability to fit existing data as well as to generalize to new data. Moreover, for the first time I also pit the recognition heuristic against models that operate on retrieval fluency, that is, the speed of retrieving and recognizing an alternative (see Chapter 3).

Table 4.1

List of competing models

Model input	Decision rule
Recognition heuristic (Experiments 11–16)	
1 recognized object R 1 unrecognized object U	Choose R.
n recognized objects R_i (i = 1, …, n) N - n unrecognized objects U_j (j = 1, …, N - n)	Rank R_i higher than U_j.
2 recognized objects R_a, R_b with recognition sources S_a, S_b 2 ecological correlations, 1 for S_a, 1 for S_b, each of strength e_a, e_b	If $e_a >> 0$ and $e_b = 0$, then choose R_a.
Cue-based Competitor 1: Take-one-cue (Experiments 11, 13)	
1 recognized object R with up to 1 cue with value $1 \leq v \leq 15$ 1 unrecognized object U 1 cut-off criterion C_1	If $v \leq C_1$, then choose R. If $v > C_1$, then choose U. If there is no cue, then choose R.
n recognized objects R_i (i = 1, …, n), each with up to 1 cue with value $1 \leq v_i \leq 24$ N - n unrecognized objects U_j (j = 1, …, N - n) 1 cut-off criterion C_2	If $v_i \leq C_2$, then rank R_i higher than U_j. If $v_i > C_2$, then rank U_j higher than R_i. If there is no cue, then rank R_i higher than U_j.
Cue-based Competitor 2: Tallying-of-negative-cues (Experiment 14)	
1 recognized object R with up to h cues with negative values $v_h = -1$ 1 unrecognized object U 1 cut-off criterion C_3	If $h \geq C_3$, then choose R. If $h < C_3$, then choose U. If there is no cue, then choose R.
Cue-based Competitor 3: Tallying-of-positive-and-negative-cues (Experiment 14)	
1 recognized object R with sum $\sum$ of up to h cues with negative values $v_h = -1$ and p cues with positive values $v_p = +1$ 1 unrecognized object U 1 cut-off criterion C_4	If $\sum \geq C_4$, then choose R. If $\sum < C_4$, then choose U. If there is no cue, then choose R.
Cue-based Competitor 4: Weighted-best-cues (Experiment 15)	
n recognized objects R_i (i = 1, …, n), each with up to 1 cue drawn from a set of i cues with values $-100 \leq v_i \leq +100$ N - n unrecognized objects U_j (j = 1, …, N - n) 1 cut-off criterion C_5	If $v_i \geq C_5$, then rank R_i higher than U_j. If $v_i < C_5$, then rank U_j higher than R_i. If there is no cue, then rank R_i higher than U_j.
Fluency-based Competitor 5: Fluency heuristic (Experiment 12)	
1 recognized object R_a with retrieval time r_a 1 recognized object R_b with retrieval time r_b	If $r_a < r_b$, then choose R_a. If $r_a > r_b$, then choose R_b.
Fluency-based Competitor 6: Weighted-fluency (Experiments 14, 16)	
1 recognized object R with retrieval time r 1 unrecognized object U 1 cut-off criterion C_6	If $r \leq C_6$, then choose R. If $r > C_6$, then choose U.

Note. All competitors except the fluency heuristic can also be thought of as weighted linear additive models with two classes of predictors, cues and recognition, or retrieval time and recognition, respectively. Different cut-off criteria, $C_{1–6}$, (free parameters) measure the weight of these predictors relative to each other. The recognition heuristic can be formally represented as a special case of these models with $C_{1–6}$ set such that the models always choose recognized objects, which is equivalent to assuming a noncompensatory recognition weight.

Toward a Theory of Strategy Selection by Default

No cognitive strategy is always relied upon. Rather, in keeping with many other frameworks (e.g., Beach & Mitchell, 1978; Busemeyer & Myung, 1992; Payne et al., 1993), I assume that the mind comes equipped with a repertoire of strategies. As explained in Chapters 1 and 3, this repertoire forms an "adaptive toolbox" of heuristics, each of which exploits how basic cognitive capacities, such as memory, represent regularities in the structure of our environment (e.g., Gigerenzer et al., 1999). In doing so, heuristics can yield accurate judgments by operating on little information, say, a sense of recognition. Often, the study of the mechanisms determining the use of a heuristic can be informed by an analysis of the heuristic's ecological rationality, that is, of the environmental structures it exploits.

Figure 4.1 illustrates the ecological rationality of the recognition heuristic in terms of three correlations. It closely resembles the ecological rationality of the fluency heuristic which I discussed in Chapter 3 (see Figures 3.1, 3.7).

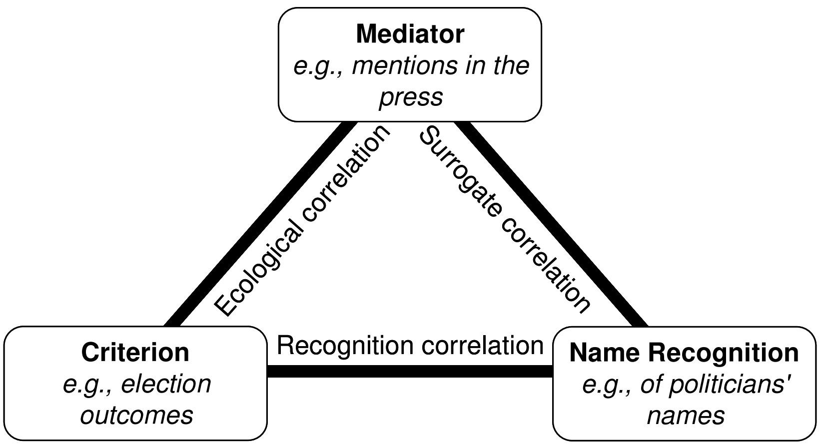

Figure 4.1. Ecological rationality of the recognition heuristic. An unknown criterion (e.g., the number of votes candidates win in an election) is reflected by a mediator (e.g., the press). The mediator makes it more likely for a person (e.g., a voter) to encounter objects with larger criterion values than those with smaller ones (e.g., the press mentions more successful candidates more frequently). As a result, the person will be more likely to recognize objects with larger criterion values than those with smaller ones, and, ultimately, recognition judgments can be relied upon to infer the criterion (e.g., the success of candidates in elections). The relations between the criterion, the mediator, and recognition can be measured in terms of correlations, or, as I have done for the fluency heuristic in Chapter 3, in terms of validities (see Figure 3.7).

There is a criterion, an environmental mediator, and a person who infers the criterion. Using the recognition heuristic is ecologically rational when there is both a substantial *ecological correlation* between the mediator and the criterion and a substantial *surrogate correlation* between the mediator and recognition. This combination can yield a substantial *recognition correlation*; that is, recognized objects tend to have higher criterion values than unrecognized ones. If either or both the ecological and surrogate correlations are zero, the use of the recognition heuristic is not ecologically rational.

How Can the Mind Judge the Ecological Rationality of Using the Recognition Heuristic?

In Chapter 3, I proposed a theory of strategy selection, arguing that the interplay between the workings of memory, the environment, and people's decision strategies (i) constrains the choice set of applicable heuristics, giving rise to what I called non-overlapping cognitive niches of heuristics. As I have shown, this interplay can (ii) also make it likely that the application of a strategy can require less effort or time when using it is also likely to result in accurate decisions. Here, I will focus less on the applicability of the recognition heuristic as a determinant of strategy selection and more on its accuracy. In the General Discussion, I will discuss in more detail how the cognitive niche of the recognition heuristic differs from that of other heuristics, constraining the choice between them.

In what follows, I argue (a) that the recognition heuristic is applied "by default," but (b) this default can be overruled by information indicating that it is not ecologically rational to use the heuristic because recognition is not predictive of the criterion. There are two kinds of evidence, behavioral and neural, indicating that the heuristic is used as a default, as opposed to being just another decision strategy. Pachur and Hertwig (2006) and Volz et al. (2006) reported response time data suggesting that recognition of an object often arises before further knowledge can be retrieved from memory. This finding is consistent with the recognition memory literature, indicating that a sense of recognition (often called "familiarity"; see Chapter 3) arrives on the mental stage earlier than recollection (Gronlund & Ratcliff, 1989; Hintzman & Curran, 1994; McElree et al., 1999; Ratcliff & McKoon, 1989). A functional magnetic resonance imaging (fMRI) study (Volz et al., 2006) showed that judgments in disagreement with the recognition heuristic required more cognitive effort (indicated by a reduction in activation in the anterior frontomedian cortex) than judgments in line with it. This study provides evidence for two processes: recognition and evaluation. The first identifies an object as recognized or not, and the second evaluates whether a default to rely on the heuristic should be overruled because recognition is not predictive of the criterion. The literature offers different hypotheses suggesting that people may rely on the contents of *semantic memory* (general knowledge about the world), *episodic memory* (personal experiences occurring at a particular place and time; Tulving, 1983), and the recognition signal itself to evaluate whether recognition is predictive of the criterion.[31]

Thesis 1: Conflicting cues. People may use semantic knowledge about objects' values on cues to evaluate the predictive accuracy of recognition. For example, a voter in Germany may recognize a candidate's name but know that he is running for one of the small parties, indicating that recognition may not be predictive of his electoral success. Thus, the voter may not rely on recognition in election forecasts for this candidate. In fact, some people appear not to rely on recognition when retrieving such cues at odds with recognition (B. R. Newell & Fernandez, 2006; Pachur et al., 2008; Richter & Späth, 2006).

Thesis 2: Recognition correlation. There is some evidence that people evaluate the predictive accuracy of recognition, using semantic knowledge to judge whether the recognition correlation (Figure 4.1) is substantial or not. When this correlation is substantial, people make inferences in accordance with the recognition heuristic (e.g., Hertwig et al., 2008; Pachur & Biele, 2007; Reimer & Katsikopoulos, 2004; Scheibehenne & Bröder, 2007; Serwe & Frings, 2006; Snook & Cullen, 2006; Volz et al., 2006). In contrast, when it is less pronounced, they tend not to do so. Pohl (2006) asked people to infer which of two cities is situated farther away from the Swiss city of Interlaken, and which of two cities is larger. Most people probably knew that their recognition of cities is not indicative of the cities' distance to Interlaken but is indicative of their size. Indeed, for the very same cities, the heuristic predicted only 54% of people's inferences when inferring spatial distance (which is no better than chance), but 89% when inferring size. Similarly, people who *always* judge recognized objects to be larger than unrecognized ones—compared to people who sometimes infer unrecognized objects to be larger—estimate the recognition correlation to be larger (Pachur et al., 2008).

Thesis 3: Episodic knowledge. In addition to semantic knowledge, episodic knowledge about the source of recognition may be used. The recognition heuristic only yields accurate inferences if recognition sources are environmental mediators that make it likely for encountered objects to have large criterion values (Figure 4.1). If a person believes an object is recognized from sources that are not linked by a substantial ecological correlation to the criterion, she might judge her recognition as not predictive of the criterion. For instance, when a voter has to infer which of two politicians is more famous and recognizes one only because he is her neighbor, she may regard

[31] The distinction between semantic and episodic knowledge is made for conceptual reasons here. In ACT-R, both could be modeled as chunks in declarative memory (see also Anderson et al., 1998).

this recognition source as unrelated to the politician's fame. She would therefore evaluate her recognition of this particular name as not being predictive. There is evidence for such source evaluations in the recognition memory literature (e.g., Jacoby, Kelley, et al., 1989; Jacoby, Woloshyn, & Kelly, 1989).

Thesis 4: Strength of the recognition signal. The recognition heuristic operates on a sense of recognition. People's reliance on it may depend on the strength of the underlying recognition signal, as captured by an object's retrieval time in Equation 4 in Chapter 3. For instance, it may take a voter much time to judge a politician as recognized, such that she lacks confidence about her recognition judgment. Thus, she may not rely on the recognition heuristic when making inferences about that politician. This can be ecologically rational: An object's recognition time can reflect the predictive accuracy of recognition for making inferences about that object (see Chapter 3). In six experiments, I tested these and related hypotheses—all of which will be further elaborated below.

Overview of the Experiments: Recognition in Political Elections

In Experiment 11, I tested whether semantic knowledge about conflicting cues leads people to overrule a possible default of using the recognition heuristic. In Experiment 12, I investigated how people use episodic knowledge to rely on the heuristic. In Experiments 13 and 14, I focused on the strengths of the recognition correlation and the recognition signal as possible determinants of strategy selection. In Experiment 15, I reconsidered cues. In Experiment 16, I returned to the strength of the recognition signal as a possible determinant of strategy selection and tested how well the heuristic predicts judgments when there is little time to execute the evaluative process that appears to be involved in overruling a default use of it. In all studies, I pitted the heuristic against other models for two-alternative (Experiments 11, 12, 14, 16) and multi-alternative (Experiments 13, 15) decisions, focusing on individual differences in strategy use in Experiments 12, 14–16.

As discussed in Chapter 3, since the publication of Woodworth's (1938) book *Experimental Psychology*, a frequent experimental practice has entailed manipulating a few variables while keeping all others constant or varying them at random. This practice can lead to the use of highly artificial stimuli in laboratory tasks. However, sometimes such tasks destroy the natural covariation of variables, making it difficult to generalize from them to a world where people exploit this covariation (Brunswik, 1955; for reviews see Dhami et al., 2004; Hoffrage & Hertwig, 2006), as is the case when they rely on recognition. Therefore, I *also* studied the recognition heuristic in more naturalistic situations where multiple objects are common: political elections.

Do Conflicting Cues Overrule the Reliance on Recognition?

(Experiment 11; Reanalysis of Earlier Election Study)

The first data set I will consider here stems from an older study that I ran in the 2004 parliamentary elections in the German federal state of Brandenburg (Marewski, 2005). Here, 2,117,145 eligible voters had the opportunity to cast their ballots. As in many other elections around the globe, several weeks before voting day, election ads were placed, and the elections started to figure prominently in the news. The election system works like this: As in most German states, every 4 years, each citizen has two votes, one for a direct candidate who will represent the person's voting district and a second for a party, representing a list of candidates. Direct candidates are typically affiliated with one of the parties competing in the election. They are voted into Parliament if they win the most votes in their voting district. If a party is voted into Parliament, then, depending on the proportion of votes the party gained, a number of the candidates from its list enter Parliament.

One question I originally asked when running this first study was whether it is actually ecologically rational to use recognition to forecast the outcomes of political elections. For this to be the case, the following needs to be true: First, the frequency of mentions of parties and candidates in the press and election advertisements *before* the election should reflect the number of votes they win. Second, these frequencies of mentions should correlate with the number of voters who have heard of the parties and candidates *before* the election. Third, it should be possible to predict election outcomes based on voters' recognition of parties and candidates alone (see Figure 4.1). To examine whether these relations hold, on two dates, 14 days and 1 day before the Brandenburg election, I invited eligible voters to fill out a questionnaire, asking them which candidates and parties they recognized. This way, I could count for each candidate and party how many voters had heard of or seen its name. In addition, during a period of 68 days to 14 days before the election I counted the number of articles in which a candidate's name appeared in the *Potsdamer Neueste Nachrichten* (*PNN*) and the *Märkische Allgemeine Zeitung* (*MAZ*). These daily newspapers focused on the federal state of Brandenburg and were the most frequently read papers in my sample of eligible voters. For parties, I computed corresponding counts in the *MAZ*. Candidates also provided me with counts of the number of campaign materials (e.g., election posters, brochures, letters; henceforth: *flyers*) in which their name had appeared and that were distributed (prior to the election) in the voting districts where the eligible voters from my sample lived.

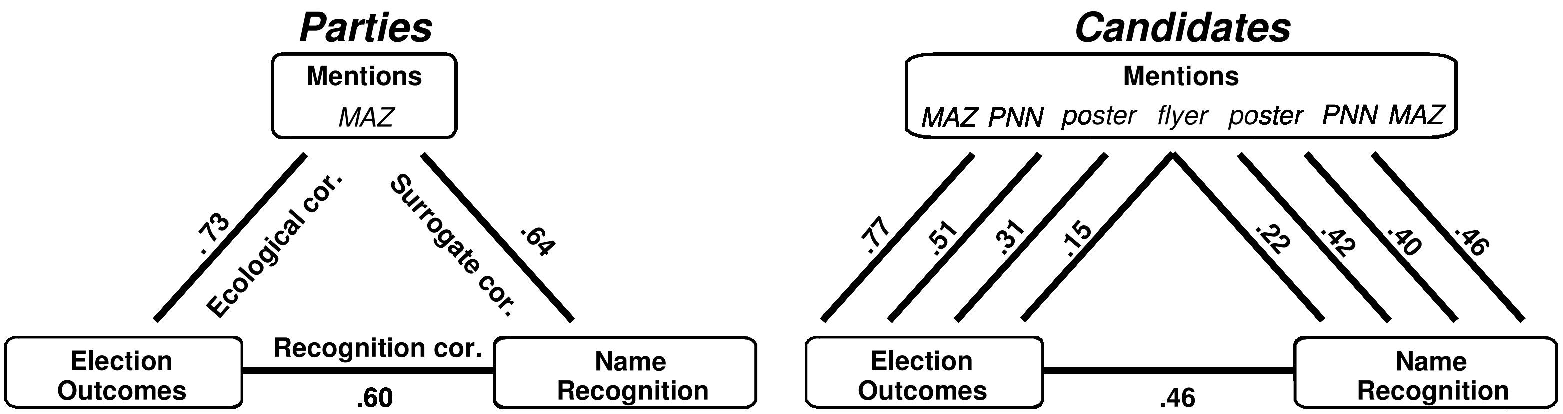

Figure 4.2. Parties: Goodman and Kruskal's (1954) gamma computed between the number of mentions of 15 parties in the newspaper *Märkische Allgemeine Zeitung* (*MAZ*), the number of votes won by 15 parties in the election, and the number of participants who recognized the name of a party. *Candidates*: Gamma computed between the number of mentions of 11 candidates in the *MAZ* and in the *Potsdamer Neueste Nachrichten* (*PNN*), the number of election posters with a candidate's name, and the number of flyers with a candidate's name, and percentages of votes won by each candidate in his or her voting district, as well as the number of participants who recognized a candidate's name. Interestingly, newspaper mentions are the best predictors of election outcomes and name recognition (cor: correlation; $N = 172$; Experiment 11).

As Figure 4.2 shows, there were in fact substantial correlations between the frequency of mentions of parties and candidates in the newspapers, the number of eligible voters who recognized them, and the election results. Importantly, my older analyses not only showed that it is ecologically rational to use the recognition heuristic to forecast electoral success, but they also provided some evidence to suggest that eligible voters use the recognition heuristic to make such forecasts. Yet, as for all other previous studies on the recognition heuristic, this one did not pit alternative models against each other. (In this study, I only counted how many of voters' election forecasts were consistent with the prediction of the recognition heuristic, ignoring competing explanations for voters' forecasts.) Moreover, even though the data then collected allowed me to examine if the presence of semantic knowledge about another cue is associated with a suspension of the recognition heuristic's use, corresponding analyses were not run at that time. In what follows, I will reanalyze data from this study, filling these two gaps. To simplify references to this reanalysis later on in the discussion sections of this dissertation, I will call it Experiment 11 throughout.

Short Summary of the Methods Employed in Experiment 11

To vote for a direct candidate from a given voting district, a citizen has to be a resident of this district. On two dates, 14 days and 1 day before the election, I invited residents of two voting districts in the downtown areas of the cities of Potsdam and Werder to fill out a questionnaire. They were paid €5 ($7). Participants had to be at least 18 years old (voting age in Germany). Of 246 recruited eligible voters (henceforth: *voters*), 172 completed the questionnaire (70%; 55% female; mean age 38 years, SD = 14.7).

The questionnaire consisted of different sections. The ones that are relevant here included a candidate recognition task and a two-alternative forced-choice task on the 11 direct candidates running in the two voting districts, and a cue knowledge task. In the recognition tasks, voters were given lists of candidates' and parties' names, respectively. For each name, they indicated whether they recognized it, that is, whether they had heard or seen it before participating in the experiment. The two-alternative task consisted of a list of 25 paired comparisons of candidates' names (complete pairings of the candidates from each voting district, respectively). Voters indicated for each pair which candidate would win more votes in the election. In the cue knowledge task, voters indicated candidates' party affiliation, which is commonly a highly valid cue for electoral success (not unique to Germany—for the U.S., see Bartels, 2000). If an affiliation was unknown, they were instructed not to guess but to leave the answer space blank. The recognition and two-alternative choice task were counterbalanced.[32] The order of appearance of names was randomized in all tasks. Filling out the questionnaire took about 10 min.

Description of Measures

To assess how well the heuristic predicts voters' forecasts in the two-alternative task, for each voter I selected the paired comparisons where one candidate was recognized but not the other. Across these comparisons, I counted how often (A) the voter had inferred a recognized candidate would win more votes than an unrecognized one and the number of times (D) the opposite was inferred. The accordance rate, k, is the proportion of inferences consistent with the heuristic:

$$k = A/(A + D). \tag{12}$$

Results and Discussion of Reanalyses

Do conflicting cues overrule the use of the recognition heuristic? A possible default use of the recognition heuristic could be overruled when semantic knowledge about cues at odds with recognition is retrieved from memory. Knowledge that a recognized candidate is running for a small party is such a cue. Party affiliations are commonly known to be highly predictive of electoral success, and according to all major opinion polls published prior to the election, the three largest parties were expected to win the most votes. In short, a candidate's party was possibly the best cue voters could rely on to overrule the default. Did they?

I had asked voters to indicate candidates' party affiliations. This allows me to examine how well the recognition heuristic predicts voters' forecasts in two situations. First, in comparisons between a recognized and an unrecognized candidate, a voter might believe that the recognized candidate is from one of the three largest parties. Thus, recognition and the party cue would make the same prediction (*nonconflicting pairs*). Second, the voter might believe the recognized candidate to be from one of the smaller parties. In this situation, the party cue would suggest that

[32] In other experiments on the recognition heuristic, order effects were not found (e.g., Goldstein & Gigerenzer, 2002; Pachur & Hertwig, 2006; see also the other experiments reported here in Chapter 4). In Marewski's (2005) study, voters (n = 87) who worked on the recognition task before the two-alternative forced-choice task recognized fewer candidates, M = 4.8, n = 87, than voters who worked on the tasks in the reverse order, M = 5.9, n = 85; 95% CI on the mean difference (.5, 1.9). However, there were no reliable differences with regard to the accordance rate, which is the variable that matters in the context of this re-analysis.

the recognized candidate would win fewer votes while recognition would suggest that this candidate would win more (*conflicting pairs*).

For the 81 voters for whom I identified at least one conflicting *and* one nonconflicting pair, the recognition heuristic predicted forecasts better on nonconflicting pairs, M_k = .87, *SE* = .03, than on conflicting pairs, M_k = .73, *SE* = .04; 95% CI on the mean difference (.07, .22). However, most voters *always* behaved in accordance with the heuristic on both types of pairs (median accordance rate in both cases = 1.00; Figure 4.3). In short, a strong contradictory cue had no impact on the majority's reliance on the recognition heuristic, but it did seem to cause a minority to overrule a possible default of using it.

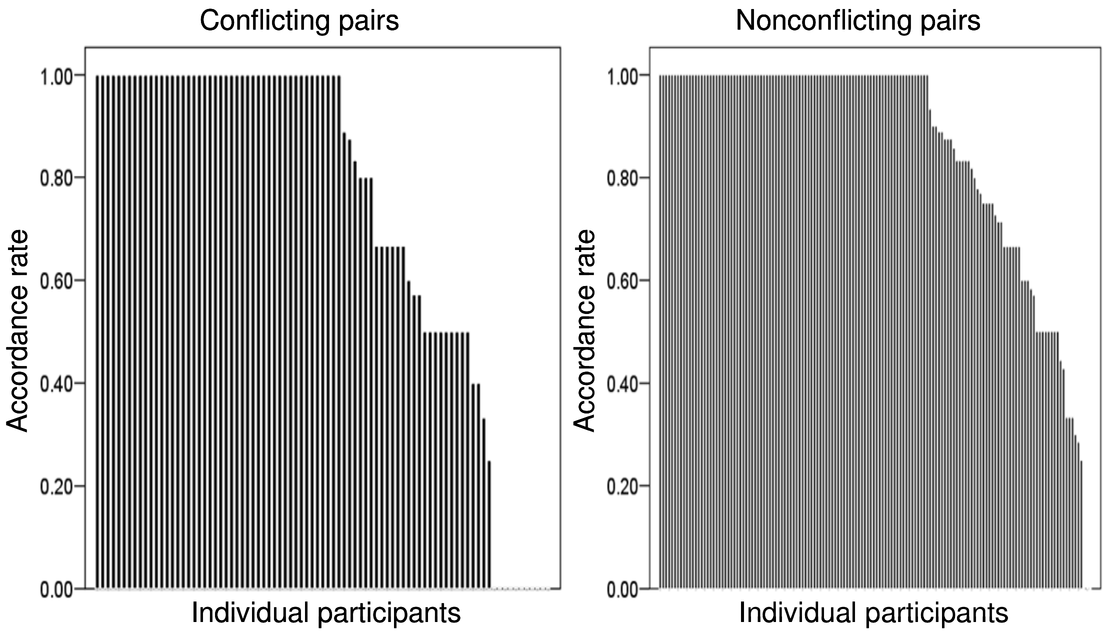

Figure 4.3. Recognition heuristic accordance rates for conflicting and nonconflicting pairs. Bars represent individual participants' accordance rates. On nonconflicting pairs, the accordance rate for 3 participants was 0; on conflicting pairs the accordance rate of 11 participants was 0. For these participants no bars are shown; dots on the figure border indicate their data ($n_{\text{nonconflicting}}$ = 146; $n_{\text{conflicting}}$ = 85; Experiment 11).

Model comparison 1. To evaluate how well the heuristic predicts behavior, I compared it to a compensatory model that integrates the party cue and recognition: According to *take-one-cue* (Table 4.1), if a voter believed a recognized candidate was affiliated to one of the smaller parties, he would infer that this candidate would win fewer votes than an unrecognized one. Conversely, if a recognized candidate was from a large party, the voter would infer that the recognized candidate would win more votes than the unrecognized one. Last, if a recognized candidate's affiliation was unknown, the voter would infer that this candidate would win more votes.

Take-one-cue is more flexible than the recognition heuristic: It can also infer larger criterion values for unrecognized objects. The model pays for this flexibility with an increase in complexity (see Chapter 2). That is, it assumes a free parameter, the *cut-off criterion*, C_1, which measures the weight of one highly predictive cue—party size—against the weight of recognition. I operationalized party size in terms of the parties' success in the Brandenburg election and classified for each participant the parties that he or she believed were a candidate's affiliation as "large" or "small" according to the number of votes the party gained. That is, to cover all possible classifications (i.e., all possible values of the parameter C_1), in a first round for $C_1 = 1$, I only classified the party that actually won the election as large; all other parties were small. In a second round, I then classified the two parties that won the most votes as large ($C_1 = 2$), and so on, until finally I classified all parties as large ($C_1 = 15$). For each participant, I computed the take-one-cue accordance rate across all pairs of a recognized and an unrecognized candidate for all values for C_1 using Equation 12 in the same way as I did for the recognition heuristic.

Formally, the recognition heuristic is a special case of take-one-cue with C_1 set to be noncompensatory, such that recognized candidates are always forecasted to win over unrecognized ones. That is, the recognition heuristic's election forecasts only differ from take-one-cue's forecasts if take-one-cue's free parameter C_1 is set to be compensatory. Does this increase in model complexity pay off?

It does not: As Figure 4.4 shows, only when take-one-cue forecasts recognized candidates to win over unrecognized ones in 99 to 100% of the pairs (from $C_1 = 6$ [99%] to $C_1 = 15$ [100%]) does it fit voters' forecasts, on average, as well as the recognition heuristic predicts them from scratch without assuming an additional parameter. From $C_1 = 1$ to $C_1 = 5$, take-one-cue makes the same predictions as the recognition heuristic in 47 to 91% of the pairs. Here, the recognition heuristic predicts voters' forecasts better. In short, this more complex, compensatory model does not outperform the recognition heuristic.

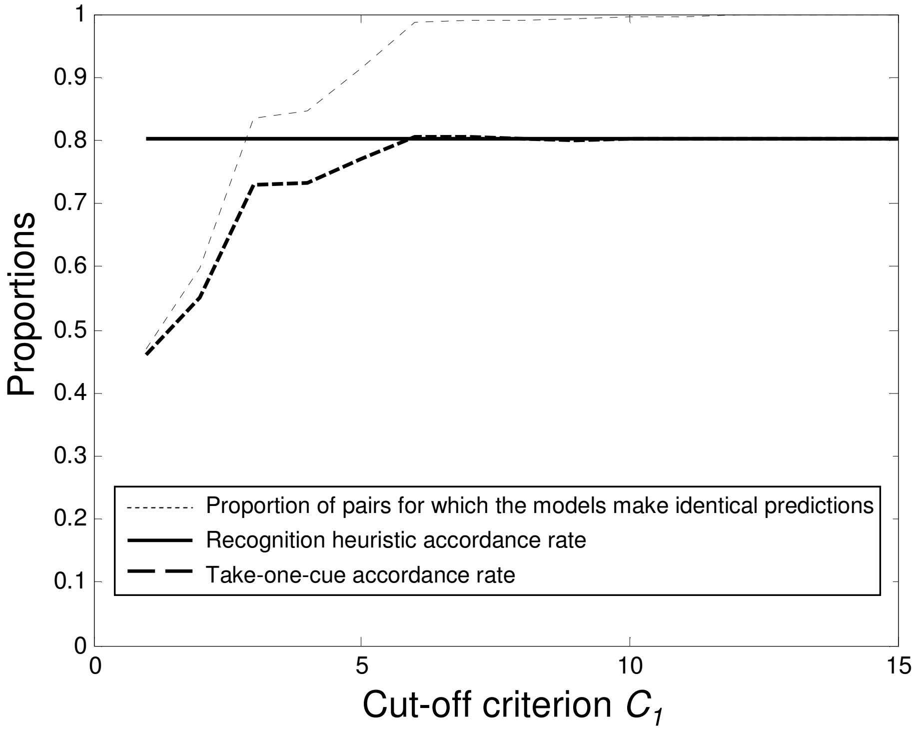

Figure 4.4. Recognition heuristic versus take-one-cue in a two-alternative forced-choice task. Lines show mean accordance rates of the two models computed for each of the possible values of the cut-off criterion C_1, as well as the mean proportion of paired comparisons in which the two models made the same predictions ($N = 164$; Experiment 11).

To summarize, when deciding which of two candidates will gain more votes, most people made forecasts that were consistent with the recognition heuristic even when a highly valid, conflicting cue was available, and in fact, a more complex compensatory model did not fit peoples' forecasts better than the heuristic predicted them.

Generalizing the Recognition Heuristic: How Episodic Knowledge Aids Using Recognition (Experiment 12)

As we sample objects in our world, recognition memory grows. Sometimes, people end up recognizing all objects of one kind, say, all parties in an election. In this case, recognition alone does not differentiate between objects, and in principle, the recognition heuristic as formulated by Goldstein and Gigerenzer (2002) cannot be used.

However, episodic memory could support the use of the recognition heuristic even then: I suspect that people treat recognized objects as unrecognized when they believe the recognition source does not represent an environmental mediator that reflects objects' criterion values (Figure 4.1)—a situation in which one has good reason to judge the predictive accuracy of recognition to be low. To illustrate, guess about whom more biographies have been written, Mr. Kissinger or Mr. Marewski. Many of you recognize the name of the graduate student who is the author of this dissertation only from reading this dissertation. Although both names are recognized, in conjunction with semantic knowledge, episodic knowledge tells us that in one case, the source is a dissertation that is not linked to the criterion, whereas in the other case, the source is the media, which may well reflect the criterion. Although recognition alone does not differentiate between the two, you could treat Mr. Marewski as unrecognized and then use the recognition heuristic to pick Mr. Kissinger. Thus, one can generalize the recognition heuristic to situations with two recognized objects (Table 4.1):[33]

> If two objects are recognized, but the source of one is unrelated to the criterion, then infer that the other has the higher value on the criterion.

Do people follow this principle based on episodic knowledge? My next experiment took place during the 2005 German national election. Here, 61,870,711 eligible voters could choose between 25 parties. In a first session prior to this election, I let eligible voters acquire recognition on parties by repeatedly exposing them to the party names in an experiment. In a second session, I could then test (a) if they treated recognized parties as unrecognized when they identified their study participation as the recognition source, a source they should have little reason to expect to reflect electoral success. I could also test (b) if they trusted recognition when they acquired it before the experiment in their natural environment, where recognition sources such as the press do reflect electoral success (Figure 4.2).

---

[33] I would not object to calling this recognition heuristic generalization a knowledge-based strategy, because in order to be applied, this generalization requires episodic knowledge about the recognition source as well as semantic knowledge about the strengths of ecological correlations to be available. However, in my view the term "recognition heuristic" adequately captures what this strategy actually does once knowledge has been used to treat one of two objects as unrecognized. That is, the strategy applies the decision rule of the recognition heuristic, inferring a recognized object to be larger than an object that is treated as unrecognized.

Method

Sixty-six residents of Berlin, Germany (52% female; mean age 26 years, *SD* = 3.7) completed a computerized experiment in the laboratories of the Max Plank Institute for Human Development. These eligible voters were at least 18 years old. They were paid €25 ($37).[34]

The first session (T1) took place about 16 days and the second session (T2) 2 days before the election (Figure 4.5). In a recognition task, in both sessions, I presented voters the name of one party at a time on a computer screen and asked them to judge whether they recognized the party, that is, whether they had seen or heard of the party prior to participating in the experiment. Participants were instructed to respond as quickly and accurately as possible. In the second session, in a recognition source task, I again presented one party name at a time and additionally asked the voters whether they had seen or heard a party name (a) before the first session, (b) during the first session, or (c) after the first session and before the second session. In both sessions, in a two-alternative forced-choice task, I presented two party names on the computer screen (one on the left side and the other on the right) and asked voters to infer which party would win more votes in the election. There were a total of 300 comparisons of parties (complete pairings of 25 parties). In all tasks, order of appearance of party names was randomized and all trials were preceded by fixation crosses for 1,000 ms. All responses were made on a standard PC keyboard. Positive responses were made with the index finger of the right hand. In each session, completing the experiment took about 50 min.

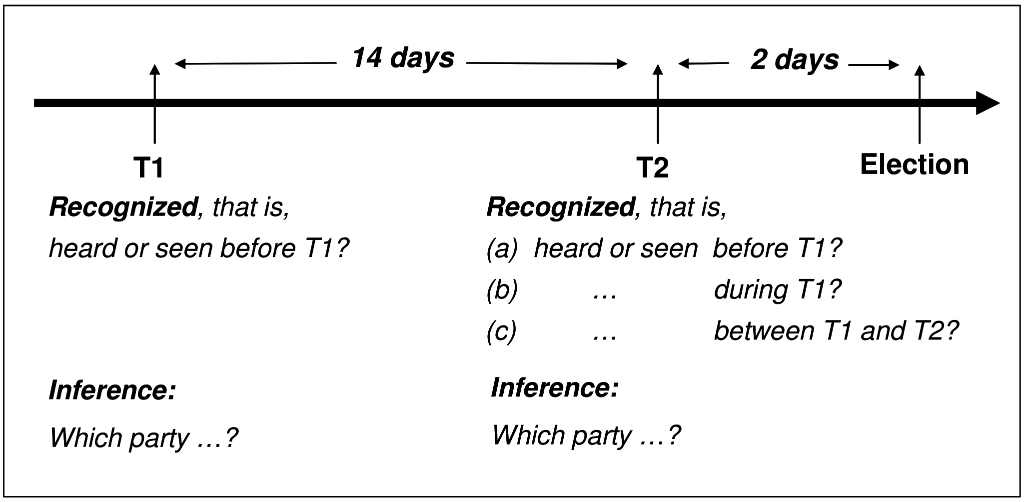

Figure 4.5. Schematic presentation of the design of Experiment 12.

[34] One participant did not return to the second session. Her data were excluded.

Results and Discussion

Do voters rely on episodic knowledge to treat parties as unrecognized? At T1, voters could apply the recognition heuristic on average on 144 pairs (SE = 2.6) consisting of one recognized and one unrecognized party. Since they had seen all parties in the experiment at T1, by T2 they could only use the recognition heuristic if they treated the parties they did not recognize at T1 as also unrecognized at T2. In this case, they could continue to forecast more votes for the parties they already recognized at T1; that is, it should be possible to predict their forecasts at T2 based on their recognition of a party name at T1.

To test this, I computed voters' accordance rates at T1 and *predicted* their accordance rates at T2 *based on recognition data from T1*. That is, I used voters' recognition data from T1 to calculate their accordance rates at T1 *and* T2, which should be similar if voters continued to forecast more votes for the party they already recognized at T1. In fact, I observed only small declines in the accordance rate from $M = .89$ at T1 to $M = .87$ at T2 (both $SE = .01$), and consistent with this result, only small drops in the accuracy of their forecasts, ($M_{\mathrm{T1}} = .84$, $M_{\mathrm{T2}} = .82$, both $SE = .01$), which should also be very similar if voters continued to forecast more votes for the previously recognized party. (The accuracy of a voter's election forecast is the proportion of correct election forecasts that voter made. It is computed separately for T1 and T2 across the comparisons of two parties the voter faced in the two-alternative task.)

Yet, in forecast for the larger parties, voters may have had an alternative to relying on episodic knowledge at T2: They could also have used semantic knowledge about cues, say, by recalling opinion poll data. To gain control over this possibility, I excluded the six largest German parties (CDU, CSU, SPD, FDP, GRÜNE, Die Linke) and, in addition, two well-known extreme right-wing parties (NPD, REP) from my analyses. These eight parties (henceforth: *larger parties*) receive news coverage even when no election is taking place and/or are regularly represented in German parliaments. Furthermore, opinion polls tend to be published only for these eight parties and not for those that are lesser known.[35] If voters relied on cues for these eight parties instead of on the recognition heuristic, for the smaller parties they could base their forecasts on random guesses. In this case, at T2 the accordance rate and accuracy of their forecast should drop to the chance level of .50 when the eight parties are excluded. However, the accordance rate did not

[35] In fact, in Experiment 13 these eight parties—except the CSU, which did not run in the election used in Experiment 13—were those parties about which, on average, 95.5% (*Mdn* = 98.3%) of the participants indicated they had knowledge. About the remaining smaller parties, only 13.4% of the participants had knowledge, on average (*Mdn* = 11.7%).

decrease to .50. It only dropped from M = .78 at T1 to M = .75 (both SE = .02) at T2, and the accuracy dropped from M = .64 (SE = .01) at T1 to M = .61 (SE = .02) at T2, suggesting that voters relied on episodic knowledge to treat some parties as unrecognized in order to use the recognition heuristic.

Model comparison 2. Yet, instead of using the recognition heuristic, voters could also rely on the fluency heuristic. As explained in Chapter 3, this heuristic has been defined in different ways (e.g., Jacoby & Brooks, 1984; Whittlesea, 1993). Here I use the term to refer to Schooler and Hertwig's (2005) model, which builds on these earlier definitions and on a long research tradition on fluency (e.g., Jacoby & Dallas, 1981) as well as related notions such as accessibility (e.g., Bruner, 1957; see Chapter 3 for details). Recall, Schooler and Hertwig implemented their fluency heuristic side by side with the recognition heuristic in the ACT-R cognitive architecture (Anderson et al., 2004; Anderson & Lebiere, 1998). In ACT-R, the same memory currency—a continuous activation trace—determines (a) whether an object will be retrieved and if so, (b) the time it takes to retrieve it. Schooler and Hertwig adopted Anderson et al.'s (1998) assumption that an object's retrieval implies recognizing it, adding the assumption that the more quickly the object is retrieved, the greater the sense of recognition. A person using the recognition heuristic can base inferences on the binary outcome of this memory process (retrieved or not). A person using the fluency heuristic, in turn, can base inferences on the more graded outcome of the same process, namely, on the speed with which the objects come to mind, that is, on their retrieval time or retrieval fluency, which I take to reflect the strength of the underlying recognition signal. (In Chapter 3, I modeled this retrieval time with Equation 4.) As explained in Chapter 3, by this token, the fluency heuristic is a computational instantiation of the version of Tversky and Kahneman's (1973) availability heuristic that bases judgments on ease of retrieval. (For a discussion of the similarities between the two heuristics and the recognition heuristic, see Schooler & Hertwig; Hertwig et al., 2008; for a discussion of different versions of availability, see Hertwig et al., 2005; Sedlmeier et al., 1998; see Jacoby & Dallas, 1981, p. 333; Jacoby, Kelley, et al., 1989, p. 328, for an articulation of the link between their fluency/familiarity concept and availability.) The fluency heuristic can be stated as follows (Table 4.1):

> If one object is more quickly retrieved than the other, infer that this object has the higher value with respect to the criterion.

Operationalizing retrieval time as recognition time, I used each voter's reaction times in the recognition task at T2 to compute the proportion of judgments consistent with the fluency heuristic at T2 (Equation 12). Recall, in Chapter 3, I successfully applied this procedure to model people's inferences with the fluency heuristic. As in the analyses reported previously in this Results section, I used the recognition data from T1 to compute the recognition heuristic accordance rate on the same pairs at T2. To gain control over voters' access to cues, I excluded the eight larger parties.

The fluency heuristic cannot be used when differences in retrieval time between two objects are very small, and therefore not detectable. Based on Fraisse's (1984) review of the timing literature, Schooler and Hertwig (2005) assumed that people are able to detect differences in time that exceed 100 ms. In Chapter 3, I provided empirical support for this assumption (see also Hertwig et al., 2008). Thus, I restricted the comparison between the two heuristics to parties that differed by more than 100 ms in recognition time at T2. On the resulting pairs, each heuristic unambiguously decides for one of the parties. Which heuristic predicts voters' decisions best?

For 49 voters, the recognition heuristic accounts for more election forecasts than the fluency heuristic. For 12, the fluency heuristic predicts judgments best, and for 4 voters the two heuristics' accordance rates are identical (Figure 4.6). This suggests that a majority of voters relied on episodic knowledge to treat recognized parties as unrecognized and applied the recognition heuristic. This finding is consistent with my experimental results and computer simulations with ACT-R's memory model reported in Chapter 3, suggesting that people are most likely to rely on the fluency heuristic when they cannot use knowledge instead. It is also consistent with studies indicating that people are less likely to rely on a sense of fluency when this sense has been manipulated experimentally (e.g., Jacoby, Kelley, et al., 1989).

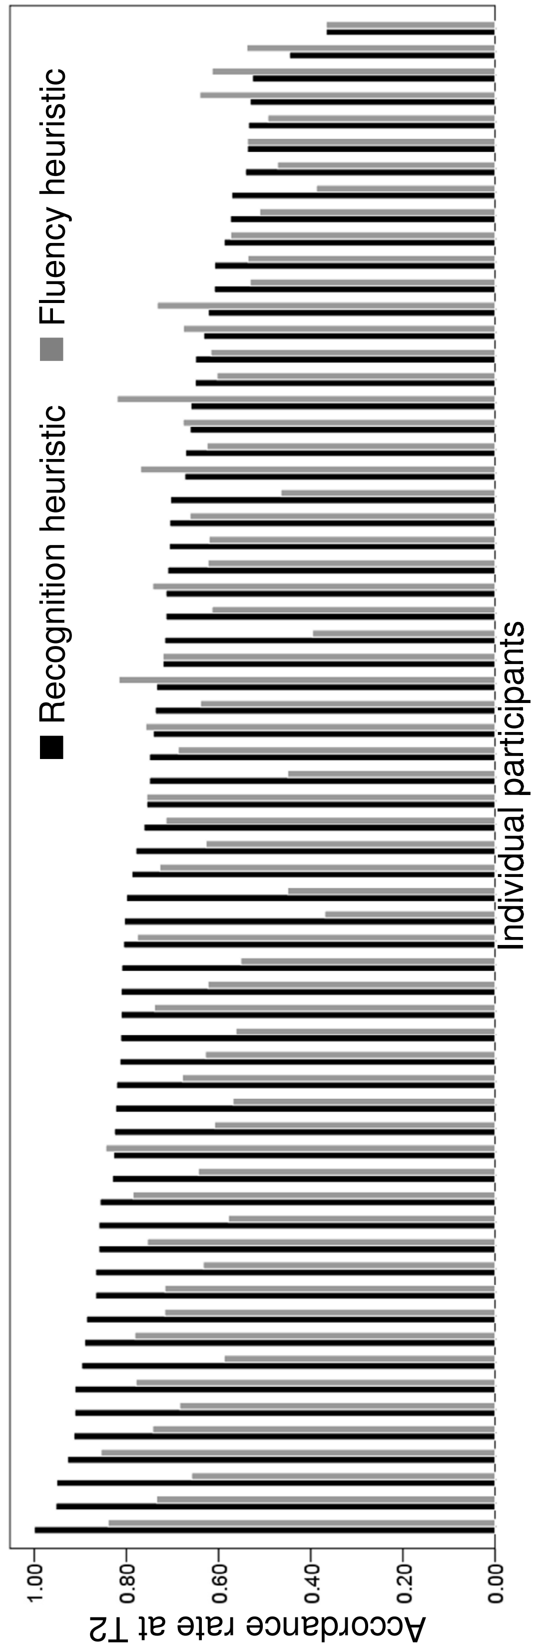

Figure 4.6. Recognition heuristic versus fluency heuristic. Pairs of bars show individual participants' accordance rates for the two models at T2 (second session of Experiment 12; $N = 65$).

Do voters use episodic knowledge to decide when to rely on recognition? At T2, voters sometimes erred when identifying recognition sources. Such source confusions constitute an interesting case. They allow for comparing voters' trust in recognition in two situations: when they (a) *correctly* identified the experiment as source, and (b) when they *falsely* believed their recognition was naturally acquired when in fact it was also only experimentally induced. Mentions of party names in studies are unlikely to reflect election outcomes; however, in the natural environment, recognition sources such as newspapers can reflect electoral success (Figure 4.2). Correspondingly, voters may have judged the predictive accuracy of recognition to be lower when they correctly tracked the experiment as source than when they falsely believed recognition was acquired in the natural environment. Therefore, voters should be more likely to treat parties as unrecognized when they have identified the experiment as the source than when they have identified the natural environment as the source.

To test this, I selected for each voter all pairs that consisted of a recognized and an unrecognized party at T1 (*RU pairs*) and that had become pairs of two recognized parties at T2 (*RR pairs*). I divided these pairs into two kinds. First, on *correct-source RR pairs*, at T2 voters correctly indicated that they only recognized the second, formerly unrecognized party from the experiment (at T1, this party was unrecognized; at T2, it was believed to be unrecognized before T1, recognized at T1, and unrecognized between T1 and T2; Figure 4.5). I expected that at T2 voters would continue to treat this second, formerly unrecognized party as unrecognized. In doing so, they would be able to use the recognition heuristic, and in this case, at T1 *and* at T2, they should forecast more votes for the party they already recognized at T1 than for the second party. As a result, the recognition heuristic accordance rates (based on recognition data from T1) should differ little between T1 and T2.

Second, *false-source RR pairs* were identical to correct-source *RR* pairs except that at T2 voters falsely believed they had recognized the second party *before* T1. On false-source *RR* pairs, voters should thus not treat the second party as unrecognized and, since the first party is also recognized from the natural environment, they should not be able to use the recognition heuristic. Instead they would have to resort to other heuristics, or guess. Therefore, the recognition heuristic accordance rate (based on recognition data from T1) should somewhat decrease from T1 to T2.

For 27 voters I identified at least one false-source *RR* pair (M = 22 pairs) *and* one correct-source *RR* pair (M = 81). To gain control over voters' access to cues, I excluded the eight larger parties. As predicted, on false-source *RR* pairs voters' recognition heuristic accordance rates dropped from T1 to T2, whereas there was little difference on correct-source *RR* pairs—a pattern I also found in the accuracy of voters' election forecasts (Tables 4.2, 4.3). When judging the magnitude of these effects, one has to keep in mind that the strategies people may have resorted to on false-source *RR* pairs may have mimicked some of the choices of the recognition heuristic, meaning that one would not necessarily expect the observed recognition heuristic accordance rate to drop to the chance level of .50 on false-source *RR* pairs.

Table 4.2

Mean (SE) recognition heuristic accordance rates and accuracy of voters' election forecasts on correct-source RR pairs and false-source RR pairs computed across pairs of smaller parties

	T1 correct	T2 correct	T1 false	T2 false
Accordance	.80 (.03)	.83 (.03)	.71 (.04)	.64 (.05)
Accuracy	.66 (.03)	.62 (.03)	.65 (.03)	.57 (.03)

Note. *n* = 27. Correct: correct-source *RR* pairs. False: false-source *RR* pairs. *RR*: two recognized parties. T1: Data collected approximately 16 days before the election. T2: Data collected approximately 2 days before the election. Note that this table depicts paired data. (For confidence intervals, see Table 4.3.)

Table 4.3

Confidence intervals (CI) on mean differences in recognition heuristic accordance rates and accuracy of voters' election forecasts on correct-source RR pairs and false-source RR pairs over time computed across pairs of smaller parties

	Mean difference	95% CI on the mean difference (Lower, upper)
Accordance	T2 correct–T1 correct	(−.03, .08)
	T2 false–T1 false	(−.16, .009)
Accuracy	T2 correct–T1 correct	(−.08, −.002)
	T2 false–T1 false	(−.13, −.02)

Note. *n* = 27. Correct: correct-source *RR* pairs. False: false-source *RR* pairs. *RR*: two recognized parties. T1: Data collected approximately 16 days before the election. T2: Data collected approximately 2 days before the election. Note that this table depicts paired data.

To summarize, Experiment 12 provided evidence for people's joint use of episodic and semantic knowledge to assess the predictive accuracy of recognition and apply the recognition heuristic even in pairs of two recognized objects: Episodic knowledge appears to have helped them identify recognition sources, and semantic knowledge informed them that one source (the lab) was unlikely to represent a mediator reflecting the criterion, while another (the natural environment) was likely to reflect it. Thus, when recognition was acquired in the natural environment, voters were more likely to trust it than when it was acquired in the lab.

Is a Lack of Knowledge Informative About the Predictive Accuracy of Recognition? (Experiment 13)

In an experiment series, Pohl (2006) asked people to categorize objects into those they recognized without knowing anything else about them (R^{-}) and those they recognized and had knowledge about (R^{+}). The recognition heuristic predicted people's inferences better on $R^{+}U$ *pairs* (i.e., where one object was recognized and there was knowledge about it and the other was unrecognized [U]) than on $R^{-}U$ pairs (i.e., where one object was recognized but there was no knowledge about it and the other was unrecognized). Does this finding indicate that people integrate cues in their inferences rather than rely on the recognition heuristic?

The recognition heuristic operates on judgments of recognition (Goldstein & Gigerenzer, 2002; Gigerenzer et al., 2008). By implementing this heuristic in ACT-R, Schooler and Hertwig (2005) provided a mathematical model of these recognition judgments (see also Chapter 3; Pleskac, 2007). They depend on an object's activation in memory, which is a function of the frequency of a person's past encounters with that object, which in turn determines the object's retrieval fluency and recognition speed, that is, the strength of its recognition signal. In the series of experiments and computer simulation studies with the ACT-R memory model reported in Chapter 3, I showed that objects about which people are unlikely to recall knowledge are apt to be less strongly activated than objects about which knowledge is likely to be available. Consequently, people are slower in retrieving and recognizing objects about which they are at the same time unlikely to retrieve knowledge. Since the recognized object's retrieval fluency tends thus to be larger in $R^{+}U$ pairs than in $R^{-}U$ pairs, it may often be harder to apply the recognition heuristic on $R^{-}U$ than on $R^{+}U$ pairs, resulting in lower recognition heuristic accordance rates on $R^{-}U$ pairs.

In fact, in Chapter 3 I also demonstrated that the way in which memory works can make it easier for a person to use a given heuristic when using it is also likely to result in accurate judgments. This may hold true for the recognition heuristic: When the recognition correlation is substantial, the probability of retrieving knowledge about an object correlates with the criterion. As a result, objects with knowledge (R^+) score on average higher on the criterion than objects without knowledge (R^-). Since both tend to have larger criterion values than unrecognized objects on average, R^+U pairs reflect larger differences on the criterion than R^-U pairs. This, in turn, may result in a stronger recognition correlation on R^+U than on R^-U pairs such that it may actually be ecologically rational to use the recognition heuristic more on R^+U than on R^-U pairs.

In short, one can formulate two competing hypotheses: (a) If Pohl's (2006) finding implies that people use compensatory strategies rather than the recognition heuristic, then such models should predict inferences better than the recognition heuristic. Alternatively, (b) if no model predicts people's inferences better than the recognition heuristic, then Pohl's finding would leave open the possibility that people rely less often on the heuristic on R^-U than on R^+U pairs because recognition is harder to assess and less predictive on R^-U than on R^+U pairs. While I continue to provide comparative model tests (for a comparison of R^-U and R^+U pairs, see Experiment 14), next I test whether the recognition correlation is smaller on R^-U than on R^+U pairs.

The setting of my experiment is the 2005 parliamentary election in North Rhine-Westphalia, which is the most populous German state. In contrast to the previously discussed elections, here voters only cast a ballot for a party and not directly for candidates. Similar to other elections, though, representatives enter Parliament as a function of the votes their parties gain.

Method

Sixty-one participants (44% female, mean age 26 years, *SD* = 3.6) filled out a questionnaire 3 to 11 days before the North Rhine-Westphalia election. About half of them completed the questionnaire in the labs of the Max Plank Institute for Human Development in Berlin and received €5 ($7) for their participation; the other half worked on it in a university class at Free University in Berlin, Germany. All participants had to be at least 18 years of age.

The questionnaire consisted of a ranking task, a recognition task, and a detailed recognition task on the 24 parties competing in the election. In the ranking task, voters were given a list of the party names and they assigned 1 of 15 ranks to each party (each rank could be assigned once) according to their forecasts of the number of votes the parties would win. The recognition task consisted of a list of the 24 parties. Voters indicated for each party whether they recognized it, that is, whether they had heard of the party name or seen it before participating in the study. In the detailed recognition task, I presented participants a list of all parties and asked them how much they knew about each party. There were three possible answers: (a) *never* heard of the party and *never* saw it before participating in the experiment (U); (b) heard of the party or saw it before but *do not know anything else* about it beyond recognizing its name (R^{-}); (c) heard of it or saw it before and know something about the party beyond simply recognizing it (R^{+}). I counterbalanced the ranking and recognition tasks; the detailed recognition task was at the end.[36] The order of appearance of parties was always randomized. Filling out the questionnaire took about 10 min.

Description of Measures

To evaluate how well the recognition heuristic predicts voters' forecasts in the ranking task, I needed a measure that would reach its maximum value if, of N parties, the n recognized parties are assigned the n highest ranks and the $(N – n)$ unrecognized parties the $(N – n)$ lower ranks. The more often the recognized parties are ranked lower than the unrecognized ones, the more the behavioral data will deviate from the model's prediction, something I also needed my measure to take into account. As it turns out, k fulfills these requirements. For each voter, I simply had to deconstruct the ranking by simulating that voter's complete (i.e., exhaustive) set of *virtual paired comparisons* consisting of one recognized and one unrecognized party and then compute the accordance rate using Equation 12.

The recognition correlation can be expressed in terms of the *recognition validity*, α, that is, the probability of a recognized object scoring higher on the criterion than an unrecognized one (Goldstein & Gigerenzer, 2002). Across virtual paired comparisons of recognized and unrecognized parties I counted for each participant the number of times T a recognized party had gained more votes than an unrecognized one and the number of instances W of the reverse happening:

$$\alpha = T/(T + W). \qquad (13)$$

[36] Two participants were excluded because they did not complete the questionnaire. Participants working on the ranking task before the recognition task recognized about as many parties, $M = 11.8$, $n = 28$, as those working on the tasks in the opposite order, $M = 13.0$, $n = 31$; 95% CI on the mean difference (-.2, 2.5). Also, the (paid) participants who filled out the questionnaire in the laboratory did not differ with regard to central variables from those participants who filled out the questionnaire in a university class.

Results and Discussion

Strength of the recognition correlation. As expected, the recognition validity was larger on R^+U pairs (M = .92, SE = .01) than on R^-U pairs (M = .68, SE = .03), 95% CI on the mean difference (.29, .19), n = 54. At the same time, the recognition heuristic accordance rate (computed across virtual comparisons between parties) was larger on R^+U pairs (M = .89, SE = .01) than on R^-U pairs (M = .62, SE = .03), 95% CI on the mean difference (.34, .21), n = 54. In short, people acted more strongly in accordance with the recognition heuristic when recognition was more predictive.

Model comparison 3. Overall, people were rather unlikely to rank recognized parties lower than unrecognized ones: When computed across all comparisons of recognized and unrecognized parties, the average accordance rate was .82 (SE = .01), which is consistent with the hypothesis that people generate consideration sets of recognized objects. To evaluate how well the recognition heuristic predicts behavior, I compared it to a generalization of take-one-cue to multiple objects (Table 4.1). Assuming that people had some knowledge about the size of the parties I classified as R^+, I implemented take-one-cue as follows. If a recognized party with knowledge (R^+) was small, then according to take-one-cue, a participant would rank this party lower than all unrecognized ones. If the R^+ party was large, the participant would rank the party higher than all unrecognized ones. The participant would rank recognized parties without knowledge (R^-) higher than all unrecognized ones.

Also in the multi-alternative case, take-one-cue assumes a free parameter, the cut-off criterion, C_2, which measures the weight of the party cue against recognition, determining which parties are small. For each participant, I classified all R^+ parties as small or large according to the results of the North Rhine-Westphalia election. That is, to cover all possible classifications (i.e., all possible values of the parameter C_2), in a first round for C_2 = 1, I only classified the party that won the election as large; all others were small. In a second round, I classified the two parties that won the most votes as large (C_2 = 2), and so on, until finally I classified all parties as large (C_2 = 24). For each participant, I computed the take-one-cue accordance (Equation 12) across virtual pairs of a recognized and an unrecognized party for all values of C_2. Does this increase in model complexity pay off?

It does not: Only when take-one-cue ranked recognized parties higher than unrecognized ones in nearly 100% of the virtual party pairs (from C_2 = 20 [99.9%] to C_2 = 24 [100%], Figure 4.7) did it fit people's forecasts, on average, as well as the recognition heuristic predicts them without assuming an additional parameter. From C_2 = 1 to C_2 = 19, take-one-cue generates the same rankings as the recognition heuristic in 35 to 98% of the pairs. Here, the recognition heuristic predicts voters' forecasts better.

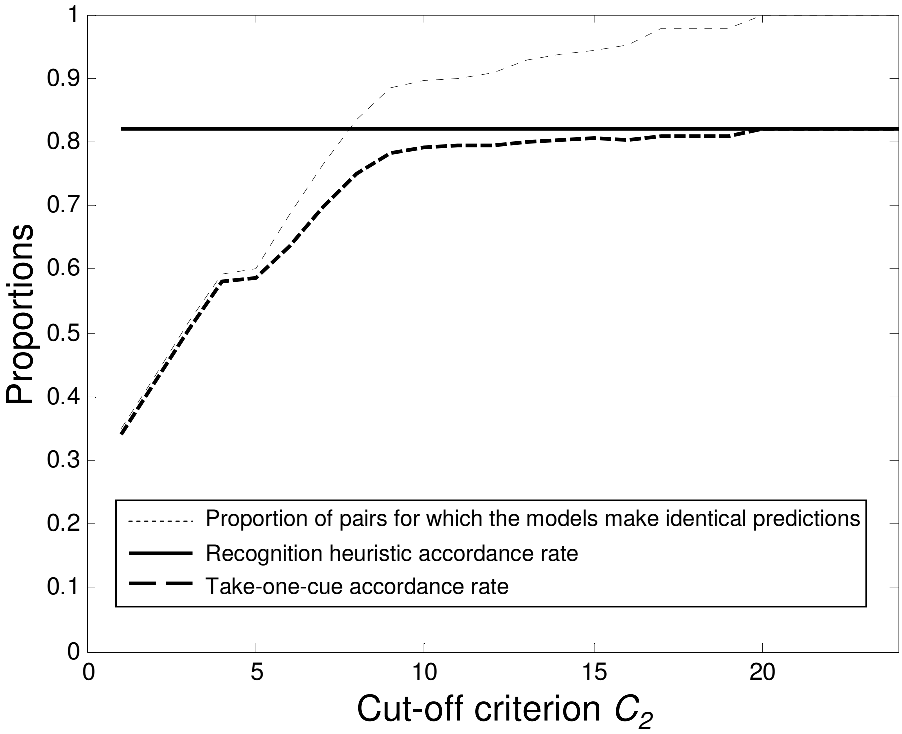

Figure 4.7. Recognition heuristic versus take-one-cue in a ranking task. Lines show mean accordance rates for the two models computed for each of the possible values of the cut-off criterion C_2 as well as the mean proportion of virtual paired comparisons in which the two models made the same predictions (*N* = 59; Experiment 13).

In short, the predictive accuracy of recognition is smaller and people act less strongly in accordance with the recognition heuristic when cues cannot be retrieved than when they are available. Consistent with the assumption that people generate consideration sets of recognized objects, participants ranked recognized parties higher than unrecognized ones. A more complex compensatory model did not fit their rankings better than the recognition heuristic predicted them.

How Well Do Compensatory Models Predict Individual Inferences?
(Experiment 14; Reanalysis of Experiment 7)

There might be individual differences in the weighting of cues, which I would have ignored in the comparison of the recognition heuristic and take-one-cue. In Experiments 14 and 15, I test the recognition heuristic against compensatory models that model individual differences. To begin, I pit it against *tallying-of-negative-cues* and *tallying-of-positive-and-negative-cues* (Table 4.1). These models compute the sum of positive and/or negative cues, weighting it against recognition. Except for the weighting, they are equivalent to tallying (Gigerenzer & Goldstein, 1996) and unit-weight linear models (e.g., Dawes, 1979; Einhorn & Hogarth, 1975).

What could be a compensatory alternative to the recognition heuristic when cues are not available (i.e., on $R^{-}U$ pairs)? Above, I argued that the speed with which the recognized object is retrieved, that is, its retrieval fluency, tends to be lower in $R^{-}U$ pairs than in $R^{+}U$ pairs, making it harder to use the recognition heuristic on $R^{-}U$ pairs than on $R^{+}U$ pairs, resulting in lower recognition heuristic accordance rates on $R^{-}U$ pairs. Yet, an alternative hypothesis is that people do not use this heuristic but systematically integrate retrieval fluency into their inferences (see Dougherty et al., 2008; B. R. Newell & Fernandez, 2006). Below, I pit the heuristic for the first time against a corresponding model: *weighted-fluency* (Table 4.1).

I reanalyzed data reported as Experiment 7 in Chapter 3 above. This data allowed me to examine inferences of city size—the task first used to test the recognition heuristic (Goldstein & Gigerenzer, 1999) and which most subsequent studies have used (e.g., B. R. Newell & Fernandez, 2006). I will remind the reader of Experiment 7 with a brief description of it here. To simplify references to this reanalysis later on in the discussion section, I will refer to it as Experiment 14 throughout.

Short Summary of the Methods Employed in Experiment 14

Forty-nine right-handed participants (43% female; mean age 24 years, *SD* = 3.1) completed a computerized experiment. They received a guaranteed payment of €13 ($17) supplemented by a performance bonus. Stimuli were 240 cities (i.e., of the 70 largest Austrian, British, French, German, Italian, Spanish, and U.S. cities, those consisting of 5 to 8 letters).

First, in a two-alternative forced-choice task, I presented two cities on a computer screen (one on the left and the other on the right) and asked participants to infer which city had more inhabitants. Pairs of cities were randomly drawn for each country without replacement such that each city could appear only once throughout the task (yielding 120 pairs of cities). On top of the guaranteed payment of €13, participants received €0.04 ($0.05) for each correct inference. Four cents was subtracted from this additional gain for each incorrect inference. (No feedback on the correctness of the responses was given until after the experiment.)

Next, in a recognition task, I gave participants the name of one city at a time and asked them to judge whether they had seen or heard of the city prior to participating in the experiment. Thereafter, in a detailed recognition task, I again presented one city at a time and asked participants how much they knew about each city. There were three possible answers: (a) never heard of it and never saw it before participating in the experiment; (b) heard of it or saw it before but do not know anything else about it beyond recognizing its name; (c) heard of it or saw it before and know something about the city beyond simply recognizing it.

Last, in a cue knowledge task, I asked participants to indicate for each city whether it (a) has an international airport, (b) has a university, (c) is a significant industry site, and (d) is a world-famous tourist site. Responses could be "yes" (positive cue value), "no" (negative cue value), or "I do not know" (unknown cue value). (Previous participants from the subject pool of the Center for Adaptive Behavior and Cognition considered these cues as the most useful ones for inferring city size; Pachur et al., 2008). Participants received €0.04 for each correct response on top of the guaranteed payment. For incorrect responses, €0.04 was subtracted from this additional gain. (No feedback on the correctness of the responses was given until after the experiment.) Participants did not receive payment nor did they lose money for "don't know" responses.

Participants took a 20-s break every other 12 trials (two-alternative choice, and recognition tasks) and every other 12 question blocks (cue knowledge task), respectively. In all tasks, each trial was preceded by a fixation cross for 1,000 ms, and participants were instructed to always fixate this cross when it appeared. In all tasks, participants were instructed to respond as quickly and accurately as possible. Positive responses were always made with the index finger of the right hand. In all tasks, the order of presentation of cities was always completely randomized.

Results and Discussion of Reanalyses

Model comparisons 4–6. Free parameters, each person's individual cut-off criteria, C_3, C_4, and C_6 (for C_5, see Experiment 15), determine when each alternative model infers unrecognized cities to be larger than recognized ones (Table 4.1). In tallying-of-negative-cues this happens when the number of negative cue values for the recognized objects falls below C_3. (There were four cues; depending on how many of them have negative values, C_3 can thus take values ranging from 0 to −4.) Tallying-of-positive-and-negative-cues decides against recognized cities when the sum of negative and positive cue values is smaller than C_4. (C_4 could thus take values ranging from −4 to +4.) In weighted-fluency this happens when the retrieval time for the recognized city falls above

C_6, that is, when it takes much time to judge a city as recognized. As in Experiment 12, I operationalized retrieval time as recognition time, which I measured in terms of participants' reaction times in the recognition task (see Hertwig et al., 2008).

Table 4.1 shows when the three models decide in favor of recognized cities. In tallying-of-negative-cues this is the case when no negative cue values can be recalled, or when the number of negative cue values falls above or is equal to C_3. In tallying-of-positive-and-negative-cues, this happens when no negative and positive cue values are available, or when the sum of positive and negative cue values falls above or is equal to C_4. Weighted-fluency decides in favor of recognized cities when the retrieval time for the recognized city falls below C_6, that is, when it takes little time to judge a city as recognized.

When comparing nested models that differ in the number of free parameters, as is the case for the recognition heuristic and its competitors, a good model evaluation criterion is their generalizability to new data (see also Chapter 2). I evaluated the models' ability to generalize to new data in a computer simulation. To pit the two cue-based models against the recognition heuristic, I selected for each participant those pairs of cities in which, according to her responses in the detailed recognition task, she had knowledge about the recognized city (R^+U pairs). To test weighted-fluency against the recognition heuristic, I chose the pairs for which there was no knowledge (R^-U pairs). I divided all R^+U pairs and all R^-U pairs, respectively, 10,000 times randomly into two halves. The first half represented the calibration set in which I calculated for each of 49 participants that person's *optimal* values for C_3, C_4, and C_6 for the three alternative models (i.e., the value at which a model's accordance rate is largest; 49 participants × 3 models = 147 parameter values; optimal values are derived by exhaustively searching the entire parameter space). I used these values to compute the proportion of inferences consistent with each model in the other half, the validation set, where the models' generalizability is evaluated. For each partition, I also computed the recognition heuristic accordance rates. For almost all participants the recognition heuristic predicted inferences better than each of the three alternative models (Figures 4.8, 4.9).

I additionally assessed the models' ability to fit existing data (see Chapter 2). To this end, I calculated the optimal values for C_3, C_4, and C_6 for each model (exhaustively searching the entire parameter space) for each participant's complete set of R^+U pairs and R^-U pairs, respectively, computing the associated accordance rates. The recognition heuristic and its competitors are nested models; however, in contrast to the competitors, it has no free parameter. This is why when the values for C_3, C_4, and C_6 are optimal, the recognition heuristic's accordance rate can *never* exceed its competitors' accordance rates; but, its competitors' accordance rates can exceed the recognition heuristic's accordance rate. Did they?

Corroborating my previous results, for most participants, recognition had a noncompensatory weight; that is, the optimal values for C_3, C_4, and C_6 resulted in the respective alternative model always inferring recognized cities to be larger than unrecognized ones just as the recognition heuristic does. Therefore, the accordance rates for the recognition heuristic and its competitors are identical for most participants (Figures 4.8, 4.9).[37]

Strength of the recognition signal. As in Experiment 13, the recognition validity was larger on R^+U pairs (M = .81, SE = .01) than on R^-U pairs (M = .74, SE = .02) while the recognition heuristic accordance rate was also larger on R^+U pairs (M = .95, SE = .01) than on R^-U pairs (M = .86, SE = .02). Consistent with the hypothesis that it is more difficult to apply this heuristic when a person using it is less likely to make accurate inferences with it (see Experiment 13; see also Chapter 3), in the two-alternative forced-choice task, inferences on R^-U pairs took on average 208 ms longer (mean of median response time, M_{Mdn} = 1,872 ms, SE = 79) than inferences on R^+U pairs (M_{Mdn} = 1,664, SE = 69), 95% CI on the mean difference (126, 290), N = 49. Similarly, as reported above in Chapter 3, in the recognition task, judging an R^- object as recognized took on average 183 ms longer (M_{Mdn} = 834 ms, SE = 33) than judging an R^+ object as recognized (M_{Mdn} = 652, SE = 12), 95% CI on the mean difference (127, 238); N = 49.

[37] This fitting exercise can also be thought of in terms of a regression analysis: Variables beyond recognition do not improve the fit for most participants.

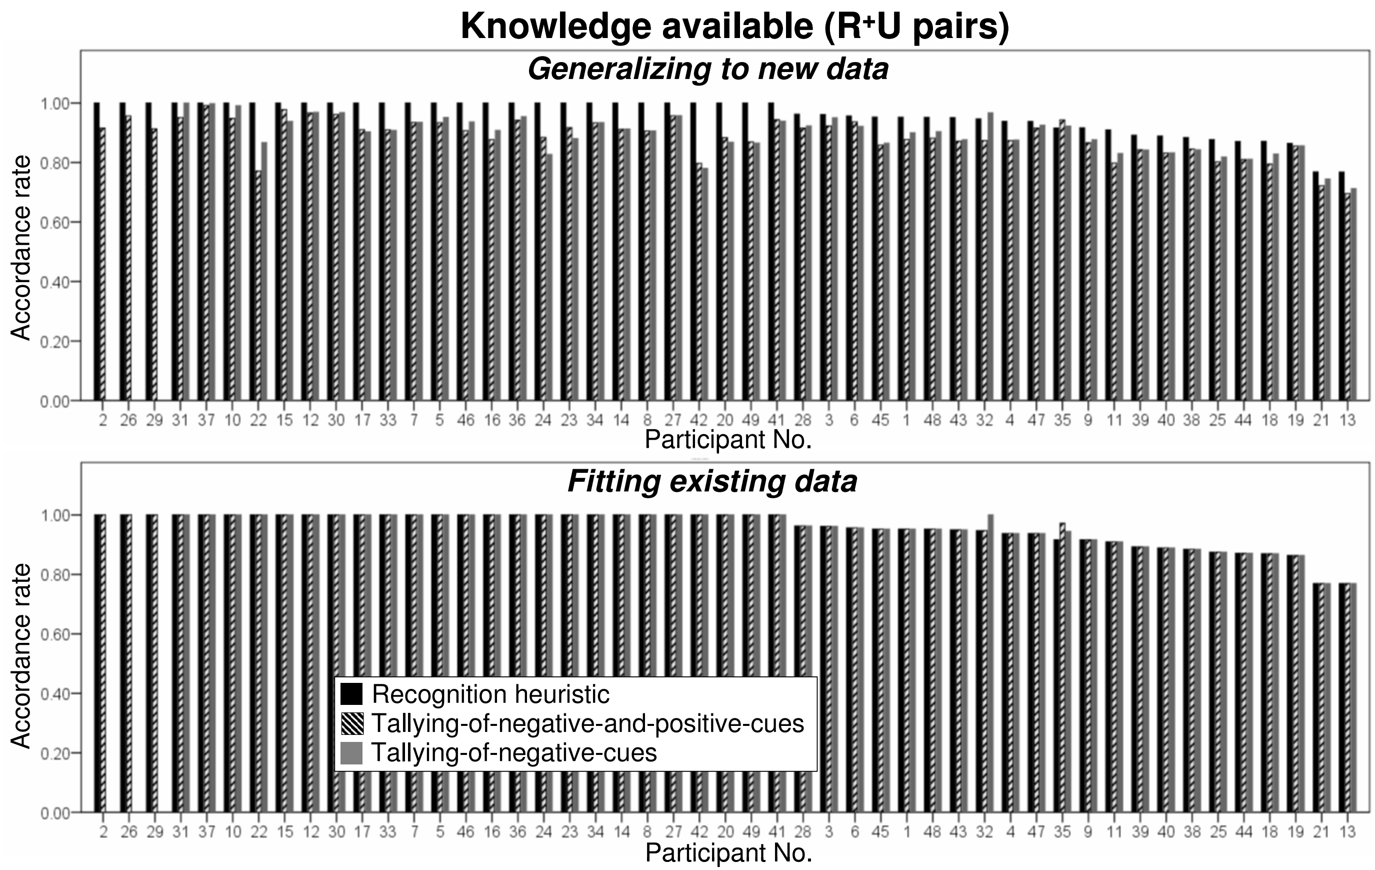

Figure 4.8. Recognition heuristic versus tallying-of-positive-and-negative-cues and tallying-of-negative-cues in a two-alternative forced-choice task. *Generalizing to new data*. Triplets of bars show means of participants' accordance rates for the three models computed in the validation set across 10,000 random partitions of each participant's data. For 2 participants (i.e., nos. 32, 35), one or both competitors predicted decisions better than the recognition heuristic; for all others the recognition heuristic predicted decisions best. *Fitting existing data*. Bars show participants' accordance rates for the three models. For 42 participants, recognition had a noncompensatory weight in *both* competitors of the recognition heuristic, resulting in the competitors always making the same decisions as the recognition heuristic does. For 2 participants (i.e., nos. 32, 35), compensatory weights resulted in the better fit for one or both of the recognition heuristic's competitors. For 2 other participants (i.e., nos. 1, 25), compensatory weights yielded the same fit as noncompensatory weights. In both panels, for participants 2, 26, and 29 no accordance rates for tallying-of-negative-cues are shown because unlike all other participants, they did not identify negative cues; their recognition weights for positive cues were noncompensatory ($N = 49$; Experiment 14).

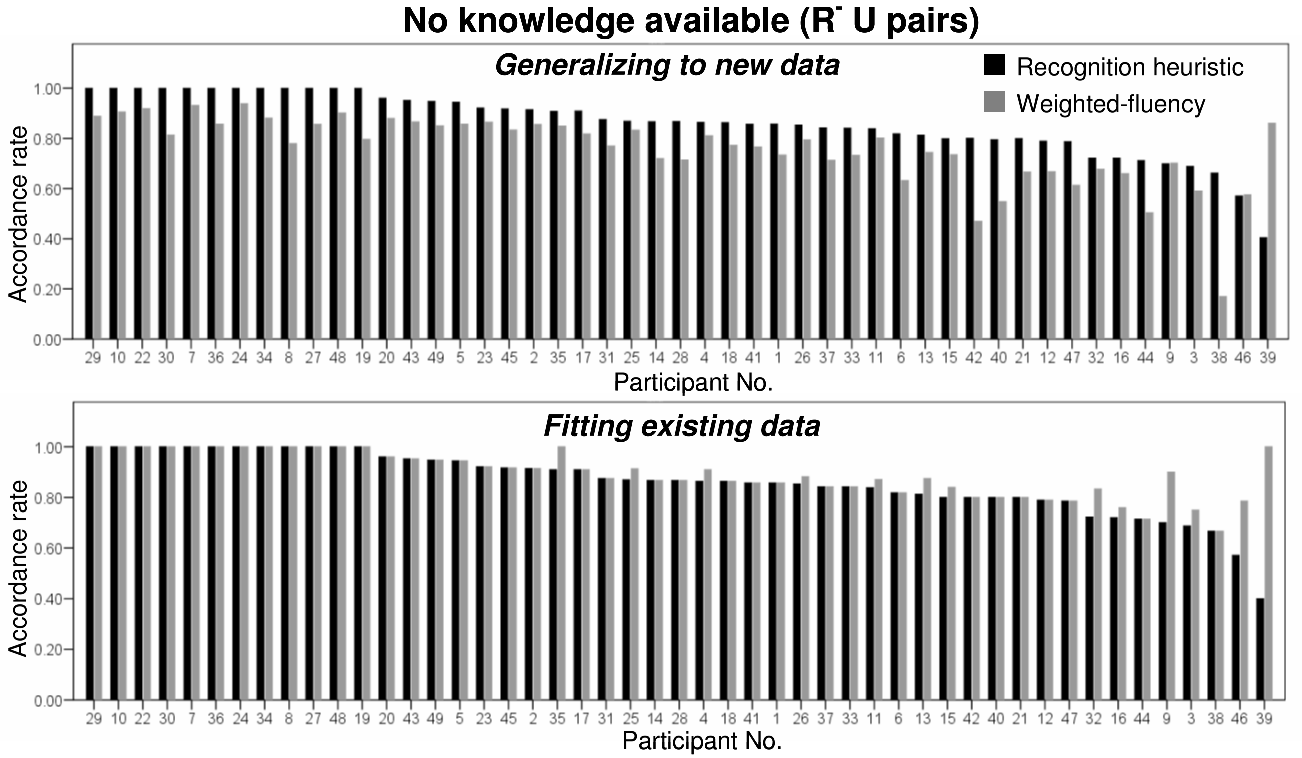

Figure 4.9. Recognition heuristic versus weighted-fluency in a two-alternative forced-choice task. *Generalizing to new data*. Pairs of bars show means of participants' accordance rates for the two models computed in the validation set across 10,000 random partitions of each participant's data. For 3 participants (i.e., nos. 9, 39, 46), weighted-fluency predicted decisions best; for all others the recognition heuristic was best. *Fitting existing data*. Bars show participants' accordance rates. For 35 participants, recognition had a noncompensatory weight, resulting in weighted-fluency always making the same inferences as the recognition heuristic does. For 13 participants, compensatory weights resulted in a better fit for weighted-fluency. For 1 participant (i.e., no. 44), compensatory weights yielded the same fit as noncompensatory ones ($N = 49$; Experiment 14).

Do People Use the Best Cues Rather Than Recognition? (Experiment 15)

In situations with many objects, people might integrate only a few cues rather than many (see Ford et al., 1989; Payne et al., 1993). Next, I pit the recognition heuristic against a corresponding model for consideration-set generation: weighted-best-cues (Table 4.1).

Method

Twenty-seven participants (44% female, mean age = 25 years, *SD* = 3.7) filled out a questionnaire in the labs of the Max Planck Institute for Human Development approximately 2 years after the 2005 German national election. I modified the design of Experiment 13 slightly: In a ranking task, participants ranked the 25 parties that ran in this election according to the election outcome they would expect if the next national election were to take place on the subsequent day. They also completed a recognition task and a detailed recognition task (identical to those in Experiment 13). Besides individual differences in the weighting of recognition (as modeled with individual values for the cut-off criteria C), the cues people use to forecast elections may differ from person to person *and* from party to party. For instance, for one party a person may consider the fact that this party has lost in past elections to be most informative, but for another party the same person may take the fact that it lacks a charismatic leader as most indicative of electoral success. Another person may consider other cues most relevant. In a cue knowledge task, I presented participants a list of the party names. Here, for each party they had some knowledge about, participants identified the cue they considered to be most relevant for inferring the election outcomes for *this particular* party. They described this cue on a blank line beside the party name and assigned a value to it (scale ranging from −100 to 100, with −100 signifying that the cue very strongly indicates that this party will win few votes, 0 signifying that the cue speaks for neither many nor few votes, and 100 signifying that the cue very strongly indicates that the party will win many votes). Filling out the questionnaire took about 30 min.

Results and Discussion

Model comparison 7. I implemented weighted-best-cues (Table 4.1) as follows: If a participant had knowledge about a recognized party (R^+), and if the relevant cue was assigned a value below this participant's individual cut-off criterion C_5, then that participant would rank the recognized party lower than all unrecognized ones (U). Conversely, if the cue had a value above or equal to C_5, the participant would rank the recognized party higher than all unrecognized ones. If the participant had no knowledge (R^{-}), she would assign the recognized party a higher rank than all unrecognized ones.

As in Experiment 14, I assessed the models' generalizability in a computer simulation. For each participant, I split the parties 10,000 times randomly into two halves, one representing the calibration set, to compute the participant's optimal value for C_5, and the other representing the validation set, with which the generalizability of the models could be assessed (see Chapter 2). For each participant, I first calculated the optimal value for C_5 in the calibration set (exhaustively searching the entire parameter space) and then used this value to compute the weighted-best-cues accordance rate in the validation set. For each participant, I also computed the recognition heuristic accordance rate in each partition of parties. The recognition heuristic is not only the simpler model; it also predicts people's election forecasts better: Computed as the average over 10,000 partitions, its accordance rate was larger than that of weighted-best-cues for 25 of 26 participants (Figure 4.10; for the 27th participant see Figure caption).

Next I assessed the models' ability to fit existing data. To this end, I calculated the optimal value for C_5 for each participant's complete data set (exhaustively searching the entire parameter space) and computed the associated weighted-best-cues accordance rates. Since the two models are nested, when C_5 is optimal, the weighted-best-cues accordance rate can *never* be smaller than the recognition heuristic accordance rate. However, the weighted-best-cues accordance rate can exceed the recognition heuristic accordance rate. Did it? This was the case for only 9 of 26 participants (Figure 4.10; for the 27th participant see Figure caption). For 16 participants, the recognition weight was noncompensatory; that is, the optimal value for C_5 resulted in weighted-best-cues always ranking recognized parties higher than unrecognized ones, just as the recognition heuristic does. Since the models' predictions are thus identical for these participants, in Figure 4.10 the accordance rates for the two models are the same.

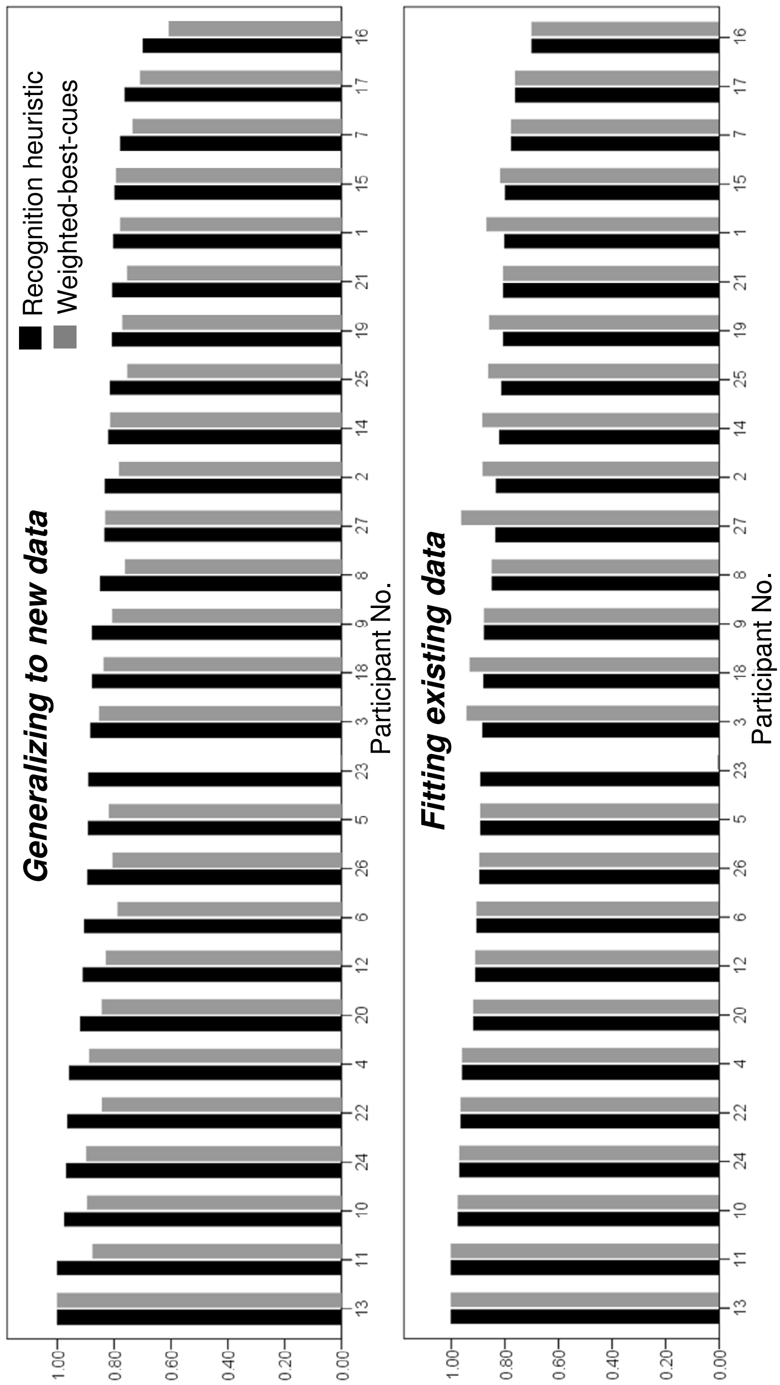

Figure 4.10. Recognition heuristic versus weighted-best-cues in a ranking task. *Generalizing to new data*. Pairs of bars show means of participants' accordance rates for the two models computed in the validation set across 10,000 random partitions of each participant's data. For all participants except 2 (i.e., nos. 13, 23), the recognition heuristic predicted participants' rankings best. Participant no. 13 assigned the weights in a way such that both models always predicted the same rankings. *Fitting existing data*. Bars show participants' accordance rates for the two models. For 16 participants, the recognition weight was noncompensatory, resulting in weighted-best-cues always ranking parties in exactly the same way as the recognition heuristic does. For 9 participants, compensatory weights resulted in the better fit for weighted-best-cues. For 1 participant (i.e., no. 7), compensatory weights yielded the same fit as noncompensatory weights. In both panels, for participant no. 23 no weighted-best-cues accordance rate is shown because unlike all other participants, this person did not identify knowledge ($N = 27$; Experiment 15).

Consideration-set generation. While cues thus played little role in the elimination of unrecognized parties, people used them to rank recognized parties *within* the consideration set: Parties with a larger cue value were ranked higher than those with a lower one in, on average, 84% (SE = 2) of all virtual comparisons between two R^+ parties that differed in the cue value.

To summarize Experiments 14 and 15, first, in generalizing to new data, the recognition heuristic predicted most participants' behavior better than four compensatory heuristics. Second, in fitting existing data, the recognition heuristic accounted for most participants' behavior as well as the competing heuristics did. Third, the recognition heuristic was the simpler model. Judged on these criteria, it is the better model.

Does Time Pressure Foster the Default of Relying on the Recognition Heuristic? (Experiment 16)

Pachur and Hertwig (2006) provided evidence to suggest that a sense of recognition arises before knowledge about an object is retrieved from memory, pointing to one reason *why* this heuristic may be used by default. In their studies, people inferred recognized diseases that are known to be virtually extinct to be more prevalent than unrecognized ones when they had to make such inferences under time pressure. Apparently, the lack of time did not allow them to complete the evaluative process that may be involved in overruling a possible default of using the recognition heuristic. In support of this thesis, in Experiment 12 I replicated another of Pachur and Hertwig's findings. Calculating median response times across all paired comparisons between two parties in the two-alternative forced-choice task (i.e., at T1), I found that inferences in agreement with the heuristic were made faster (M_{Mdn} = 1,409 ms, SE = 40) than those in disagreement with it, which were most likely produced by other strategies (M_{Mdn} = 2,919, SE = 157), 95% CI on the mean difference (-1,794, -1,225), N = 65.

Next, I will examine whether the recognition heuristic's advantage over other strategies extends to heuristics operating on retrieval fluency rather than knowledge. To this end, I will pit the recognition heuristic against weighted-fluency. If the recognition heuristic is used by default, and executing the evaluative process needed to overrule this default is time consuming, then one would expect time pressure to increase the percentage of participants whose inferences can be better modeled by the recognition heuristic than by weighted-fluency.

Method

An experiment was completed by 173 right-handed participants (64 % female, mean age 25 years, *SD* = 3.9) in the labs of the Max Planck Institute for Human Development. They received a guaranteed payment of €8 ($12) supplemented by a performance bonus.

The tasks differed from those reported in Experiment 14 only in the following respects. Participants were randomly assigned to one of two conditions: under time pressure (*n* = 87) or not under time pressure (*n* = 85) in the two-alternative forced-choice task. To put participants under time pressure, each presentation of a pair of cities started with a blank screen. After 900 ms, an acoustic signal (Tone 1) sounded, followed by a second signal (Tone 2) 900 ms later that coincided with the presentation of a fixation cross. Again 900 ms later, upon the sound of a third signal (Tone 3), the fixation cross was replaced by a pair of cities. Participants were instructed to respond as quickly and accurately as possible but not later than a fourth, imagined signal (Tone 4, imagined), 900 ms after the third signal and stimulus onset. I used an imaginary signal to avoid possible interference of a real signal with the ongoing processing of a city pair. (A similar procedure is employed in lexical decision tasks, e.g., Wagenmakers et al., 2004, and has been used by Pachur & Hertwig, 2006.) If a response was markedly delayed, that is, 1,200 ms after the third tone, an aversive tone sounded and the message "too late" appeared for 900 ms on the screen. Responses that were made in time (i.e., < 1,200 ms) were followed by the presentation of a blank screen for 900 ms. Then the next trial started. On top of the guaranteed flat fee of €8, participants were paid €0.04 ($0.06) for each correct response. Incorrect responses resulted in a subtraction of €0.04 from this additional payment. (No feedback on the correctness of the responses was given until after the experiment.) Regardless of whether the response was correct or wrong, responses that were followed by the message "too late" always resulted in a loss of €0.04. In the experimental condition without time pressure, participants had the same amount of time between each presentation of a pair of cities (3,600 ms) as in the condition with time pressure.

To acquaint participants with the two-alternative forced-choice task, they took 20 practice trials, each consisting of a pair of arrows (">" and "<", randomized). The task was to indicate whether the ">" appeared on the left or right side of the screen. Depending on the experimental condition, these practice decisions had to be made under time pressure or not.

To test whether there was an order effect, I randomly chose one third of the participants in each experimental condition and assigned them to complete the recognition task before the two-alternative forced-choice task; all others worked on the task in the opposite order.[38] In all tasks, the fixation cross preceding a trial was shown for 900 ms.

Results and Discussion

Model comparison 8. First, I evaluated the models' ability to generalize to new data in a computer simulation. As in Experiment 14, for each participant I selected those pairs of cities in which, according to the responses in a detailed recognition task, she had no knowledge about the recognized city ($R^{-}U$ pairs). I divided each participant's $R^{-}U$ pairs 10,000 times randomly into two halves, one representing the calibration set and the other the validation set. I calculated for each participant the optimal value for C_6 in the calibration set (exhaustively searching the entire parameter space), using this value to compute the proportion of inferences consistent with weighted-fluency in the validation set. For each partition, I also computed the recognition heuristic accordance rate.

Corroborating my results in Experiment 14, the recognition heuristic predicted most participants' inferences better than weighted-fluency, regardless of whether participants were set under time pressure in the two-alternative forced-choice task. As expected, the percentage of participants whose inferences were better predicted by the recognition heuristic than by weighted-fluency was larger under time pressure (Figure 4.11).

Second, I assessed the models' ability to fit existing data, calculating the optimal value for C_6 (exhaustively searching the entire parameter space) on each participant's complete set of $R^{-}U$ pairs and computing the associated accordance rates. Recall that the recognition heuristic and weighted-fluency are nested models; however, in contrast to the latter, the recognition heuristic has no free parameter. This is why when the value for C_6 is optimal, the recognition heuristic's accordance rate can *never* exceed weighted-fluency's accordance rate; but weighted-fluency's accordance rate can exceed the recognition heuristic's accordance rate.

[38] Participants working on the two-alternative forced-choice task before the recognition task recognized about as many cities (M = 123.7; n = 112) as those working on the tasks in the opposite order (M = 123.3; n = 61). Also, the percentage of participants whose inferences were better modeled by weighted-fluency as opposed to the recognition heuristic was about the same in both situations.

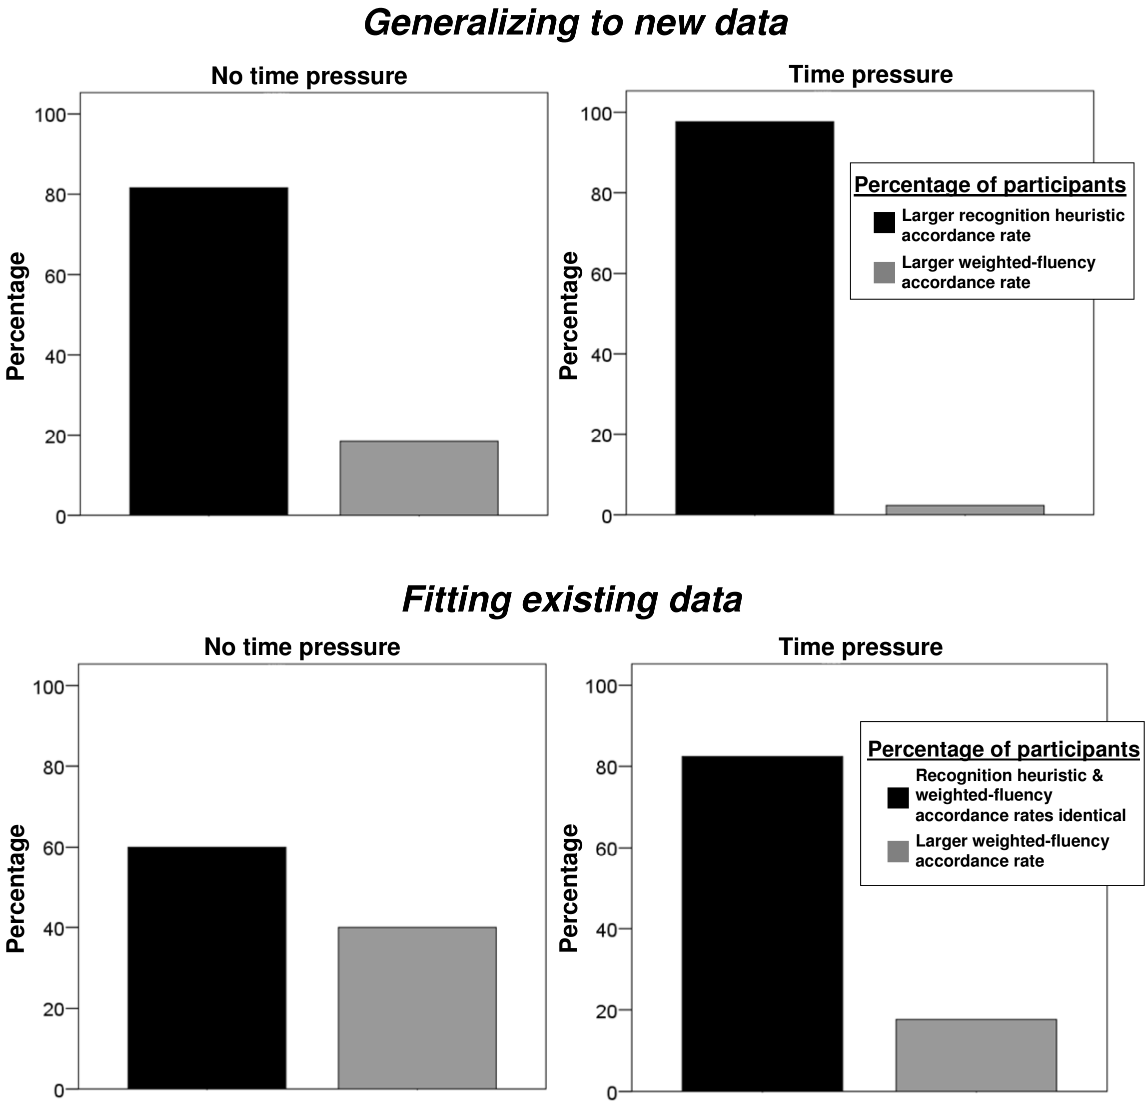

Figure 4.11. Recognition heuristic versus weighted-fluency in a two-alternative forced-choice task when inferences were or were not made under time pressure. *Generalizing to new data.* The black bars show the percentage of participants who had a larger recognition heuristic accordance rate than weighted-fluency accordance rate. The grey bars show the percentage of participants who had a larger weighted-fluency accordance rate than recognition heuristic accordance rate. Accordance rates are means computed in the validation set across 10,000 random partitions of each participant's data. *Fitting existing data.* The black bars show the percentage of participants for whom the recognition weight was noncompensatory, resulting in weighted-fluency always making the same inferences as the recognition heuristic does. The grey bars show the percentage of participants who had a larger weighted-fluency accordance rate than recognition heuristic accordance rate. Data from one participant are excluded in both panels, because unlike all other participants, this one did not make inferences on pairs of cities about which he had no knowledge ($R^{-}U$ pairs; $n_{\text{time pressure}} = 87$; $n_{\text{no time pressure}} = 85$; Experiment 16).

Mirroring my results in the generalization test, for most participants, recognition had a noncompensatory weight; that is, the optimal value for C_6 resulted in weighted-fluency inferring recognized cities to be larger than unrecognized ones just as the recognition heuristic does. Therefore, the accordance rates of the two models are identical for most participants. As expected, the percentage of participants for whom the models' accordance rates were the same was larger when participants were set under time pressure (Figure 4.11). These results suggest that a possible default of applying the recognition heuristic is less likely to be overruled by low retrieval fluency when there is little time to execute the associated evaluative processes.

General Discussion

In this chapter, I (i) re-formulated the recognition heuristic for tasks with multiple objects, and (ii) for tasks in which all objects are recognized. I (iii) formally specified a range of compensatory and noncompensatory models as alternatives to the recognition heuristic. Assuming that the recognition heuristic is used by default, I (iv) specified under what conditions people employ it and when they instead rely on other strategies. In eight model comparisons in six experiments, I tested the recognition heuristic for the first time against alternative models.

Forecasting Political Elections with Mere Recognition

Before addressing the main points of this chapter, I would like to comment on the domain in which I investigated the recognition heuristic: political elections. Around the globe, much money is spent on forecasting elections by interviewing voters about their party preferences, political attitudes, or opinions of candidates. Intriguingly, my findings indicate that asking voters which candidates or parties they recognize, or even simply counting how often a candidate or party is mentioned in the press, could sometimes be sufficient to forecast elections. In fact, further bootstrapping analyses show that under certain circumstances recognition may even help to predict elections more accurately than traditional voting polls (Marewski, Gaissmaier, Schooler, & Gigerenzer, 2008).

Voters rely on simple rules of thumb to make decisions (Gigerenzer, 1982, 2007; Jackman & Sniderman, 2002; Kelley & Mirer, 1974; Regenwetter, Ho, & Tsetlin, 2007; Sniderman, 2000; Wang, 2008). However, while I believe that recognition predicts electoral success because it increases the chance for a party or a candidate to be included in voters' consideration sets, I would not conclude that people vote for parties or candidates solely because they recognize them. Often voters will rely on more information to decide which of possibly many recognized parties or candidates to vote for, and here one can think of strategies operating on voters' likes of candidates, their attitudes about political issues, candidates' party affiliations, or their perceived competence (Campbell, Converse, Miller, & Stokes, 1960; Kelley & Mirer, 1974; Todorov, Mandisodza, Goren, & Hall, 2005; Wang, 2008; see also Brady & Sniderman, 1985). At the same time, empirical evidence suggests that voters not only take into account the desirability of candidates but also their likelihood of being elected (Stone & Abramowitz, 1983), and, as my findings show, the recognition heuristic helps voters eliminate unrecognized candidates and parties that have low chances of being elected from the consideration sets of candidates and parties potentially worth a vote.

Generalization of the Recognition Heuristic

I formulated two generalizations of the recognition heuristic. In what follows, I will focus on the generalization to multiple objects. (I will discuss the generalization to two recognized objects in the section on strategy selection further below.) The recognition heuristic generalization to multiple objects generates consideration sets of recognized objects with large criterion values, which can be evaluated using other heuristics. By eliminating all unrecognized objects, this simple strategy reduces complexity, and since people are unlikely to have knowledge about unrecognized objects, putting them aside or assigning their criterion values at random does not necessarily imply a loss of accuracy. In Experiments 13 and 15, I provided evidence that people rely on the heuristic when forecasting elections in sets of up to 25 parties. This result complements another, which suggested that the heuristic may be used when choosing among up to four objects (Frosch, Beaman, & McCloy, 2007).

When scanning through objects, recognition is highly accessible information (Pachur & Hertwig, 2006). My thesis that people rely on it by using the recognition heuristic is consistent with theories of consideration-set generation that assume noncompensatory heuristics (e.g., Kohli & Jedidi, 2007). It is also consistent with work in consumer choice showing that priming a familiar brand increases the probability that it will be considered for purchase (Coates et al., 2004), while at the same time only a single exposure can lead people to consider buying a novel brand (Coates et al., 2006). Moreover, my work complements recent advances in model building: The recognition heuristic was initially proposed as the first step in *take-the-best* (Gigerenzer & Goldstein, 1996), a noncompensatory heuristic for two-alternative decisions. Hogarth and Karelaia (2005) generalized take-the-best to the multi-alternative case.[39] The recognition heuristic could be the first step in their model in which one would eliminate all unrecognized objects. In a second step, one would evaluate recognized objects based on the best cues. In fact, while the cues people considered to be best for forecasting elections played little role in the elimination of unrecognized parties, they predicted their rankings within the consideration set of recognized parties (Experiment 15).

Toward a Theory of Strategy Selection by Default

Much work in decision making, memory, and beyond has tackled the problem of how different strategies, operators, routines, or forms of processing are selected (e.g., Busemeyer & Myung, 1992; Gray et al., 2006; Logan, 1988; Lovett & Anderson, 1996; Payne et al., 1993; Rieskamp & Otto, 2006; see Chapter 3 for details). Here, I have sketched four ways in which the use of the recognition heuristic may hinge on the retrieval of knowledge, as well as objects' retrieval fluency. In essence, all address the ecological rationality of overruling a possible default application of this heuristic. First, semantic knowledge about cues in conflict with recognition appears to lead some people not to use the recognition heuristic. Second, semantic knowledge suggesting that the recognition correlation is weak may overrule the default. In the literature, related findings have been interpreted as challenging the plausibility of the recognition heuristic, rather than as an indicator of heuristic selection (e.g., Richter & Späth, 2006). While I have discussed these two ways of overruling the default above, here I will focus on the third and fourth way.

Episodic knowledge. A long research tradition in cognitive psychology has stressed the importance of source evaluation processes for behavior (see Johnson, Hashtroudi, & Lindsay, 1993). As I have shown, a possible default use of the recognition heuristic can be overruled when episodic knowledge about the surrogate correlation between recognition and the mediator, combined with semantic knowledge about the ecological correlation between the mediator and the criterion (Figure 4.1), suggests that recognition is not predictive of the criterion. In Experiment 12, the recognition heuristic modeled people's inferences better when they identified the natural environment as recognition source (a mediator correlating with the criterion) than when they took

[39] Hogarth and Karelaia (2005) called their take-the-best generalization *deterministic elimination by aspects*. Note that take-the-best differs from *elimination by aspects* (Tversky, 1972). The latter is a model of preferential choice that has no deterministic rule to order cues (i.e., attributes), and it is not specified how to compute cues' weights. Instead, it has an aspiration level for each cue and cues are quantitative. Take-the-best, in turn, is a model of inference operating on cues with binary values and a deterministic, specified order of cues.

the lab to be the source (a mediator uncorrelated with the criterion).[40] Importantly, Dougherty et al. (2008) criticized the recognition heuristic for not being useful in comparisons of *two* recognized objects. Providing another generalization of this heuristic, I demonstrated that this combination of episodic and semantic knowledge allows people to use the recognition heuristic even when both objects are recognized.

My findings explain earlier work: In an experiment by Oppenheimer (2003) people were unlikely to rely on recognition to infer the size of cities when they knew that they recognized cities for reasons dissociated with the criterion. My data are also consistent with Jacoby, Kelley, et al.'s (1989) "overnight fame" experiments in which people seemed to rely on experimentally induced recognition to infer other people's fame when they were unable to recall the source of recognition. At the same time, they tended to ignore recognition when they could trace it to stemming exclusively from the experiment—a mediator that is unlikely to reflect the criterion fame. Similarly, Goldstein and Gigerenzer (2002) induced false recognition and observed how after the passing of several weeks, people would more often infer objects with induced recognition to be larger than novel, unrecognized ones. It is likely that as time passed people were unable to track the experiment as the recognition source, and so they relied on induced recognition even though it was not predictive of the criterion.

In this experiment, Goldstein and Gigerenzer (2002) also showed that the accuracy of people's inferences decreases when they are unable to use the recognition heuristic because all objects from a set end up being recognized. In such situations people have to rely on cues or fluency, or guess instead, which can be less accurate than relying on recognition. Recently, much research has investigated whether and when such *less-is-more effects* occur (Dougherty et al., 2008; Gigerenzer et al., 2008; Pachur & Biele, 2007; Pohl, 2006; Reimer & Katsikopoulos, 2004; Snook & Cullen, 2006). My findings have implications for this work: When a person can correctly recall recognition sources and treat some recognized objects as unrecognized, less-is-more effects may not emerge (see also Footnote 27 in Chapter 3).

[40] The recognition heuristic is a model for situations where people make inferences based solely on information retrieved from memory (as opposed to *inferences from givens*; Gigerenzer & Goldstein, 1996), and where recognition is acquired in the world prior to participating in a study (Goldstein & Gigerenzer, 2002). As I have shown (Experiment 12), when people believe recognition is acquired prior to participating in a study, they can have good reasons to believe that environmental mediators link recognition to the criterion. However, when they track a study as the source, they are less likely to trust recognition. In two experiments (Bröder & Eichler, 2006; B. R. Newell & Shanks, 2004), recognition was experimentally induced shortly before a two-alternative forced-choice task in such a way that the study could be easily tracked as the recognition source. Bröder and Eichler found that people were less likely to trust such induced recognition in inferences from memory when the experimenter taught them cues at odds with it. B. R. Newell and Shank's data indicated that people were less likely to use induced recognition when such cues could be read off a computer screen, that is, in inferences from givens.

Strength of the recognition signal. In keeping with theories that assume memory processes as major determinants of decision behavior (e.g., Dougherty et al., 1999; Schooler & Hertwig, 2005) and strategy selection (Chapter 3 above), I suggested that a possible default use of the recognition heuristic will be overruled when recognized objects' retrieval fluency is low, which in turn indicates that recognition is unlikely to be predictive of the criterion. Using ACT-R's memory model to predict people's recognition judgments, retrieval and recognition time distributions, and knowledge in nine experiments, in Chapter 3 I found strong correlations between the probability of a person recognizing an object, the object's retrieval and recognition times, and the probability of retrieving knowledge cues about the object. That is, objects about which people are likely to recall some knowledge (R^+) tend also to be more strongly activated in memory and more quickly retrievable than recognized objects about which no knowledge is available (R^-). As a result, it will often be easier to apply the recognition heuristic in pairs that include an R^+ object and an unrecognized object (R^+U pairs) than in R^-U pairs. In fact, not only are R^+ objects recognized more quickly than R^- objects, but also inferences are made faster on R^+U pairs than on R^-U pairs (Experiment 14). This way, memory makes it easier for a person to rely on the recognition heuristic when using it is also likely to result in accurate inferences: As I have shown for inferences about electoral success and city size, recognition validities tend to be larger on R^+U than on R^-U pairs (Experiments 13, 14). As a side note, by reanalyzing data from Chapter 3, I have replicated this result, for instance, for inferences about companies' market capitalization and countries' gross domestic product (Table 4.4). In other words, it is ecologically rational to use the recognition heuristic more on R^+U than on R^-U pairs, and indeed, I observed that people's inferences are more likely to agree with this heuristic on R^+U than on R^-U pairs (Experiments 13, 14).

Table 4.4

Recognition validities when people can recall knowledge and when they cannot

Alternative; *criterion*	Knowledge recalled (R^+U pairs)		Knowledge lacking (R^-U pairs)		95% CI on the mean difference in validity	
	M	*SE*	*M*	*SE*	Lower	Upper
164 countries; *GDP*	.91	.02	.72	.02	.16	.22
80 companies; *market capitalization*	.83	.01	.76	.03	.01	.13

Note. Data taken from Experiments 1 and 2, reported in Chapter 3. GDP: gross domestic product in 2006. Companies' market capitalization is as of May 31, 2007. This table depicts paired data.

Importantly, in Experiments 13 to 16, I found that the recognition heuristic predicts people's decisions better than fluency-based and knowledge-based strategies on R^+U *and* R^-U pairs (see also Experiment 12). The outcomes of these experiments provided evidence against plausible alternative explanations for my findings (see Pohl, 2006), namely, that systematic differences in recognition heuristic accordance rates on R^+U and R^-U pairs are produced by people's overall reliance on fluency- or knowledge-based strategies. Rather, my data are consistent with the hypothesis that they are a result of suspensions of the recognition heuristic, triggered by a weak recognition signal and low recognition validity in R^-U pairs.

My results are consistent with findings by B. R. Newell and Fernandez (2006; see also Hertwig et al., 2008) suggesting that objects' retrieval fluency impacts on the use of the recognition heuristic. This thesis is also supported by Volz et al.'s (2006) fMRI data. They found that inferences in accordance with the recognition heuristic correlated with higher activation in brain areas that had previously been associated with greater recognition confidence. My data additionally indicate that a possible default to rely on the recognition heuristic is less likely to be overruled by low retrieval fluency when there is little time to execute the associated evaluative processes (Experiment 16).

How Do the Cognitive Niches of the Recognition Heuristic and the Fluency Heuristic Differ?

In Chapter 3, I proposed a theory of strategy selection, arguing that the interplay between the workings of memory, the environment, and people's decision strategies (i) constrain the choice set of applicable heuristics, giving rise to what I called non-overlapping cognitive niches of heuristics. This mind–environment interplay can (ii) make it faster and easier to rely on a given heuristic when using it is also likely to result in accurate inferences. In Chapter 4, I focused less on the applicability of the recognition heuristic as a determinant of strategy selection but more on its accuracy. In what follows, I will discuss the cognitive niche of the recognition heuristic, explaining how it differs from the niche of the fluency heuristic.

In what respects are the two heuristics similar? The recognition heuristic and the fluency heuristic are models of inference. They resemble each other in at least three ways. First, in contrast to many strategies receiving input from knowledge cues, both heuristics operate on the accessibility of memories rather than on knowledge. Second, the two heuristics share very similar ecological rationales, characterized by a set of correlations between an unknown criterion, an environmental mediator, and memory (compare Figures 3.1 and 4.1). Correspondingly, it is likely that the ecological rationality of using either one of the heuristics can be informed by similar information. For instance, I argued that people are unlikely to apply the recognition heuristic when they know that the correlative link between recognition and the criterion is not substantial; it is to be expected that the same rationale would also hold true for the fluency heuristic. In fact, the results of Experiment 12 might be interpreted in this light, suggesting that people do not rely on the fluency heuristic when retrieval times have been manipulated experimentally, that is, when they have been influenced by an environmental mediator that is not related to the criterion to be inferred. Third, as I have shown in my computer simulation studies and experiments, for both heuristics a basic principle seems to hold: The interplay between memory and the environment can make it easier and faster to apply them when relying on these heuristics is also most likely to result in accurate inferences.

In what respects are the two heuristics different? While there are thus some similarities between the heuristics, their cognitive niches seem to differ fundamentally in at least two regards. First, as pointed out in Chapter 3, the two heuristics' cognitive niches essentially do not overlap. Unless discriminating source knowledge can be retrieved (see Table 4.1; p. 94 and Footnote 33), when deciding between two objects, the recognition heuristic is only applicable when a person recognizes one object but not the other. The fluency heuristic, in contrast, requires representations of both objects to be stored in memory, simplifying strategy selection between these two heuristics. Second, the cognitive niches of the two heuristics differ with respect to their overlap with the cognitive niches of knowledge-based heuristics. Let me first illustrate this point for the recognition heuristic, and then for the fluency heuristic.

Consider the situation when just one object is recognized in a pair. Here, the recognition heuristic is applicable in two cases, one being characterized by a weak recognition signal for the recognized object and a lack of knowledge about that object (i.e., as in $R^{-}U$ pairs), the other entailing a strong recognition signal and the likely accessibility of further knowledge (i.e., as in $R^{+}U$ pairs). In this second situation, the heuristic is most likely to result in fast and accurate inferences. This is also the situation in which the recognition heuristic's cognitive niche overlaps with that of knowledge-based strategies, over which it tends to be selected. In the first situation, in turn, namely, when the recognition signal is weak and knowledge tends to be unavailable, the recognition heuristic has no competition with knowledge-based strategies. However, since the recognition signal is weaker, using the heuristic also takes more time and is less likely to aid making accurate inferences.

Now consider the situation when both objects are recognized in a pair. In Chapter 3, I showed that the fluency heuristic is most likely applicable when a person has little or no knowledge about the two objects in question (i.e., as in $R^{-}R^{-}$ or $R^{-}R^{+}$ pairs). In such situations of limited knowledge, differences in retrieval times tend to be large and easier to detect, favoring the applicability of the fluency heuristic. A person is less likely to be able to apply the fluency heuristic when knowledge is abundant, because in this case differences in retrieval times tend to be small (i.e., as in $R^{+}R^{+}$ pairs). When both knowledge-based strategies and the fluency heuristic are applicable, that is, when their cognitive niches overlap, knowledge-based strategies tend to result in more accurate inferences than the fluency heuristic, which depends on noisy retrieval times. In fact, when both types of strategies can be applied, people seem to rely on knowledge-based strategies rather than the fluency heuristic. Figure 4.12 provides a schematic summary of these results, depicting the applicability of the fluency heuristic, recognition heuristic and knowledge-based strategies.

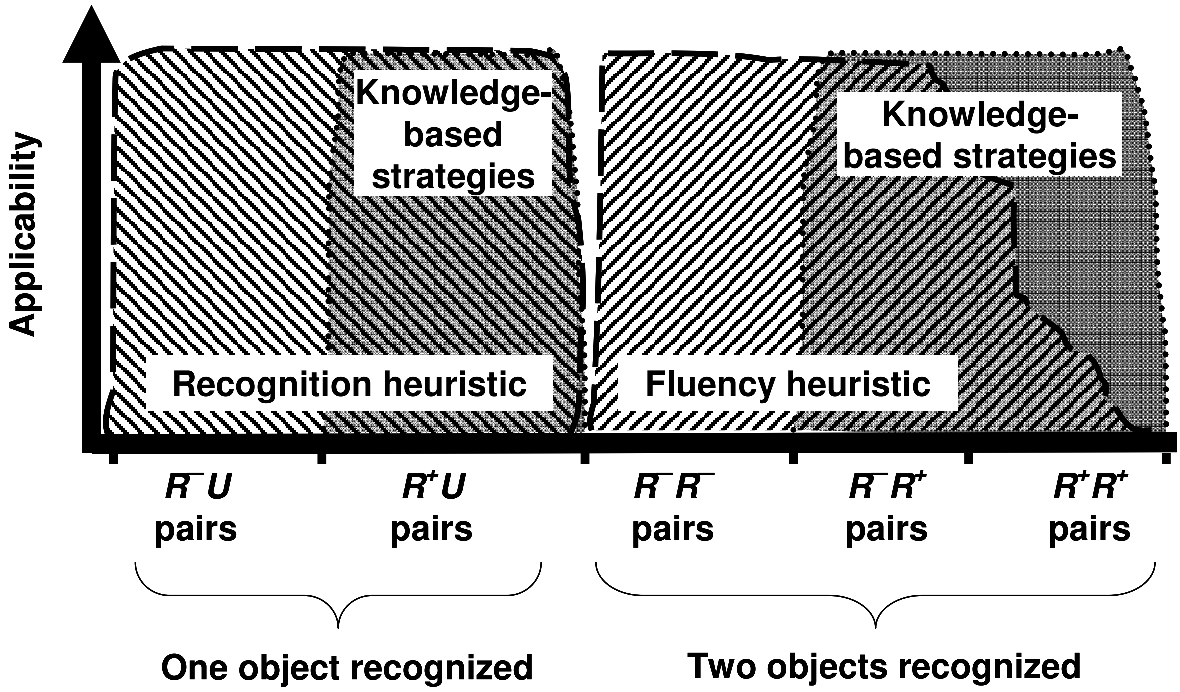

Figure 4.12. Schematic representation of the applicability of different strategies as a function of recognition and knowledge. In the case of knowledge-based strategies, applicability can be thought of as the probability of retrieving discriminating cues; in the case of the recognition heuristic, applicability is the likelihood of recognizing one object but not the other; and in the case of the fluency heuristic, applicability is the probability of detecting a difference in retrieval times between two objects. The grey shaded areas depict the applicability of knowledge-based strategies; the striped areas illustrate the applicability of the recognition heuristic and the fluency heuristic, respectively. Where the grey shaded areas and the striped areas intersect, the cognitive niches of different strategies overlap. The letter *R* indicates that an object is recognized. About recognized objects knowledge may be available (R^+) or not (R^-). The letter *U* indicates that an object is unrecognized (see also text above).

In short, the recognition heuristic and the fluency heuristic differ in two respects. First, with the strategies essentially being applicable in different situations, their cognitive niches tend not to overlap. Second, in contrast to the fluency heuristic, the cognitive niche of the recognition heuristic is not characterized by a lack of knowledge—while the fluency heuristic is dominated by knowledge-based strategies, the opposite holds true for the recognition heuristic, making it a default strategy.

Why do the recognition heuristic and the fluency heuristic differ? Let me speculate about possible answers to this question by first contrasting the recognition heuristic and knowledge-based strategies, then comparing the recognition heuristic to the fluency heuristic, and finally pitting the fluency heuristic against knowledge-based strategies.

There seem to be at least three reasons why the recognition heuristic dominates its knowledge-based counterparts. First, recognition, the input information of the recognition heuristic, is often more quickly accessible in memory than knowledge, a hypothesis that is supported by both previous experimental work (Pachur & Hertwig, 2006; see also Experiment 16) and my modeling reported in Chapter 3: A comparison of Figures 3.3 and 3.4 shows that the activation of chunks representing knowledge tends to be smaller than the activation of chunks representing recognition, resulting in faster retrieval times for recognition than for knowledge. As a result, it may often be easier and faster to execute the recognition heuristic than to rely on knowledge. Second, the cognitive operations required to use recognition can be boiled down to fewer productions than the cognitive operations that may be necessary to instantiate many knowledge-based strategies: Often these require the adding of different cues (e.g., tallying strategies), or the ordering of many cues (e.g., lexicographic strategies); relying on the recognition heuristic, in contrast, requires just using one piece of information. Third, relying on knowledge will often not enable a person to make more accurate inferences than relying solely on recognition.

The amount of time and effort necessary to execute the fluency heuristic, in turn, also depends on the speed with which fluency information can be accessed in memory. This speed is identical to the speed of unfolding recognition and can be modeled in terms of Equations 4 and 5. However, in contrast to the recognition heuristic, which simply operates on a binary recognition judgment, the fluency heuristic requires making a comparative judgment of two continuous variables, namely, of the speed of recognizing each of two objects. Successfully executing this comparative judgment may be relatively more time consuming and difficult than making a simple binary recognition judgment. To give an analogy: We can rather quickly and easily decide whether a lightbulb in a room is turned on or off; however, it may be more difficult to decide which of two lightbulbs in a room is shining brighter. The on–off judgment can be thought of as the binary recognition judgment; the comparative judgment of brightness may translate into a judgment of relative fluency. In fact, in Experiment 16, I found that the recognition heuristic seems to be more quickly executable than weighted-fluency, another fluency-based heuristic.

As my findings in Chapter 3 suggest, successfully executing a comparative judgment of relative fluency may also often be more time-consuming and difficult than retrieving and relying on knowledge: Recall, it is only when available knowledge is sparse and differences in retrieval times are very large and easy to detect that an application of the fluency heuristic seems to involve less time than applying knowledge-based strategies (see Experiment 9). At the same time, as I have shown in Experiment 7, relying on knowledge will often enable a person to make more accurate inferences than relying solely on fluency, which depends on noisy retrieval times.

In short, it may be the greater accessibility of recognition, the simplicity of the operations required to use the recognition heuristic, and the high validity of recognition that make the recognition heuristic a default heuristic compared to knowledge-based strategies. The fluency heuristic, in turn, depends on the successful execution of a comparatively difficult judgment of relative retrieval fluency, which moreover is subject to noise in retrieval times, letting this heuristic often be dominated by its more accurate, and more easily executable, knowledge-based competitors.

Note that these results seem counterintuitive: One might have believed that the recognition heuristic and the fluency heuristic would be very similar, and in fact in the literature, many have argued that the recognition heuristic and the fluency heuristic are instantiations of the same notion, availability, doubting that one should assume different heuristics (see B.R. Newell & Fernandez, 2006; Dougherty et al., 2008). My work highlights that such claims warrant detailed models of memory in order to be substantiated. In fact, Schooler and Hertwig (2005) deliberately decided to conceptualize the fluency heuristic and the recognition heuristic in terms of two different sets of ACT-R production rules, speculating that these two heuristics might differ and hoping that by treating them separately, they would be able to retain the freedom to identify such differences. As my findings indicate, they were right.

To summarize, for situations in which using the recognition heuristic and the fluency heuristic are ecologically rational, my results in Chapters 3 and 4 suggest the following sequence of strategy choices: If two objects are unrecognized, then a person is likely to guess. If one object is recognized, and no knowledge can be retrieved about it, then a person is likely to apply the recognition heuristic. If one object is recognized, and knowledge is available, then a person is also likely to rely on recognition. If both objects are recognized and no knowledge is available, a person is likely to apply the fluency heuristic. If both objects are recognized and knowledge can be retrieved about only one, then a person is likely to apply knowledge-based strategies unless the difference in retrieval time between the two objects is very large and judgments have to be made rather quickly. If both objects are recognized and knowledge is also available about both, then people are most likely to use this knowledge. As explained above, this sequence is also likely to result in fast, effortless, and accurate decisions. In short, the strategies in our repertoire are sorted into different cognitive niches, and it is easier and faster to use a given strategy when applying it is also most likely to lead to making accurate inferences, simplifying strategy selection this way.

Alternative Models of Inference: A Competition

Addressing concerns about the adequacy of the recognition heuristic as a model of behavior (e.g., Dougherty et al., 2008; B. R. Newell & Shanks, 2004; Pohl, 2006), I showed that this heuristic predicts inferences better than alternative models operating on cues or retrieval fluency (Experiments 11–16). This is the first time that such formal model comparisons have been conducted for the recognition heuristic. I would like to point out that prior to carrying out the model competition few people would have anticipated that the recognition heuristic would fare so well. For example, at least one of my coauthors on a corresponding journal article (Lael J. Schooler) declared that he would have bet on the fluency-based models. My modeling exercise highlights an important methodological point. Often, informal reasoning may not suffice to evaluate a model; instead, a formal comparison of competitors may be more informative.

I hasten to add that some people are better fitted by alternative models than by the recognition heuristic (Figures 4.6, 4.8–4.11). In a series of experiments and in reanalyses of earlier studies (i.e., B. R. Newell & Fernandez, 2006; Richter & Späth, 2006), Pachur et al. (2008) provided evidence to suggest that such inter-individual variability represent a recurring phenomenon in people’s reliance on recognition. In fact, individual differences are common in many decision strategies. It has repeatedly been found that strategy use depends on the characteristics of the individual, the task environment, and the strategy itself (e.g., Bergert & Nosofsky, 2007; Bröder & Gaissmaier, 2007; Bröder & Schiffer, 2006; Gaissmaier et al., 2006; von Helversen & Rieskamp, 2008; Mata et al., 2007; Payne et al., 1993; Rieskamp & Hoffrage, 1999, 2008; Rieskamp & Otto, 2006).

To conclude, I would like to stress once more that much research has investigated how people make decisions based on a sense of recognition, fluency, availability, familiarity, or accessibility (e.g., Bruner, 1957; Dougherty et al., 1999; Hertwig et al., 2005; Jacoby, Kelley, et al., 1989; Tversky & Kahneman, 1973; Winkielman et al., 2003). In pitting the recognition heuristic against competing models, I do not mean to imply that one or the other strategy will always be relied upon. Rather, in response to work claiming that the recognition heuristic is an inadequate model of behavior, I provide evidence that *in certain situations* recognition is relied upon in the ways specified by the recognition heuristic. I identified some of these situations and generalized the recognition heuristic to others. Our minds do not come equipped with just one or two strategies. In keeping with the fast and frugal heuristics (e.g., Todd & Gigerenzer, 2000) and other frameworks (e.g., Beach & Mitchell, 1978; Hogarth & Karelaia, 2007; Payne et al., 1988, 1993), I believe the mind makes use of a repertoire of strategies. This adaptive toolbox contains compensatory, noncompensatory, as well as many other decision-making mechanisms that compete for their adaptive use in strategy selection. The important question is what determines the outcomes of this competition—that is, when do people rely on which cognitive strategy?

Chapter 5

Summary and Conclusion

People make decisions under the constraints of limited time, knowledge, and information-processing capacity—be they about the likely performance of stocks, which movie to watch in the cinema, or for which job to apply. According to the fast and frugal heuristics research program (e.g., Gigerenzer et al., 1999; Todd & Gigerenzer, 2000), such decisions can be made successfully because people can rely on a repertoire of simple rules of thumb, called heuristics. These simple decision strategies can perform well because they exploit the structure of information in the environment in which a decision maker acts and build on the ways that evolved cognitive capacities, such as the human memory system, work. Together, the heuristics form an adaptive toolbox of the cognitive system, allowing decision makers to respond adaptively to different decision situations by relying on the heuristics that are appropriate for a task.

I started this dissertation with a look at three closely interrelated questions of the fast and frugal heuristics framework, namely, (i) what heuristics people use, (ii) when they use them, and (iii) to which environmental structure the heuristics are adapted (e.g., Gigerenzer et al., 2008). Specifically, I focused on the problem of how people choose among strategies when making decisions by retrieving all relevant information from memory (i.e., inferences from memory, Gigerenzer & Goldstein, 1996). My research brought together different psychological theories—about decision strategies, decision makers' environments, memory, and strategy selection. While my work on decision strategies and the environment was grounded in the fast and frugal heuristics framework, parts of this dissertation linked this research program to other ecological approaches to psychology, in particular to John R. Anderson and colleagues' ACT-R cognitive architecture (e.g., Anderson et al., 2004; Anderson & Lebiere, 1998). This helped me to develop a quantitative model of the interplay between memory and the environment.

In Chapter 1, I provided an overview of the fast and frugal heuristics research program and the ACT-R theory of cognition. Setting some of the methodological preliminaries for my empirical work, I gave a short summary of different model selection criteria in Chapter 2. In the experiments and simulation studies in Chapters 3 and 4, I then tackled the three questions from the fast and frugal heuristics research program, showing which heuristics people use in inferences from memory, when they use them, and how they nestle into the structure of the environment. In what follows, I will provide a short summary of the empirical work I conducted.

How Memory Aids Strategy Selection

Many theories explicitly or implicitly assume two determinants of strategy selection: the strategies' costs (e.g., operationalized in terms of the time or effort involved in using them), and their accuracy. For instance, a relatively old assumption in the decision-making literature is that people engage in cost–benefit trade-offs when selecting between available strategies (see Beach & Mitchell, 1978; Christensen-Szalanski, 1978; Payne et al., 1988, 1993), say, by using a meta-strategy to trade the cognitive effort involved in relying on a strategy against its accuracy in making decisions. Also more recent ideas formulated in the fast and frugal heuristics framework have focused on accuracy, information costs, or the time and effort required to execute a strategy (e.g., Bröder 2003; Bröder & Schiffer, 2006; Mata et al., 2007; B. R. Newell & Shanks, 2003; Pachur & Hertwig, 2006; Payne et al., 1988; Rieskamp & Hoffrage, 1999; Rieskamp & Otto, 2006; Volz et al., 2006).

Complementing these earlier proposals, in Chapter 3 I argued that (i) the workings of the human cognitive system *constrains* the choice set of strategies that can be applied to solve a given task, giving rise to what I called different *cognitive niches* of the heuristics in the adaptive toolbox. In particular, I proposed that these niches emerge from how memory represents regularities in the environment. This mind–environment interplay (ii) can favor particular strategies that are not only easy and fast to execute, but also result in accurate inferences—an instance of ecologically rational strategy selection and a situation in which the speed–accuracy trade-off relations typically assumed in theories of strategy selection are violated.

To elaborate my thesis, I proposed and tested a formal model of memory that quantifies the ways in which the human memory system interacts with the environment. This ACT-R memory model makes systematic quantitative predictions about memory retrieval based on environmental data. Specifically, the model allows predicting a person's recognition of an object and knowledge about it, as well as the associated retrieval time distribution.[41] This information, in turn, can be used to predict what decision strategies that person will employ.

In 4 simulation studies and 10 experiments with a total of 500 participants, I considered the choice between strategies that depend on knowledge, that is, on the content of what is retrieved, and strategies that—by operating on the accessibility of memories, such as retrieval time—depend on the characteristics of the retrieval rather than on the retrieved contents. One such strategy that operates on retrieval time is the fluency heuristic (Schooler & Hertwig, 2005), a simple rule of thumb for

[41] As is important in the context of modeling reaction time distributions, my ACT-R model actually allows for predicting retrieval time distributions that are characterized by an increasing skew and spread as a function of memory activation, which in turn depends on the environmental pattern of occurrence of objects (see Figure 3.5).

inferring which of two objects (e.g., cars) scores higher on a given criterion (e.g., car quality). According to this heuristic, the object that can be retrieved faster from memory is likely to score higher. Thus, the fluency heuristic's applicability depends on a decision maker's ability to tell which of two objects is retrieved faster. In my experiments and computer simulation studies, I showed that this heuristic is most likely applicable when a person has little or no knowledge about the objects in question. In such situations of limited knowledge, differences in retrieval times tend to be large and easier to detect, favoring the applicability of the fluency heuristic. A person is less likely to be able to apply the fluency heuristic when knowledge is abundant, because in this case differences in retrieval times tend to be small. Knowledge-based strategies, in turn, can only be relied upon when knowledge is available. Correspondingly, people will most likely be able to use the fluency heuristic when they cannot rely on knowledge instead, and vice versa, illustrating how different cognitive niches can simplify the choice between decision strategies. By submitting people's reliance on different heuristics to tests, my experiments also helped to answer the question of what heuristics people use, pitting the particularly simple fluency heuristic for the first time against strategies operating on knowledge.

In Simulation Study 1, I calibrated my ACT-R memory model to the behavioral data from one experiment, showing that the model can fit people's recognition and knowledge about objects in the world, as well as the associated retrieval time distributions. In a generalization test, I then found that the model accurately predicts these data in my other experiments. In Simulation Study 2, I used this predicted data to show how the magnitude of differences in retrieval times between two objects varies as a function of people's knowledge about these objects, providing evidence that the memory model can accurately predict different cognitive niches for the fluency heuristic and knowledge-based strategies. In Simulation Studies 3 and 4, I used the predicted recognition, knowledge, and retrieval time data once more to show that the model accurately predicts when a person using the fluency heuristic would be most likely to make accurate inferences. This is the case when available knowledge is sparse and differences in retrieval times between two objects are large, which is also the situation when the fluency heuristic is likely to allow for making rapid, effortless judgments. Experiments 1–10 provided the behavioral data necessary to calibrate and test the ACT-R model. Experiments 7–9 additionally yielded tests of people's reliance on the fluency heuristic, suggesting that this heuristic best predicts people's inferences in a two-alternative choice task when no knowledge is available (Experiments 7, 8), when differences in retrieval times are large so that the heuristic is easy to use (Experiment 8), or when inferences need to be made under time pressure while available knowledge is sparse, and differences in retrieval times are large (Experiment 9).

Models of Recognition-based Multi-alternative Inference

The focus of Chapter 4 was the recognition heuristic (Goldstein & Gigerenzer, 2002), a simple rule of thumb for inferring which of two objects (e.g., two cities), one recognized and the other not, has a larger value on a quantitative criterion (e.g., city size). When there is a strong positive correlation between a person's recognition of objects and the criterion, a person using this heuristic would predict that recognized objects are likely to score higher on the criterion. That is, the heuristic operates on a binary sense of prior encounter, which is how recognition is defined here. In doing so, the heuristic assumes a noncompensatory use of recognition: Even when a person could rely on knowledge cues (e.g., facts about a city) to complement her recognition of an object, when the heuristic is used to make inferences about that object, these cues are ignored.

Recently, the recognition heuristic has been criticized for being a model of inference between a recognized and an unrecognized object only (Dougherty et al., 2008); that is, in principle the heuristic cannot be relied upon when all objects in a set are recognized. Moreover, findings that not all people always make inferences consistent with the heuristic have raised doubts about its adequacy as a model of behavior (e.g., Bröder & Eichler, 2006; B. R. Newell & Shanks, 2004; Oppenheimer, 2003; Pohl, 2006). For instance, Richter and Späth (2006) ran a series of studies and—observing that fewer decisions were consistent with the recognition heuristic when cues that contradicted recognition were available—concluded that there was no evidence of a noncompensatory use of recognition. According to them, there was clear evidence that recognition is integrated with knowledge. To illustrate, one could hypothesize that people evaluate objects by weighting and adding their values on a range of knowledge cues, such that an object's low value on one cue can be compensated for by a high value on another cue. In a related vein, based on another set of studies, B. R. Newell and Fernandez (2006) suggested that people might rely on a more graded sense of recognition (e.g., retrieval fluency or ease of retrieval) rather than on the recognition heuristic. Yet, such conclusions may be premature since in both study series the corresponding alternative hypotheses—compensation by integration, or people's reliance on ease of retrieval—were not formally specified and formulated as testable models. Moreover, contradicting Richter and Späth's findings, other studies—including reanalyses of Richter and Späth's data (Pachur et al., in press)—did provide evidence suggesting that people rely on the noncompensatory recognition heuristic (e.g., Goldstein & Gigerenzer, 2002; Pachur & Biele, 2007; Pachur et al., 2008). Two of these studies even reported fMRI and reaction time data suggesting that this heuristic might actually be used by "default" because it is easily and quickly executable (Pachur & Hertwig, 2006; Volz et al., 2006). Importantly, however, in no study was the heuristic ever tested extensively against alternative models, which is what is needed to evaluate any model's ability to account for behavior and is essential for determining whether and when people embrace the noncompensatory recognition heuristic and whether and when other noncompensatory or compensatory strategies are at play.

In the research outlined in Chapter 4, I tested the recognition heuristic for the first time against alternative models, including more complex compensatory and noncompensatory ones that operate on knowledge cues, or on objects' retrieval fluency. In response to Dougherty et al.'s (2008) critique, I also tested generalizations of the heuristic to situations in which all objects are recognized, and to situations with multiple objects. In doing so, I essentially examined how the recognition heuristic allows people to form consideration sets (Alba & Chattopadhyay, 1985; Hauser & Wernerfelt, 1990; Howard & Sheth, 1969), that is, how they single out objects from a multitude that are worth further information search—a key problem in the marketing literature. In addition, I addressed the problem of strategy selection, investigating the conditions under which a possible default of using this heuristic might be overruled. These conditions hinge on the retrieval of knowledge about the source of recognition (e.g., about the context in which one has heard of a car brand), knowledge about cues in conflict with recognition (e.g., a consumer protection report suggesting that a recognized car brand has chronic mechanical problems), objects' retrieval fluency, and other information indicating whether recognition is predictive of the criterion to be inferred, that is, whether relying on the recognition heuristic is ecologically rational. For instance, a person may recognize a car brand's name but recall that these cars are known for chronic mechanical problems, indicating that recognition may not be predictive of this particular brand's quality and should not be relied on. Similarly, a person may recognize a car brand only because she has read the name in this dissertation. If she believes that mentions of brand names in dissertations do not reflect the quality of the brands, then she may not trust her recognition of this brand name when making inferences about its quality. To give a final example, it may take a person much time to judge a brand as recognized, such that she lacks confidence about her recognition judgment. Thus, she may not rely on the recognition heuristic when making inferences about that brand. This can be ecologically rational, because recognition times can reflect the predictive accuracy of recognition for making inferences.

In eight model comparisons and six experiments with a total of over 540 participants, the recognition heuristic predicted people's inferences—including potential voters' forecasts of three German political elections—better than six more complex compensatory and noncompensatory models operating on knowledge cues or retrieval fluency (Experiments 11–16). Showing how the heuristic may be used to form consideration sets, voters' inferences about up to 30 candidates and parties in political elections were consistent with the proposed generalizations of the recognition heuristic (Experiments 12, 13, 15). Reaction time data as well as choice data in a signal-to-response paradigm provided further evidence to suggest that the heuristic is used by default (Experiments 12, 14, 16). For instance, under time pressure the recognition heuristic seems to be easier and faster to execute than weighted-fluency, a strategy integrating objects' retrieval fluency and recognition into judgments (Experiment 16). At the same time, these experiments indicated that the retrieval of conflicting cues (Experiment 11), knowledge about the source of recognition (Experiment 12), and other information might lead to ecologically rational suspensions of a default of using the recognition heuristic. To illustrate, in Experiments 13, 14, and two reanalyses of Experiments 1 and 2 (see General Discussion in Chapter 4), I found that the recognition heuristic is less likely to yield accurate inferences when objects' retrieval fluency is low, that is, when it is hardest to use the heuristic. As it turns out, people seem in fact to be less likely to embrace the heuristic in such situations. In short, the behavioral data from Chapter 4 supported the thesis proposed in Chapter 3: The interplay between memory and the environment may result in cognitive niches of heuristics that can make it easier and faster to use a heuristic when applying it is also most likely to result in accurate inferences, aiding strategy selection in this way. At the close of this dissertation, I would like to comment on a methodological theme of the empirical work reported here, namely, why people's use of heuristics should be studied comparatively, guided by theories of strategy selection.

Why Is It Important to Study Heuristics Comparatively, Guided by Theories of Strategy Selection?

In Chapters 3 and 4, I conducted a number of comparative tests of formal models, pitting different heuristics against each other. There are many reasons why such modeling efforts can be beneficial for theory development (see e.g., Fum et al., 2007; Gigerenzer & Brighton, 2008; Hintzman, 1991; Jacobs & Grainger, 1994; Sedlmeier & Renkewitz, 2007). Rather than giving an exhaustive summary of the arguments, in what follows, I consider only four closely interrelated points that are particularly important for the development of the theory of the adaptive toolbox. I will highlight that comparative model tests (i) lead to the identification of better models of behavior, (ii) provide a yardstick for model evaluation, (iii) should be accompanied by a theory of strategy selection, and (iv) can inform the building of theories of strategy selection.

First, following Gigerenzer and Brighton (2008), research on heuristics should not be about testing just one model in isolation, proclaiming whether it fits the data or not, as has been done with the recognition heuristic on numerous occasions (e.g., Bröder & Eichler, 2006; B. R. Newell & Fernandez, 2006; B. R. Newell & Shanks, 2004; Pohl, 2006; Richter & Späth, 2006). Rather, research should be about identifying better models of behavior than those that already exist, aiding scientific progress in developing the theory of the adaptive toolbox by building better theories of heuristics. For instance, the comparative model tests reported in Chapter 4 could have shown that weighted-fluency is a better model than the recognition heuristic, which in turn would have added a new model to the adaptive toolbox. As it turns out, the recognition heuristic is the better model—at least in a majority of participants, but notably not in all. On a related note, assessments of people's reliance on different heuristics have really progressed, as research has shifted from asking questions such as whether people use *one* heuristic in *all* situations, to testing heuristics comparatively, examining *when* a given heuristic might be applied (see Bröder, in press; Gigerenzer & Brighton, 2008). To illustrate, Bröder (2000) started out by asking the question whether all people use the take-the-best heuristic in probabilistic inferences. But, as he has later pointed out, "hypothesis rejections at the group level may throw out the baby with the bath water if individual strategy differences are not taken into account" (Pachur et al., 2008, p. 204), and in fact, almost all studies on take-the-best suggested that varying proportions of participants rely on this heuristic, depending, for instance, on the characteristics of the decision task (e.g., Bröder & Gaissmaier, 2007; Rieskamp & Hoffrage, 1999). In the light of such findings, the focus of research on take-the-best has shifted toward explorations of variables that might guide strategy use (e.g., Bröder, 2003; Bröder & Schiffer, 2003, 2006), and, as discussed in Chapter 3, much progress has been made recently in understanding strategy selection in heuristics such as take-the-best, for instance, in terms of reinforcement learning (e.g., Rieskamp & Otto, 2006). This dissertation thesis is meant to complement these advances in theory development.

Second, formal model comparisons establish yardsticks for evaluating the descriptive adequacy of competing theories, with the formal instantiations of the theories being each other's benchmarks in theory evaluation. When just one model is tested, a seemingly large discrepancy between the model's predictions and the observed data might lead a researcher to reject that model. With a comparison, in turn, the researcher could learn that all models suffer, enabling him or her to find out which model suffers least. Sometimes it might actually be theoretically interesting sources of variation, such as memory variables, that affect all models. This point is illustrated by the set of model comparisons reported in Chapter 4, which also show what dramatically different conclusions one can make from experimental results, depending on whether alternative models are formally specified and tested or just verbally sketched without proper comparative testing. As discussed in detail in Chapter 4, previous findings that people do not always make decisions consistent with the recognition heuristic not only raised doubts about the adequacy of this heuristic as a model of behavior, but they were also used to propose that people rely on compensatory strategies instead (e.g., Richter & Späth, 2006). Yet, no study—including one I coauthored (Pachur et al., 2008)—tested a corresponding alternative model against the heuristic. Instead, the authors of previous work only provided verbally formulated alternative hypotheses of how people could make their decisions if they did not use the recognition heuristic. While I was able to replicate several of the previous findings, namely, that the heuristic does not always predict people's decisions, I also showed that for most people, it predicted behavior better than each of six alternative models that implemented some of the verbal alternative hypotheses. In doing so, I provided evidence to suggest that memory variables such as the strength of the recognition signal—that is, objects' retrieval fluencies—are responsible for systematic variations in the frequency of inferences consistent with the recognition heuristic, pointing to mechanisms of strategy selection rather than to shortcomings in the descriptive adequacy of the heuristic.

Third, if one assumes that people select from a repertoire of strategies, it is important that comparative model tests come accompanied with theories of strategy selection (see also Cooper, 2000; Feeney, 2000; Gigerenzer et al., 2008; Luce, 2000; Payne et al., 1993; Rieskamp & Otto, 2006; Wallin & Gärdenfors, 2000). As my findings for the recognition heuristic highlight in Chapter 4, unless one has a theory of strategy selection, it is problematic to reject a model of a decision strategy simply because it does not predict behavior in a certain situation. In Chapter 4, I explained that there can be two reasons why a decision strategy does not predict behavior. One is that the strategy is not used because people (or the corresponding selection mechanisms) *choose* not to use it in a particular situation; a completely different reason is that the decision strategy per se is not a good model of behavior. This point can also be illustrated with my findings about the applicability of the fluency heuristic in Chapter 3. Before the theory of strategy selection by different cognitive niches was developed, one could have reasonably assumed that this heuristic is equally applicable in all situations, that is, when no knowledge is available, when knowledge can be retrieved about just one object in a pair, and when knowledge is available about both objects (see Hertwig et al., 2008, who did not distinguish between these situations). Comparative model tests, in which these situations are not examined separately, would have shown that knowledge-based strategies predict people's decisions systematically better than the fluency heuristic. Yet, it would have been a mistake to conclude from such model comparisons that the fluency heuristic is not a good model of behavior: As I have shown, the fluency heuristic predicts people's decisions best when knowledge is sparse or unavailable, representing an instance of strategy selection, and this prediction can be made from models of memory, such as the ACT-R memory model elaborated in Chapter 3.[42] In fact, I would like to stress that such predictions about strategy selection require a detailed model of memory. As has been pointed out in Chapter 3, when I first verbally formulated the niche hypotheses, I realized that it was impossible to informally reason my way through how the distributional characteristics of recognition, knowledge, retrieval time, and the frequency of objects in the world might relate, all of which shape in one way or another the cognitive niche of the fluency heuristic. To make a more general point: If one assumes that the successful completion of a decision task is the product of many different processes, for instance, entailing vision, memory, and motor responses in addition to decision making, then it is important that these processes be included in models of people's behavior in that task (see A. Newell, 1973b, 1990; Schooler & Hertwig, 2005). Here, detailed theories of cognition, such as the ACT-R architecture, impose precise theoretical constraints on how to model these processes, facilitating precise predictions.[43] In my view, the memory model developed in Chapter 3 represents only a first, moderate step along these lines—future research should implement all known heuristics in ACT-R (or another overarching architecture), furthering progress in uncovering how the heuristics' cognitive niches aid strategy selection.

Fourth, if one assumes that people select from a repertoire of strategies, it is useful to examine the descriptive adequacy of a given heuristic in comparison to that of others, because it is possible that several strategies might *equally* be able to produce adaptive behavior, for instance, when there is considerable overlap between the cognitive niches of different heuristics. To illustrate, in certain situations, take-the-best, tallying heuristics, and other knowledge-based strategies could result in equally accurate, effortless, and fast decisions so that it matters little which strategy a decision maker

[42] Admittedly, theories assuming a repertoire of strategies face a dilemma: If one finds that a strategy does not predict behavior where it should predict behavior according to the chosen strategy selection model, then it can be hard to distinguish which model needs improvement: the strategy selection model, the decision strategy model, or both.

[43] On a related note, the fast and frugal heuristics framework has in fact been criticized for not specifying the memory processes on which heuristics operate: According to Dougherty et al. (2008), a lack of specification can result in models of heuristics that are incompatible with the functioning of the underlying processes. I view the experiments and simulation studies reported in Chapters 3 and 4 as a response to this critique, showing how models of memory, heuristics, and strategy selection interplay (see also Gigerenzer et al., 2008, for another response).

employs—applying any of them would yield ecologically rational decisions. Possibly, in such situations of strong overlap between cognitive niches, one might find the greatest variability in people's use of different decision strategies, because then the mechanisms that could otherwise systematically push strategy selection in the majority of individuals toward one strategy or another might not be at work. It might be worth speculating whether in such situations of strongly overlapping cognitive niches, social practices, individual preferences, habits, or even personality dispositions channel people's choice of different decision strategies, giving rise to large individual differences in strategy use.

In short, the comparative study of heuristics can aid in identifying better models of behavior than those that already exist, for instance, by establishing criteria for evaluating the descriptive adequacy of competing models. Guided by theories of strategy selection, this study can enhance our understanding of the ecological rationality of a decision maker's strategy choices.

Conclusion: Ecologically Rational Strategy Selection

Within the fast and frugal heuristics framework, much research has focused on what is termed the strategies' ecological rationality, that is, whether and how different strategies can exploit environmental structure (e.g., Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2007; Katsikopoulos & Martignon, 2006b; Martignon & Hoffrage, 1999). Structure has often been characterized in terms of the objective properties of environments, such as cue validities and cue intercorrelations, rather than how these properties are represented by the mind. In this dissertation, I have pushed this emphasis on the interplay between the cognitive system and the environment further, arguing that one should aim to build formal models of how the environment is *represented* by the cognitive system, and by memory in particular. That is, in my view, it is not so much the objective environmental structure per se that is important for decision-making behavior, but how this structure is embodied by the memory system. As I have argued, it is these mental representations of the environment, coupled with the cognitive capacities of the organism that define the cognitive niche of a strategy.

In short, with this dissertation thesis I hope to contribute to understanding the ecological rationality of the heuristics in the adaptive toolbox, by modeling memory retrieval based on environmental data (e.g., Anderson & Schooler, 1991, 2000; Burgess, & Lund, 1997; Griffiths et al., 2007; Landauer & Dumais, 1997; Lund & Burgess, 1996; Schooler & Anderson, 1997), and by integrating models of memory and decision strategies (e.g., Dougherty et al., 1999; Gray et al., 2006;

Schooler & Hertwig, 2005). In doing so, I build on many researchers' work, following a long research tradition of studying how the cognitive system nestles into the structure of the environment (e.g., Anderson, 1990; Brunswik, 1943, 1955; Gibson, 1979; Gigerenzer et al., 1999; Simon, 1956, 1990). One of these researchers, Herbert A. Simon, once argued that memory acts as an extension of the environment in which human thought processes take place. This dissertation has spun this thread of thought. I would therefore like to conclude with Simon's (1996) own words, taken from his book *The Sciences of the Artificial*:

> We can think of the memory as a large encyclopaedia or library, the information stored by topics (nodes), liberally cross-referenced (associational links), and with an elaborate index (recognition capability) that gives direct access through multiple entries to the topics. Long-term memory operates like a second environment, parallel to the environment sensed through eyes and ears, through which the problem solver can search and to whose contents he can respond. (p. 88)

References

Adelman, J. S., & Brown, G. D. A. (2008). Modeling lexical decision: The form of frequency and diversity effects. *Psychological Review, 115,* 214–227.

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrox & F. Caski (Eds.), *Second international symposium on information theory* (pp. 267–281). Budapest, Hungary: Akademiai Kiado.

Alba, J. W., & Chattopadhyay, A. (1985). Effects of context and part-category cues on recall of competing brands. *Journal of Marketing Research, 22,* 340–349.

Alter, A. L., & Oppenheimer, D. M. (2008). Effects of fluency on psychological distance and mental construal (or why New York is a large city, but New York is a civilized jungle). *Psychological Science, 19,* 161–167.

Anderson, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin and Review, 8,* 629–647.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036–1060.

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language, 38,* 341–380.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston and Sons.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review, 96,* 703–719.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408.

Anderson, J. R., & Schooler, L. J. (2000). The adaptive nature of memory. In E. Tulving & F. I. M. Craik (Eds.), *Handbook of memory* (pp. 557–570). New York: Oxford University Press.

Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, R. D. Lucek, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 1, pp. 243–293). San Francisco: Freeman.

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition, 32,* 1369–1378.

Bartels, L. M. (2000). Partisanship and voting behavior, 1952–1996. *American Journal of Political Science, 44,* 35–50.

Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review, 3,* 439–449.

Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General, 121,* 446–458.

Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 331,* 107–129.

Berretty, P. M., Todd, P. M., & Martignon, L. (1999). Categorization by elimination. Using few cues to choose. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 235–254). New York: Oxford University Press.

Bishop, M. (2006). Fast and frugal heuristics. *Philosophy Compass, 1,* 201–223.

Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 59–72). New York: Oxford University Press.

Box, G. E. P. (1979). Robustness in the strategy of scientific model-building. In R. L. Launer & G. N. Wilkinson (Eds.), *Robustness in statistics* (pp. 201–236). New York: Academic Press.

Brady, H. E., & Sniderman, P. M. (1985). Attitude attribution: A group basis for political reasoning. *American Political Science Review, 79,* 1061–1078.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review, 113,* 409–432.

Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & B. Wray (Eds.), *Between a rock and a hard place: Cognitive science principles meet AI-hard problems. Papers from the AAAI Spring Symposium* (AAAI Tech. Rep. No. SS-06-03, pp. 17–22). Menlo Park, CA: AAAI Press.

Bröder, A. (2000). Assessing the empirical validity of the "take-the-best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1332–1346.

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning Memory, and Cognition, 29,* 611–625.

Bröder, A. (in press). The quest for take the best: Insights and outlooks from experimental research. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, *Ecological rationality: Intelligence in the world.* New York: Oxford University Press.

Bröder, A., & Eichler, A. (2006). The use of recognition information and additional cues in inferences from memory. *Acta Psychologica, 121,* 275–284.

Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multi-attribute decisions. *Psychonomic Bulletin and Review, 14,* 895–900.

Bröder, A., & Newell, B. R. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox. *Judgment and Decision Making, 3,* 205–214.

Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General, 132,* 227–293.

Bröder, A., & Schiffer, S. (2006). Stimulus format and working memory in fast and frugal strategy selection. *Journal of Behavioral Decision Making, 19,* 361–380.

Browne, M. W. (2000). Cross-validation methods. *Journal of Mathematical Psychology, 44,* 108–132.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, *64*, 123–152.

Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological Review, 50,* 255–272.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review, 62,* 193–217.

Burgess, C., & Lund, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12,* 177–210.

Burnham, K., & Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach* (2nd ed.). New York: Springer.

Busemeyer, J. R. (1993). Violations of the speed-accuracy tradeoff relation: Decreases in decision accuracy with increases in decision time. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 181–193). New York: Plenum Press.

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General, 121,* 177–184.

Campbell, A., Converse, P., Miller, W., & Stokes, D. (1960). *The American voter*. Chicago: University of Chicago Press.

Christensen-Szalanski, J. J. (1978). Problem solving strategies: A selection mechanism, some implications, and some data. *Organizational Behavior & Human Performance, 22,* 307–323.

Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society, 362,* 1585–1599.

Coates, S. L., Butler, L. T., & Berry, D. C. (2004). Implicit memory: A prime example for brand consideration and choice. *Applied Cognitive Psychology, 18,* 1195–1211.

Coates, S. L., Butler, L. T., & Berry, D. C. (2006). Implicit memory and consumer choice: The mediating role of brand familiarity. *Applied Cognitive Psychology, 20,* 1101–1116.

Cooper, R. (2000). Simple heuristics could make us smart: But which heuristics do we apply when? (Open peer commentary). *Behavioral and Brain Sciences, 23,* 746.

Coyle, T. R., Read, L. E., Gaultney, J. F., & Bjorklund, D. F. (1998). Giftedness and variability in strategy processing on a multitrial memory task. Evidence for stability in gifted cognition. *Learning & Individual Differences, 10,* 273–290.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.

Davis, J. N., & Todd, P. M. (1999). Parental investment by simple decision rules. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 309–324). New York: Oxford University Press.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34,* 571–582.

Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin, 81,* 95–106.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin, 130,* 959–988.

Dougherty, M. R. P. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General, 130,* 579–599.

Dougherty, M. R. P., Franco-Watkins, A. N., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review, 115,* 199–213.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). Minerva-DM: A memory processes model for judgments of likelihood. *Psychological Review, 106*, 180–209.

Dror, I. E., Busemeyer, J. R., & Basola, B. (1999). Decision making under time pressure: An independent test of sequential sampling models. *Memory and Cognition, 27,* 713–725.

Dutton, J. M., & Starbuck, W. H. (Eds.). (1971). *Computer simulation of human behavior: A history of an intellectual technology*. New York: Wiley.

Ebbinghaus, H. (1964). *Memory: A contribution to experimental psychology*. Mineola, NY: Dover. (Original work published 1885)

Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 51, 380–417.

Einhorn, H. J. (1970). The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin, 73,* 221–230.

Einhorn, H. J., & Hogarth, R. M. (1975). Unit weighting schemes for decision making. *Organizational Behavior and Human Performance, 13,* 171–192.

Einhorn, H. J., &, Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgement and choice. *Annual Review of Psychology, 32,* 53–88.

Eisenberg, P., & Becker, C. A. (1982). Semantic context effects in visual word recognition, sentence processing, and reading: Evidence for semantic strategies. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 739–756.

Erev, I., & Roth, A. E. (2001). Simple reinforcement learning models and reciprocation in the prisoner's dilemma game. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 215–231). Cambridge, MA: MIT Press.

Fechner, G. T. (1966). *Elements of psychophysics*. (H. E. Adler, Trans.). New York: Holt, Rinehart and Winston. (Original work published in 1860)

Feeney, A. (2000). Simple heuristics: From one infinite regress to another? (Open peer commentary). *Behavioral and Brain Sciences, 23,* 749–750.

Fischer, J. E., Steiner, F., Zucol, F., Berger, C., Martignon, L., Bossart, W., Altwegg, M., & Nadal, D. (2002). Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia. *Archives of Pediatrics and Adolescent Medicine, 156,* 1005–1008.

Fishburn, P. C. (1974). Lexicographic orders, utilities and decision rules: A survey. *Management Science*, *20*, 1442–1471.

Fishburn, P. C. (1980). Lexicographic additive differences. *Journal of Mathematical Psychology, 21,* 191–218.

Ford, J. K., Schmitt, N., Schechtman, S. L., Hults, B. H., & Doherty, M. L. (1989). Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Decision Processes, 43,* 75–117.

Forster, M. R. (2000). Key concepts in model selection: Performance and generalizability. *Journal of Mathematical Psychology, 44,* 205–231.

Fraisse, P. (1984). Perception and estimation of time. *Annual Review of Psychology, 35,* 1–36.

Frosch, C. A., Beaman, C. P., & McCloy, R. (2007). A little learning is a dangerous thing: An experimental demonstration of recognition-driven inference. *The Quarterly Journal of Experimental Psychology, 60,* 1329–1336.

Fu, W.-T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General, 135,* 184–206.

Fum, D., Del Missier, F., & Stocco, A. (2007). The cognitive modeling of human behavior: Why a model is (sometimes) better than 10,000 words. *Cognitive Systems Research, 8,* 135–142.

Gaissmaier, W. (2007). *The mnemonic decision maker: How search in memory shapes decision making.* Unpublished doctoral dissertation, Freie Universität Berlin, Germany.

Gaissmaier, W., Schooler, L. J., & Mata, R. (2008). An ecological perspective to cognitive limits: Modeling environment-mind interactions with ACT-R. *Judgment and Decision Making, 3,* 278–291.

Gaissmaier, W., Schooler, L. J., & Rieskamp, J. (2006). Simple predictions fuelled by capacity limitations: When are they successful? *Journal of Experimental Psychology*: *Learning, Memory, and Cognition, 32,* 966–982.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Gibson, J. J. (2001). Survival in a world of probable objects. In K. R. Hammond & T. R. Stewart (Eds.), *The essential Brunswik: Beginnings, explications, applications* (pp. 244–246). New York: Oxford University Press. (Original work published in 1957)

Gigerenzer, G. (1982). Der eindiminsionale Wähler. *Zeitschrift für Sozialpsychologie, 13,* 217–236.

Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review, 98,* 254–267.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103,* 592–596.

Gigerenzer, G. (1998). Surrogates for theories. *Theory & Psychology, 8,* 195–204.

Gigerenzer, G. (2001). Ideas in exile: The struggles of an upright man. In K. R. Hammond & T. R. Stewart (Eds.), *The essential Brunswik: Beginnings, explications, applications* (pp. 445–453). New York: Oxford University Press.

Gigerenzer, G. (2002). *Calculated risks: How to know when numbers deceive you.* New York: Simon & Schuster.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious.* New York: Viking.

Gigerenzer, G., & Brighton, H. (2008). *Homo heuristicus: Why biased minds make better inferences.* Manuscript submitted for publication.

Gigerenzer, G., & Edwards, A. (2003). Simple tools for understanding risks: From innumeracy to insight. *British Medical Journal, 327,* 741–744.

Gigerenzer, G., & Engel, C. (Eds.). (2006). *Heuristics and the law*. Cambridge, MA: MIT Press.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 104,* 650–669.

Gigerenzer, G., & Goldstein, D. G. (1999). Betting on one good reason. The take the best heuristic. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 75–95). New York: Oxford University Press.

Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, & Thomas (2008). *Psychological Review, 115,* 230–239.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98,* 506–528.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Ginossar, Z., & Trope, Y. (1987). Problem solving in judgment under uncertainty. *Journal of Personality and Social Psychology, 52,* 464–474.

Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 37–48). New York: Oxford University Press.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review, 109,* 75–90.

Gonzalez, R., & Nelson, T. O. (1996). Measuring ordinal association in situations that contain tied scores. *Psychological Bulletin, 119,* 159–165.

Goodman, L. A., & Kruskal, W. H. (1954). Measures of association for cross classifications. *Journal of the American Statistical Association, 49,* 732–764.

Grainger, J., & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review, 103,* 518–565.

Gray, W. D., & Fu, W.-T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science, 28,* 359–382.

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review, 113,* 461–482.

Green, L., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *The Journal of Family Practice, 45,* 219–226.

Greeno, J. G. (1994). Gibson's affordances. *Psychological Review, 101,* 336–342.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review, 114,* 211–244.

Gronlund, S. D., & Ratcliff, R. (1989). The time course of item and associative information: Implications for global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 846–858.

Grünwald, P. D. (2007). *The minimum description length principle*. Cambridge, MA: MIT Press.

Hammond, K. R. (1996). *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. Oxford: Oxford University Press.

Haruvy, E., & Stahl, D. O. (2007). Equilibrium selection and bounded rationality in symmetric normal-form games. *Journal of Economic Behavior and Organization, 62,* 98–119.

Haruvy, E., Stahl, D. O., & Wilson, P. W. (2001). Modeling and testing for heterogeneity in observed strategic behavior. *Review of Economics and Statistics, 83,* 146–157.

Hauser, J. R., & Wernerfelt, B. (1990). An evaluation cost model of consideration sets. *Journal of Consumer Research, 16,* 393–408.

Helversen, B. von, & Rieskamp, J. (2008). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General, 137,* 73–96.

Hertwig, R., Gigerenzer, G., & Hoffrage, U. (1997). The reiteration effect in hindsight bias. *Psychological Review, 104,* 194–202.

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 1191–1206.

Hertwig, R., Hoffrage, U., & Martignon, L. (1999). Quick estimation: Letting the environment do the work. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 209–234). New York: Oxford University Press.

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 621–642.

Hertwig, R., & Todd, P. M. (2003). More is not always better: The benefits of cognitive limits. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making* (pp. 213–231). Chichester, UK: Wiley.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford Press.

Hintzman, D. L. (1988). Judgment of frequency and recognition memory in a multiple trace memory model. *Psychological Review, 95,* 528–551.

Hintzman, D. L. (1991). Why are formal models useful in psychology? In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 39–56). Hillsdale, NJ: Erlbaum.

Hintzman, D. L., & Curran, T. (1994). Retrieval dynamics of recognition and frequency judgments: Evidence for separate processes of familiarity and recall. *Journal of Memory and Language, 33,* 1–18.

Hoffrage, U., & Hertwig, R. (2006). Which world should be represented in representative design. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (pp. 381–408). New York: Cambridge University Press.

Hoffrage, U., Lindsey, S., Hertwig, R., & Gigerenzer, G. (2000). Communication of statistical information. *Science, 290,* 2261–2262.

Hogarth, R. M. (1987). *Judgment and choice: The psychology of decision*. Chichester, UK: Wiley.

Hogarth, R. M., & Karelaia, N. (2005). Simple models for multiattribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. *Management Science, 51,* 1860–1872.

Hogarth, R. M., & Karelaia, N. (2006). Regions of rationality: Maps for bounded agents. *Decision Analysis*, *3*, 124–144.

Hogarth, R. M., & Karelaia, N. (2007). Heuristics and linear models of judgment: Matching rules and environments. *Psychological Review, 114,* 733–758.

Howard, J. A., & Sheth, J. N. (1969). *The theory of buyer behavior.* New York: Wiley.

Hume, D. L. (1978). *A treatise of human nature*. Oxford: Clarendon Press. (Original work published in 1740)

Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes, 69,* 87–124.

Jackman, S., & Sniderman, P. M. (2002). The institutional organization of choice spaces: A political conception of political psychology. In K. Monroe (Ed.), *Political psychology* (pp. 209–224). Mahwah, NJ: Erlbaum.

Jacobs, A. M., & Grainger, J. (1994). Models of visual word recognition. Sampling the state of the art. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 1311–1334.

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30,* 513–541.

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 18, pp. 1–47). New York: Academic Press.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110,* 306–340.

Jacoby, L. L., Kelley, C. M., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology, 56,* 326–338.

Jacoby, L. L., Woloshyn, V., & Kelly, C. M. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General, 118,* 115–125.

Johnson, M. K., Hashtroudi, S., & Lindsay, S. D. (1993). Source monitoring. *Psychological Bulletin, 114*, 3–28.

Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science, 26,* 563–607.

Kahneman, D. (2003). A perspective on judgment and choice. Mapping bounded rationality. *American Psychologist, 58,* 697–720.

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). New York: Cambridge University Press.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology, 3,* 430–454.

Katsikopoulos, K. V., & Martignon, L. (2006a). Naïve heuristics for paired comparisons: Some results on their relative accuracy. *Journal of Mathematical Psychology, 50,* 488–494.

Katsikopoulos, K. V., & Martignon, L. (2006b). On the accuracy of lexicographic strategies for pair comparison. *Journal of Mathematical Psychology, 50,* 116–122.

Keller, F., & Lapata, M. (2003). Using the web to obtain frequencies for unseen bigrams. *Computational Linguistics, 29,* 459–484.

Kelley, S., & Mirer, T. W. (1974). The simple act of voting. *American Political Science Review, 68,* 572–591.

Kohli, R., & Jedidi, K. (2007). Representation and inference of lexicographic preference models and their variants. *Marketing Science, 26,* 380–399.

Koriat, A. (1993). How do we know that we know? The accessibility model of feeling of knowing. *Psychological Review, 100,* 609–639.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104,* 211–240.

Lewandowsky, S. (1993). The rewards and hazards of computer simulations. *Psychological Science, 4,* 236–243.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527.

Lopes, L. L. (1991). The rhetoric of irrationality. *Theory & Psychology, 1,* 65–82.

Lopes, L. L. (1992). Three misleading assumptions in the customary rhetoric of the bias literature. *Theory & Psychology, 2,* 231–236.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology, 31,* 168–217.

Luce, D. R. (2000). Fast, frugal, and surprisingly accurate heuristics. (Open peer commentary). *Behavioral and Brain Sciences, 23,* 757–758.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, and Computers, 28,* 203–208.

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review, 87,* 252–271.

Marewski, J. N. (2005). *Is ignorance useful and used? Applying the recognition heuristic to political elections.* Diploma thesis, Freie Universität Berlin, Germany.

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2008). *Complex problems do not require complex cognition.* Manuscript submitted for publication.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., & Gigerenzer, G. (2008). *Fast and frugal election forecasts.* Manuscript in preparation, Max Planck Institute for Human Development, Berlin.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2008). *Models of recognition-based multi-alternative inference.* Manuscript submitted for publication.

Marewski, J. N., Galesic, M., & Gigerenzer, G. (in press). Fast and frugal media choices. In T. Hartmann (Ed.), *Media choice: A theoretical and empirical overview*. New York: Routledge.

Marewski, J. N., & Olsson, H. (in press). Beyond the null ritual: Formal modeling of psychological processes. *Zeitschrift für Psychologie / Journal of Psychology*.

Marewski, J. N., & Schooler, L. J. (2008a). *How memory aids strategy selection.* Manuscript submitted for publication.

Marewski, J. N., & Schooler, L. J. (2008b). *Using ACT-R to model how people make inferential judgments from recognition information.* Unpublished manuscript, Max Planck Institute for Human Development, Berlin.

Marr, D. (1982). *Vision*. New York: Freeman.

Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 119–140). New York: Oxford University Press.

Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology and Aging, 22,* 796–810.

McAllister, D. W., Mitchell, T. R., & Beach, L. R. (1979). The contingency model for the selection of decision strategies: An empirical test of the effects of significance, accountability, and reversibility. *Organizational Behavior and Human Performance, 24,* 228–244.

McElree, B., Dolan, P. O., & Jacoby, L. L. (1999). Isolating the contributions of familiarity and source information to item recognition: A time course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 563–582.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3–65.

Myung, I. J., Balasubramanian, V., & Pitt, M. A. (2000). Counting probability distributions: Differential geometry and model selection. *Proceedings of the National Academy of Sciences of the United States of America, 97,* 11170–11175.

Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review, 4,* 79–95.

Navarro, D. J., Pitt, M. A., & Myung, I. J. (2004). Assessing the distinguishability of models and the informativeness of data. *Cognitive Psychology, 49,* 47–84.

Nellen, S. (2003). The use of the "take-the-best" heuristic under different conditions, modelled with ACT-R. In F. Detje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the fifth international conference on cognitive modelling* (pp. 171–176). Bamberg, Germany: Universitätsverlag Bamberg.

Newell, A. (1973a). Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing* (pp. 463–526). New York: Academic Press.

Newell, A. (1973b). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–310). New York: Academic Press.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Newell, B. R., & Fernandez, D. (2006). On the binary quality of recognition and the inconsequentiality of further knowledge: Two critical tests of the recognition heuristic. *Journal of Behavioral Decision Making, 19,* 333–346.

Newell, B. R., Rakow. T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of "success". *Journal of Behavioral Decision Making, 17,* 117–137.

Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 53–65.

Newell, B. R., & Shanks, D. R. (2004). On the role of recognition in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 923–935.

Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast and frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes, 91,* 82–96.

Nilsson, H., Olsson, H., & Juslin, P. (2005). The cognitive substrate of subjective probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 600–620.

Oaksford, M., & Chater, N. (Eds.). (1998). *Rational models of cognition*. Oxford: Oxford University Press.

Oppenheimer, D. M. (2003). Not so fast! (and not so frugal!): Rethinking the recognition heuristic. *Cognition, 90,* B1–B9.

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Science*, *12*, 237–241.

Pachur, T., & Biele, G. (2007). Forecasting from ignorance: The use and usefulness of recognition in lay predictions of sports events. *Acta Psychologica, 125,* 99–116.

Pachur, T., Bröder, A., & Marewski, J. N. (2008). The recognition heuristic in memory-based inference: Is recognition a non-compensatory cue? *Journal of Behavioral Decision Making, 21,* 183–210.

Pachur, T., & Hertwig, R. (2006). On the psychology of the recognition heuristic: Retrieval primacy as a key determinant of its use. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 983–1002.

Pachur, T., Todd, P. M., Gigerenzer, G., Schooler, L. J., & Goldstein, D. G. (in press). Is ignorance an adaptive tool? A review of recognition heuristic research. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior & Human Performance, 16,* 366–387.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 534–552.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.

Payne, J. W., Bettman, J. R., & Luce, M. F. (1996). When time is money: Decision behavior under opportunity-cost time pressure. *Organizational Behavior and Human Decision Processes, 66,* 131–152.

Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review, 113,* 57–83.

Pitt, M. A., & Myung, I. J. (2000). When a good fit can be bad. *Trends in Cognitive Sciences, 6,* 421–425.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method for selecting among computational models for cognition. *Psychological Review, 109,* 472–491.

Pleskac, T. J. (2007). A signal detection analysis of the recognition heuristic. *Psychonomic Bulletin & Review, 14,* 379–391.

Pohl, R. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making, 19,* 251–271.

Provost, F. J., & Buchanan, B. G. (1995). Inductive policy. The pragmatics of bias selection. *Machine Learning, 20,* 35–61.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93–134.

Rapoport, A., & Wallsten, T. S. (1972). Individual decision behavior. *Annual Review of Psychology, 21,* 131–176.

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review, 9,* 278–291.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology, 21,* 139–155.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review, 106,* 261–300.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review, 8,* 364–382.

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science, 9,* 45–48.

Reder, L. M. (1987). Strategy selection in question answering. *Cognitive Psychology, 19,* 90–138.

Regenwetter, M., Ho, M.-H., & Tsetlin, I. (2007). Sophisticated approval voting, ignorance priors, and plurality heuristics: A behavioral social choice analysis in a Thurstonian framework. *Psychological Review, 114,* 994–1114.

Reimer, T., & Katsikopoulos, K. V. (2004). The use of recognition in group decision-making. *Cognitive Science, 28,* 1009–1029.

Richter, T., & Späth, P. (2006). Recognition is used as one cue among others in judgment and decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 150–162.

Rieskamp, J. (2006). Perspectives on probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1355–1370.

Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics, and how well can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 141–167). New York: Oxford University Press.

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica, 127,* 258–276.

Rieskamp, J., & Otto, P. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207–236.

Ritter, S., Anderson, J. R., Koedinger, K. R., & Corbett, A. (2007). Cognitive tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review, 14,* 249–255.

Roberts, J. H., & Lattin, J. M. (1991). Development and testing of a model of consideration set composition. *Journal of Marketing Research, 28,* 429–440.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107,* 358–367.

Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. I). Cambridge, MA: MIT Press.

Runeson, S., Juslin, P., & Olsson, H. (2000). Visual perception of dynamic properties: Cue heuristics versus direct-perceptual competence. *Psychological Review, 107,* 525–555.

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors, 48,* 362–380.

Scheibehenne, B., & Bröder, A. (2007). Predicting Wimbledon 2005 tennis results by mere player name recognition. *International Journal of Forecasting, 23,* 415–426.

Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology, 32,* 219–250.

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review, 112,* 610–628.

Schooler, L. J., Shiffrin, R. M., & Raaijmakers, J. G. W. (2001). A Bayesian model of implicit effects in perceptual identification. *Psychological Review, 108,* 257–272.

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics, 6,* 461–464.

Sedlmeier, P., Hertwig, R., & Gigerenzer, G. (1998). Are judgments of the positional frequencies of letters systematically biased due to availability? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 754–770.

Sedlmeier, P., & Renkewitz, F. (2007). *Forschungsmethoden und Statistik in der Psychologie*. Munich: Pearson Education.

Serwe, S., & Frings, C. (2006). Who will win Wimbledon? The recognition heuristic in predicting sports events. *Journal of Behavioral Decision Making, 19,* 321–332.

Shephard, R. N. (2001). Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences, 24,* 581–601.

Siegler, R. S. (1988). Strategy choice procedures and the development of multiplication skill. *Journal of Experimental Psychology: General, 117,* 258–275.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics, 69,* 99–118.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129–138.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology, 41,* 1–19.

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

Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65,* 167–196.

Sniderman, P. M. (2000). Taking sides: A fixed choice theory of political reasoning. In A. Lupia, M. D. McCubbins, & S. L. Popkin (Eds.), *Elements of reason: Cognition, choice, and the bounds of rationality* (pp. 74–84). New York: Cambridge University Press.

Snook, B., & Cullen, R. M. (2006). Recognizing national hockey league greatness with an ignorance-based heuristic. *Canadian Journal of Experimental Psychology, 60,* 33–43.

Stahl, D. O. (1999). Evidence based rule learning in symmetric normal-form games. *International Journal of Game Theory, 28,* 111–130.

Stahl, D. O. (2000). Rule learning in symmetric normal-form games: Theory and evidence. *Games and Economic Behavior, 32,* 105–138.

Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization, 25,* 309–327.

Stevens, S. S. (1957). On the psychophysical law. *Psychological Review, 64,* 153–181.

Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions (with discussion). *Journal of the Royal Statistical Society, Series B, 36,* 111–147.

Stone, M. (1977). Asymptotics for and against cross-validation. *Biometrika, 64,* 29–35.

Stone, W. J., & Abramowitz, A. I. (1983). Winning may not be everything, but it's more than we thought: Presidential party activists in 1980. *American Political Science Review, 77,* 945–956.

Svenson, O. (1979). Process descriptions of decision making. *Organizational Behavior & Human Performance, 23,* 86–112.

Taatgen, N., & Wallach, D. (2002). Whether skill acquisition is rule or instance based is determined by the structure of the task. *Cognitive Science Quarterly, 2,* 163–204.

Taatgen, N., van Rijn, H., & Anderson, J. R. (2007). An integrated theory of prospective time interval estimation: The role of cognition, attention, and learning. *Psychological Review, 114,* 577–598.

Tobii Technology. (2005, June). *Tobii eye tracker & ClearView analysis software. User Manual.* Danderyd, Sweden: Author.

Todd, P. M., & Gigerenzer, G. (2000). Précis of *simple heuristics that make us smart*. *Behavioral and Brain Sciences, 23,* 727–741.

Todd, P. M., & Miller, G. F. (1999). From pride and prejudice to persuasion: Realistic heuristics for mate search. In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 287–308). New York: Oxford University Press.

Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. (2005). Inferences of competence from faces predict election outcomes. *Science, 308,* 1623–1626.

Toth, J. P., & Daniels, K. A. (2002). Effects of prior experience on judgments of normative word frequency: Automatic bias and correction. *Journal of Memory and Language, 46,* 845–874.

Townsend, J. T. (1975). The mind-body problem revisited. In C. Cheng (Ed.), *Philosophical aspects of the mind-body problem* (pp. 200–218). Honolulu, HI: Honolulu University Press.

Tulving, E. (1983). *Elements of episodic memory.* Oxford: Clarendon Press.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior, 5,* 381–391.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79,* 281–299.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5,* 207–232.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124–1130.

Vicente, K. J. (2003). Beyond the lens model and direct perception. Toward a broader ecological psychology. *Ecological Psychology, 15,* 241–267.

Vilalta, R., & Drissi, Y. (2002). A perspective view and survey of meta-learning. *Artificial Intelligence Review, 18,* 77–95.

Volz, K. G., Schooler, L. J., Schubotz, R. I., Raab, M., Gigerenzer, G., & von Cramon, D. Y. (2006). Why you think Milan is larger than Modena: Neural correlates of the recognition heuristic. *Journal of Cognitive Neuroscience, 18,* 1924–1936.

Wagenmakers, E.-J. (2003). How many parameters does it take to fit an elephant? *Journal of Mathematical Psychology, 47,* 580–586.

Wagenmakers, E.-J., & Brown, S. (2007). On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review, 114,* 830–841.

Wagenmakers, E.-J., Zeelenberg, R., Steyvers, M., Shiffrin, R. M., & Raijmakers, J. G. W. (2004). Nonword repetition in lexical decision: Support for two opposing processes. *Quarterly Journal of Experimental Psychology, 57A,* 1191–1210.

Walker, M. (2000). An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *Journal of Artificial Intelligence Research, 12,* 387–416.

Wallin, A., & Gärdenfors, P. (2000). Smart people who make simple heuristics work. (Open peer commentary). *Behavioral and Brain Sciences, 23,* 765.

Wang, X. T. (2008). Decision heuristics as predictors of public choice. *Journal of Behavioral Decision Making, 21,* 77–89.

Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1235–1253.

Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation increases positive affect. *Journal of Personality and Social Psychology, 81,* 989–1000.

Winkielman, P., Knutson, B., Paulus, M. P., & Trujillo, J. T. (2007). Affective influence on decisions: Moving towards the core mechanisms. *Review of General Psychology, 11,* 179–192.

Winkielman, P., Schwarz, N., & Belli, R. F. (1998). The role of ease of retrieval and attribution in memory judgments: Judging your memory as worse despite recalling more events. *Psychological Science, 9,* 124–126.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 189–217). Mahwah, NJ: Erlbaum.

Woodworth, R. (1938). *Experimental psychology*. New York: Holt.

Yee, M., Dahan, E., Hauser, J., & Orlin, J. (2007). Greedoid-based non-compensatory inference. *Marketing Science, 26,* 532–549.

Appendices

The appendices belong to Chapter 3. Together with this chapter, they were submitted for publication to the journal Psychological Review. An extended version of this chapter has received strong encouragement for resubmission to this journal. A copy the original dissertation chapter, including the appendices, is available from Julian Marewski and can be requested at marewski[AT]mpib-berlin[DOT]mpg[DOT]de.

Marewski, J. N. & Schooler, L. J. (2010). Cognitive niches. An ecological model of emergent strategy selection.

Deutsche Zusammenfassung

Ganz gleich, ob es darum geht, welche Aktien man kauft, welchen Film man im Kino sieht oder auf welche Stellenangebote man sich bewirbt: Menschen treffen ihre Entscheidungen oft unter Zeitdruck, mit relativ begrenztem Wissen und unter Rückgriff auf verhältnismäßig geringe Informationsverarbeitungskapazitäten. Die Arbeitsgruppe um Gerd Gigerenzer (Gigerenzer et al., 1999; Todd & Gigerenzer, 2000) vertritt die Überzeugung, dass solche Entscheidungen dennoch sehr erfolgreich getroffen werden können, weil Menschen über ein Repertoire von einfachen, sparsamen und schnellen Entscheidungsstrategien verfügen. Diese Entscheidungsstrategien – oder *Heuristiken* – sind an die Struktur von Entscheidungsumwelten angepasst und greifen auf im Laufe der Evolution entstandene Fähigkeiten zurück (z. B. die Fähigkeit, sich an Vergangenes zu erinnern oder die Fähigkeit zu sehen). Gigerenzer und Kollegen beschreiben das Repertoire einfacher Heuristiken, das dem kognitiven System zur Verfügung steht, mit der Metapher einer Werkzeugkiste. Die gezielte Auswahl von Heuristiken aus dieser Werkzeugkiste erlaubt es Menschen, angemessen auf unterschiedliche Entscheidungssituationen zu reagieren. Sie müssen lediglich jene Heuristik auswählen, die an die entsprechende Entscheidungsumwelt angepasst und somit gut zur Lösung des Entscheidungsproblems geeignet ist, d. h. deren Verwendung in der gegebenen Situation *ökologisch rational* ist („ecological rationality"; siehe z. B. Todd & Gigerenzer, 2000).

In dieser Dissertation habe ich drei eng zusammenhängende Fragestellungen des *Forschungsprogramms der schnellen und sparsamen Heuristiken* („fast and frugal heuristics research program"; Gigerenzer et al., 2008) untersucht: (i) Welche Heuristiken verwenden Menschen? (ii) Wann wird welche Heuristik verwendet? (iii) An welche Umweltstrukturen sind die Heuristiken angepasst, d. h. wann ist die Verwendung einer Heuristik ökologisch rational? Der Schwerpunkt meiner Arbeit lag dabei auf dem *Problem der Strategieselektion*: Wie wird unter den zur Verfügung stehenden Heuristiken ausgewählt, wenn Menschen alle für die Entscheidung relevanten Informationen aus dem Gedächtnis abrufen müssen („inferences from memory"; siehe Gigerenzer & Goldstein, 1996)? Bei der Ausarbeitung eines Lösungsvorschlags habe ich verschiedene psychologische Theorien über menschliche Entscheidungsstrategien, Entscheidungsumwelten, das menschliche Gedächtnis und Strategieselektion herangezogen. So stellen Teile dieser Dissertation Verbindungen zwischen dem Forschungsprogramm der schnellen und sparsamen Heuristiken und anderen ökologischen Ansätzen in der Psychologie her. Hierbei ist insbesondere die kognitive Architektur ACT-R (Anderson et al., 2004; Anderson & Lebiere, 1998) zu nennen, die es mir ermöglichte, ein quantitatives Modell des Zusammenspiels zwischen Gedächtnis und

Entscheidungsumwelten zu entwickeln. Da die Prüfung der vorgeschlagenen Heuristiken eine Durchführung von formalen Modellvergleichen erforderte, enthalten die einleitenden Teile dieser Dissertation auch einen kurzen Überblick über das Thema Modellselektion.

Kapitel 1 führt in das Forschungsprogramm der schnellen und sparsamen Heuristiken sowie die ACT-R Architektur ein. Als erweiterte Einleitung zu meiner empirischen Arbeit gibt Kapitel 2 einen kurzen Überblick über verschiedene Modellselektionskriterien. In den beiden Hauptteilen der Dissertation, den in Kapitel 3 und 4 dargestellten Experimenten und Simulationsstudien, formuliere ich Antworten auf die drei oben genannten Fragen des Forschungsprogramms der schnellen und sparsamen Heuristiken. Ich zeige, welche Heuristiken Menschen in der gedächtnisbasierten Entscheidungsfindung verwenden, wann sie diese anwenden und wie diese an die Struktur der Umwelt angepasst sind. In den folgenden Zeilen finden Sie eine kurze Zusammenfassung dieser empirischen Arbeiten. Ich beginne mit Kapitel 3.

Wie hilft das Gedächtnis bei der Selektion von Strategien?

In der Entscheidungsforschung gehen viele Theorien explizit oder implizit von zwei Determinanten der Strategieauswahl aus. Es handelt sich hierbei um die Kosten der Strategieanwendung (z. B. operationalisiert als Zeit oder Anstrengung und Energie, die aufzuwenden sind, um eine Strategie auszuführen) und die Güte der Entscheidungen, die durch die Anwendung einer Strategie erreicht werden kann (z. B. operationalisiert als die Genauigkeit von Vorhersagen). Eine relativ alte Annahme in der Entscheidungsforschungsliteratur ist dabei, dass Menschen bei der Strategieauswahl zwischen den Kosten und dem Nutzen der Anwendung verschiedener Strategien *abwägen* (siehe Beach & Mitchell, 1978; Christensen-Szalanski, 1978; Payne et al., 1988, 1993). Beispielsweise könnten sie mit Hilfe einer Metastrategie die erwarteten Anwendungskosten der erwarteten Genauigkeit einer Strategie gegenüberstellen. Auch neuere Ideen konzentrieren sich weiterhin auf die Genauigkeit von Strategien oder die Zeit und Anstrengungen, die für ihre Anwendung aufgebracht werden müssen (siehe z. B. Bröder 2003; Bröder & Schiffer, 2006; Mata et al., 2007; B.R. Newell & Shanks, 2003; Pachur & Hertwig, 2006; Payne et al., 1988; Rieskamp & Hoffrage, 1999; Volz et al., 2006).

Als Ergänzung zu diesen Theorien habe ich in Kapitel 3 gezeigt, wie die Funktionsweise des menschlichen kognitiven Systems bei der Auswahl zwischen verschiedenen Strategien behilflich sein kann. So habe ich vorgeschlagen, dass die Heuristiken der adaptiven Werkzeugkiste in unterschiedlichen „kognitiven Nischen“ situiert sind, die durch die Art und Weise entstehen, wie das menschliche Gedächtnis die Informationsstruktur der Entscheidungsumwelten repräsentiert. Beispielsweise rufen Menschen Informationen über Objekte, die sie häufig (oder vor relativ kurzer Zeit) in der Umwelt angetroffen haben, besonders schnell aus dem Gedächtnis ab. Dieses Zusammenspiel von Gedächtnis und Umwelt (i) begrenzt einerseits die Anzahl der Strategien, die zur Lösung eines Entscheidungsproblems verwandt werden können und (ii) erleichtert andererseits den Zugriff auf jene Strategien, deren Verwendung auch tatsächlich gute Ergebnisse liefert. So kann das erwähnte Zusammenwirken von Gedächtnis und Umwelt dazu führen, dass eine Strategie genau dann einfach und schnell ausführbar ist, wenn ihre Anwendung auch gute Vorhersageurteile liefern würde. Es handelt sich hierbei somit um einen Fall ökologisch rationaler Strategieauswahl, der der häufig in der Entscheidungsforschungsliteratur vertretenen Annahme einer Kostennutzenabwägung entgegensteht.

Um meine Thesen zu untermauern, habe ich ein formales ACT-R Modell vorgeschlagen und getestet. Dieses ACT-R Gedächtnismodell erklärt, wie das menschliche Gedächtnis und die Umwelt zusammenwirken. Dabei macht es quantitative Vorhersagen über die Abrufbarkeit von Gedächtnisinhalten in Abhängigkeit von der Informationsstruktur der Umwelt. Basierend auf diesen Gedächtnisinhalten können Entscheidungen getroffen werden.

In 4 Simulationsstudien und 10 Experimenten mit insgesamt 500 Teilnehmern beschäftigte ich mich einerseits mit der Auswahl zwischen Entscheidungsstrategien, die auf Wissen beruhen, d. h. auf den Inhalten abgerufener Informationen, und andererseits mit solchen Strategien, die auf den Charakteristika des Informationsabrufs an sich, wie beispielsweise der Abrufgeschwindigkeit, basieren. Eine sehr einfache Strategie, die sich auf die Geschwindigkeit des Informationsabrufs stützt, ist die *Fluency-Heuristik* („fluency heuristic“; Schooler & Hertwig, 2005). Diese Heuristik erlaubt es, zu entscheiden, welches von zwei Objekten (z. B. Autos) einen höheren Wert hinsichtlich eines gegebenen Kriteriums erzielt (z. B. die Qualität von Autos). Die Fluency-Heuristik setzt darauf, dass das Objekt, das schneller aus dem Gedächtnis abgerufen werden kann, auch durch einen höheren Kriteriumswert gekennzeichnet ist. Wissensbasierte Strategien hingegen treffen solche Entscheidungen basierend auf Attributen von Objekten (z. B. Fakten über Autos, wie Wiederverkaufswerte oder Herstellergarantien). So könnte eine wissensbasierte Strategie darauf setzen, dass das Objekt, das mehr positive Attribute aufweist (z. B. einen höheren Wiederverkaufswert und eine umfangreichere Herstellergarantie hat), auch über einen höheren Kriteriumswert verfügt. In meinen Simulationsstudien konnte ich zeigen, dass die Lokalisation der Strategien in verschiedenen kognitiven Nischen die Auswahl zwischen ihnen vereinfacht.

Es folgt ein kurzer Überblick über meine Simulationsstudien und Experimente. In Simulationsstudie 1 kalibrierte ich die freien Parameter meines ACT-R Gedächtnismodells anhand der Verhaltensdaten meiner Versuchsteilnehmer. Dabei konnte ich aufzeigen, dass mein Modell in der Lage ist, vorherzusagen, welche Objekte Menschen wiedererkennen, über welche Objekte sie Wissen abrufen können und wie die Verteilungen der dazugehörigen Abrufzeiten aussehen. In Simulationsstudie 2 benutzte ich die vom Modell vorhergesagten Daten um zu demonstrieren, wie Unterschiede zwischen den Abrufzeiten zweier Objekte mit der Wahrscheinlichkeit des Abrufens von Wissen über diese Objekte zusammenhängen. Auf diese Art und Weise konnte ich nachweisen, dass das Gedächtnismodell in der Tat unterschiedliche kognitive Nischen für die Fluency-Heuristik und für wissensbasierte Strategien vorhersagt. Demnach hängt die Anwendbarkeit der Fluency-Heuristik von der Fähigkeit einer Person ab zu bestimmen, welches von zwei Objekten schneller aus dem Gedächtnis abgerufen wurde. In meinen Simulationsstudien konnte ich zeigen, dass die Heuristik am ehesten anwendbar ist, wenn wenig oder kein Wissen über die Objekte vorhanden ist. In diesen Situationen sind Unterschiede zwischen den Abrufzeiten zweier Objekte in der Regel größer und damit einfacher zu bemerken, was wiederum die Anwendbarkeit der Fluency-Heuristik erhöht. Umgekehrt ist die Wahrscheinlichkeit, dass die Fluency-Heuristik anwendbar ist, geringer, wenn reichlich Wissen vorhanden ist. In solchen Situationen sind Unterschiede zwischen den Abrufzeiten zweier Objekte eher schwierig zu entdecken. Wissensbasierte Strategien dagegen können nur dann eingesetzt werden, wenn Wissen vorhanden ist. Dementsprechend werden Menschen am ehesten die Fluency-Heuristik einsetzen, wenn sie nicht auf Wissen zurückgreifen können. Anderseits werden sie sich am ehesten auf ihr Wissen verlassen, wenn die Fluency-Heuristik nicht verwendbar ist. In den Simulationsstudien 3 und 4 benutzte ich die von meinem ACT-R Modell vorhergesagten Wiedererkennungs-, Wissens- und Abrufzeitdaten um zu demonstrieren, dass das ACT-R Modell auch vorhersagt, unter welchen Umständen eine Verwendung der Fluency-Heuristik am ehesten zu genauen Entscheidungen führt (z. B. die Qualität von Autos *gut* vorhersagt werden kann). Das ist genau dann der Fall, wenn die Heuristik auch am ehesten anwendbar ist; nämlich wenn wenig Wissen verfügbar ist und die Unterschiede zwischen den Abrufzeiten zweier Objekte groß sind.

Experimente 1 bis 10 lieferten die für die Kalibrierung und Evaluation des ACT-R Modells notwendigen Verhaltensdaten. Experimente 7, 8 und 9 erbrachten zudem Verhaltensdaten, die dazu dienten, weiter zu untersuchen, wann Menschen die Fluency-Heuristik tatsächlich verwenden. Im Einklang mit den in den Simulationsstudien zuvor generierten Vorhersagen des ACT-R Modells zeigte sich hier, dass Menschen die Fluency-Heuristik in der Tat dann am ehesten verwenden, wenn nur wenig Wissen verfügbar ist und Unterschiede zwischen den Abrufzeiten zweier Objekte groß sind, so dass die Heuristik einfach anzuwenden ist und zu schnellen und genauen Entscheidungen führen kann.

Modelle wiedererkennungsbasierter Entscheidungen zwischen multiplen Objekten

Der Schwerpunkt von Kapitel 4 liegt auf der *Rekognitionsheuristik* („recognition heuristic"; Goldstein & Gigerenzer, 2002). Diese ermöglicht es, sich zu entscheiden, welches von zwei Objekten (z. B. Städte) einen höheren Wert hinsichtlich eines gegebenen Kriteriums erzielt (z. B. Einwohnerzahl). In Situationen, in denen die Wiedererkennung von Objekten positiv mit deren Kriteriumswerten korreliert, setzt die Heuristik darauf, dass in einem Vergleich zwischen einem wiedererkannten und einem unbekannten Objekt das wiedererkannte Objekt einen höheren Kriteriumswert aufweist. Die Heuristik basiert solche Entscheidungen dabei auf einem binären Wiedererkennungsurteil und geht von einer nicht-kompensatorischen Nutzung von Wiedererkennungsinformation aus. Nicht-kompensatorisch bedeutet hier, dass eine Person Wissen über Objekte (z. B. Fakten über Städte) ignoriert, wenn die Rekognitionsheuristik verwandt wird.

Unlängst wurde die Heuristik dafür kritisiert, dass sie nur ein Modell für Entscheidungen zwischen einem wiedererkannten und einem unbekannten Objekt sei. Im Prinzip könne daher nicht auf die Heuristik zurückgegriffen werden, wenn alle Objekte einer Objektmenge wiedererkannt würden (Dougherty et al., 2008). Befunde, die zeigen, dass nicht alle Menschen immer Entscheidungen treffen, die den Vorhersagen der Heuristik entsprechen, haben zudem Zweifel an deren Güte als Modell für menschliches Verhalten aufkommen lassen (siehe z. B. Bröder & Eichler, 2006; B. R. Newell & Shanks, 2004; Oppenheimer, 2003; Pohl, 2006). Beispielsweise führten Richter and Späth (2006) eine Serie von Studien zur Verwendung der Rekognitionsheuristik durch. Sie stellten fest, dass weniger Entscheidungen im Einklang mit den Vorhersagen der Rekognitionsheuristik getroffen werden, wenn zusätzliches Wissen und Wiedererkennungsinformation einander widersprechen (z. B. wenn Fakten über wiedererkannte Städte andeuten, dass es sich um kleine Städte handelt). Basierend auf solchen Beobachtungen schlussfolgerten diese Autoren, dass es keinen Beleg für die nicht-kompensatorische Nutzung von Wiedererkennungsinformation gäbe. Den beiden Forschern zufolge gäbe es hingegen klare Nachweise, dass Wiedererkennungsinformation und Wissen integriert würden. Dementsprechend könnte man beispielsweise annehmen, dass Menschen Wissen und Wiedererkennungsinformation gewichten und addieren, so dass beide Variabeln einander kompensieren können. In eine ähnliche

Richtung zielt eine Studie von B. R. Newell und Fernandez (2006). Diese Autoren schlagen vor, dass Menschen Entscheidungen eher an einem kontinuierlicheren Gefühl der Wiedererkennung (z. B. der Abrufgeschwindigkeit; „fluency“; „availability“) ausrichten als an der Rekognitionsheuristik, nach der Entscheidungen auf binären Wiedererkennungsurteilen basieren.

In meiner Dissertation habe ich dargelegt, dass solche Schlüsse als voreilig gelten müssen, weil in sämtlichen Studien zur Rekognitionsheuristik die jeweils formulierten Alternativhypothesen – beispielsweise Kompensation durch Integration oder der Rückgriff auf eher kontinuierliche Wiedererkennungsurteile – nicht als präzise (und damit testbare) Modelle formuliert wurden. Zudem gibt es eine Reihe von Studien mit Ergebnissen, die nicht im Einklang mit den Schlussfolgerungen von Richter and Späth (2006) stehen und stattdessen Evidenz dafür liefern, dass Menschen die nicht-kompensatorische Rekognitionsheuristik verwenden (siehe z. B. Goldstein & Gigerenzer, 2002; Pachur et al., 2008; Pachur & Biele, 2007). Zwei dieser Studien führen sogar Reaktionszeit- und fMRI-Daten an, die darauf hinweisen, dass die Rekognitionsheuristik per „Default” verwendet wird (Pachur & Hertwig, 2006; Volz et al., 2006). Demnach handelt es sich hier um eine Strategie, die sehr einfach und schnell ausführbar ist. Jedoch gilt auch für diese Studien, dass keine Alternativmodelle mit der Rekognitionsheuristik verglichen wurden. Genau dieses ist jedoch notwendig, um die Güte der Heuristik sowie die entsprechender Alternativvorschläge als Modelle für Verhalten zu evaluieren und um zu verstehen, wann Menschen sich der nicht-kompensatorischen Rekognitionsheuristik bedienen und wann sie andere Strategien bevorzugen.

In meinen Experimenten testete ich die Rekognitionsheuristik zum ersten Mal gegen alternative Modelle. Es handelte sich dabei um kompensatorische und nicht-kompensatorische Heuristiken, die Wissen und Wiedererkennungsinformation systematisch integrieren oder Entscheidungen basierend auf der Abrufgeschwindigkeit von Objekten – d. h. auf einem eher kontinuierlichen Gefühl der Wiedererkennung – treffen. Als Antwort auf die Kritik von Dougherty et al. (2008) untersuchte ich auch Situationen, in denen alle Objekte wiedererkannt werden sowie Situationen mit mehreren Objekten. Dabei erweiterte ich die Rekognitionsheuristik so, dass sie als Modell für die Generierung von *Vorauswahlen* verstanden werden kann („consideration sets“; Alba & Chattopadhyay, 1985; Hauser & Wernerfelt, 1990; Howard & Sheth, 1969; in der Marketingliteratur wird oft angenommen, dass Menschen solche Vorauswahlen treffen, um sich – beispielsweise beim Einkauf im Supermarkt – zwischen einer Vielzahl von Produkten zu entscheiden.) Ferner formulierte ich die Bedingungen, unter denen ein möglicher „Default“ die Rekognitionsheuristik zu verwenden außer Kraft gesetzt wird. Das kann der Fall sein, wenn

Menschen wissen, dass die Ursache (d. h. die Quelle) der Bekanntheit eines Objektes nicht den Kriteriumswert des Objektes widerspiegelt. Ferner kann der „Default" auch dann aufgehoben werden, wenn Menschen über Faktenwissen verfügen, das der Wiedererkennungsinformation stark widerspricht oder wenn ein Objekt nur sehr langsam wiedererkannt wird. In all diesen Situationen ist die Verwendung der Rekognitionsheuristik in der Regel auch nicht ökologisch rational, d. h. die Heuristik liefert keine guten Vorhersagen über die Kriteriumswerte von Objekten.

In 8 Modellvergleichen in 6 Experimenten mit einer Gesamtanzahl von über 540 Teilnehmern sagte die Rekognitionsheuristik Entscheidungen der untersuchten Personen besser vorher als 6 Alternativmodelle, die Wissen oder die Abrufgeschwindigkeit von Objekten systematisch in den Entscheidungsprozess integrieren (Experimente 11 – 16). Bei meinen Studien handelte es sich dabei sowohl um computergestützte Laborexperimente, in denen Menschen entscheiden mussten, welche von zwei Städten größer ist, als auch um Fragebogenstudien, die vor drei deutschen Landtags- bzw. Bundestagswahlen durchgeführt wurden. In diesen Fragebogenstudien sagten die Versuchsteilnehmer vorher, welche von zwei oder mehreren wiedererkannten Parteien (bzw. Kandidaten) die meisten Stimmen bei der anstehenden Wahl gewinnen würden. In diesen Studien konnte ich darlegen, dass die Rekognitionsheuristik in der Tat dazu verwandt werden kann, Vorauswahlen zu generieren; und tatsächlich stimmten die Wahlvorhersagen der Versuchsteilnehmer mit den Vorhersagen meiner erweiterten Rekognitionsheuristikmodelle überein (Experimente 12, 13, 15). Sowohl Reaktionszeitdaten (Experimente 12, 14) als auch die Ergebnisse eines Modellvergleichs in einem Experiment, in dem die Versuchsteilnehmer Entscheidungen unter hohem Zeitdruck treffen mussten (Experiment 16), erbrachten weitere Evidenz dafür, dass die Rekognitionsheuristik per „Default" verwandt wird. So konnte ich beispielsweise in Experiment 16 aufzeigen, dass die Rekognitionsheuristik einfacher und schneller anwendbar ist als die Weighted-Fluency-Heuristik, eine Strategie, die die Wiedererkennung von Objekten und deren Abrufgeschwindigkeit systematisch bei der Entscheidungsfindung integriert. Meine Befunde weisen auch darauf hin, dass Wissen über die Quelle der Bekanntheit eines Objektes (Experiment 12) und Faktenwissen, das im Widerspruch zur Wiedererkennungsinformation steht (Experiment 11), zu einer ökologisch rationalen Nichtanwendung der Rekognitionsheuristik führen können. Außerdem stellte ich in den Experimenten 13 und 14 sowie in weiteren Analysen der Experimente 1 und 2 fest, dass die Rekognitionsheuristik am ehesten dann zu genauen Entscheidungen führt, wenn die Abrufgeschwindigkeit von Objekten groß ist und damit die Ausführung der Heuristik verhältnismäßig einfach und schnell erfolgen kann. Ist die Abrufgeschwindigkeit hingegen gering und somit die Anwendung der Heuristik eher schwierig und zeitintensiv, dann ist es wahrscheinlich, dass eine Anwendung der Heuristik auch zu weniger genauen Entscheidungen führt. In der Tat scheinen Menschen die Heuristik in solchen Situationen mit einer geringeren Wahrscheinlichkeit zu verwenden. Dieser Befund lässt sich mit dem von mir im Kapitel 3 untersuchten ACT-R Gedächtnismodell vorhersagen. Meine Ergebnisse in Kapitel 4 stehen somit im Einklang mit der von mir in Kapitel 3 formulierten These, nach der die Funktionsweise des menschlichen Gedächtnisses bedingen kann, dass eine Strategie in genau jenen Situationen einfacher und schneller ausführbar ist, in denen ihre Anwendung auch zu genauen Ergebnissen führt. Kurz, es handelt sich auch hier um einen Fall ökologisch rationaler Strategieauswahl.

Fazit: Ökologisch rationale Strategieauswahl

Weite Teile der Forschung zu schnellen und sparsamen Heuristiken haben die ökologische Rationalität der Heuristiken untersucht, d. h. wie die Heuristiken an die Struktur von Entscheidungsumwelten angepasst sind (siehe z. B. Gigerenzer & Goldstein, 1996; Hogarth & Karelaia, 2007; Katsikopoulos & Martignon, 2006b; Martignon & Hoffrage, 1999). Diese Struktur wurde dabei oft durch die objektiven Eigenschaften von Umwelten und nicht so sehr durch deren Repräsentation im kognitiven System charakterisiert. In meiner Dissertation habe ich dargelegt, wie wichtig es ist, nicht nur die objektive Umweltstruktur an sich, sondern auch die Art und Weise, wie diese Struktur durch das menschliche Gedächtnis *repräsentiert* wird, zu verstehen. Meines Erachtens sind es diese mentalen Repräsentationen, die für die Auswahl zwischen den verschiedenen Heuristiken verantwortlich sind. Sie definieren die kognitiven Nischen der Heuristiken in der adaptiven Werkzeugkiste.

Insgesamt hoffe ich, mit dieser Dissertation zum Verständnis der ökologischen Rationalität der Heuristiken in der adaptiven Werkzeugkiste beizutragen, indem ich einerseits Gedächtnisinhalte basierend auf Umweltdaten modelliere (siehe auch Anderson & Schooler, 1991, 2000; Burgess, & Lund, 1997; Griffiths et al., 2007; Landauer & Dumais, 1997; Lund & Burgess, 1996; Schooler & Anderson, 1997) und andererseits, indem ich an der Integration von Gedächtnis- und Entscheidungsmodellen arbeite (siehe auch Dougherty et al., 1999; Gray et al., 2006; Schooler & Hertwig, 2005). Ich baue dabei auf den Forschungsarbeiten vieler Wissenschaftler auf, die sich mit dem Zusammenspiel des kognitiven Systems und der Umwelt beschäftigt haben (siehe z. B. Anderson, 1990; Brunswik, 1943, 1955; Gibson, 1979; Gigerenzer et al., 1999; Simon, 1956, 1990). Einer dieser Wissenschaftler, Herbert A. Simon (1996), argumentierte einmal, dass das Gedächtnis eine Erweiterung der Umwelt darstelle, in der menschliches Denken stattfinde. Da sich diese Idee durch meine Dissertation zieht, möchte ich die deutsche Zusammenfassung mit Simons (1996) eigenen Worten (auf Englisch) beenden:

> We can think of the memory as a large encyclopaedia or library, the information stored by topics (nodes), liberally cross-referenced (associational links), and with an elaborate index (recognition capability) that gives direct access through multiple entries to the topics. Long-term memory operates like a second environment, parallel to the environment sensed through eyes and ears, through which the problem solver can search and to whose contents he can respond. (p. 88)

Erklärung

Hiermit versichere ich, dass ich die vorgelegte Arbeit „Ecologically Rational Strategy Selection" selbstständig verfasst und andere als die angegebenen Hilfsmittel nicht verwendet habe. Die Arbeit ist in keinem früheren Promotionsverfahren angenommen oder abgelehnt worden.

Die Arbeit ist nicht als Ganzes veröffentlicht. Kapitel 1 basiert in Teilen auf einem begutachteten Buchkapitel (Marewski, Galesic, & Gigerenzer, in press) und einem Artikel, der in der Fachzeitschrift *Zeitschrift für Psychologie / Journal of Psychology* erscheinen wird (Marewski & Olsson, in press). Auch Kapitel 2 enthält ausgewählte Inhalte des Artikels, der in der Fachzeitschrift *Zeitschrift für Psychologie / Journal of Psychology* erscheinen wird (Marewski & Olsson, in press). Kapitel 3 ist die überarbeitete Fassung eines Artikels, der bei der Fachzeitschrift *Psychological Review* zur Publikation eingereicht worden ist (Marewski & Schooler, 2008a). Basierend auf den Gutachten hat der zuständige Editor bei *Psychological Review* mir die Möglichkeit des Einreichens einer revidierten Version gegeben, d. h. der augenblickliche Status dieses Artikels ist „revise resubmit" bei *Psychological Review*. Die Arbeit, auf der Kapitel 3 beruht, ist außerdem im Jahr 2007 mit dem *Brunswik New Investigator Award* ausgezeichnet worden. Kapitel 4 ist eine erweiterte und gleichzeitig gekürzte Fassung eines Artikels (Marewski, Gaissmaier, Schooler, Goldstein, & Gigerenzer, 2008), der bei der Fachzeitschrift *Cognitive Psychology* zur Publikation eingereicht worden ist und dort augenblicklich begutachtet wird. Weiterhin sind Teile von Kapitel 1 und 2 in stark überarbeiteter und ergänzter Form bei der Fachzeitschrift *Cognitive Processing—International Quarterly of Cognitive Science* zur Publikation eingereicht worden (Marewski, Gaissmaier, & Gigerenzer, 2008). Ferner habe ich Teile von Kapitel 3 und 4 auf verschiedenen Tagungen vorgestellt (siehe Curriculum Vitae).

In näherer Zukunft ist vorgesehen, bisher unveröffentlichte Teile der Dissertation in überarbeiteter und erweiterter Form zur Veröffentlichung bei weiteren Fachzeitschriften einzureichen. Drei Artikel sind in diesem Zusammenhang geplant. Kapitel 2 und 5 bilden die Grundlage für einen Artikel zum Thema „Methoden der Erforschung einfacher Heuristiken", bei dem Lael Schooler und Gerd Gigerenzer als meine Koautoren fungieren werden. Das in Kapitel 3 vorgestellte ACT-R Modell soll auf die Beschreibung der kognitiven Nische der Rekognitionsheuristik angewandt werden. Mein Koautor für einen entsprechenden Artikel wird Lael Schooler sein. Teile von Kapitel 4 sollen ferner in einen Artikel zum Thema „Wahlvorhersagen mit

Rekognition“ eingehen. Meine Koautoren werden hier Wolfgang Gaissmaier, Lael Schooler und Gerd Gigerenzer sein.

Alle angeführten Koautoren werden bestätigen, dass ich der Hauptverantwortliche für die Ideen, die Planung und die Durchführung der Projekte, die Datenanalysen und für das Schreiben der Kapitel war.

Julian Marewski

Berlin, 30. November 2008

Publication of Dissertation Chapters – Update December 2010

Revised and extended portions of work on which Chapters 1 and 2 were partially based, has been published as follows:

- Marewski, J. N., Galesic, M., & Gigerenzer, G. (2009). Fast and frugal media choices. In T. Hartmann (Ed.),*Media choice: A theoretical and empirical overview* (pp. 107-128). New York & London: Routledge.
- Marewski, J. N., & Olsson, H. (2009). Beyond the null ritual: Formal modeling of psychological processes. *Zeitschrift für Psychologie / Journal of Psycholoy*, *217*, 49–60. doi: 10.1027/0044-3409.217.1.49

A revised and extended version of Chapter 3 has received strong encouragement for resubmission to Psychological Review.

- Marewski, J. N. & Schooler, L. J. (2010). Cognitive niches. An ecological model of emergent strategy selection.

A revised, extended and shortened version of Chapter 4 has been published as follows:

- Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2010). From Recognition to Decisions: Extending and Testing Recognition-Based Models for Multi-Alternative Inference. *Psychonomic Bulletin and Review*, *17*, 287- 309. doi: 10.3758/PBR.17.3.287

In addition, a revised section of Chapter 4 that has not been published in the just mentioned Psychonomic Bulletin and Review article has been published as follows:

- Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2009). Do Voters Use Episodic Knowledge to Rely on Recognition? In N.A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2232-2237). Austin, TX: Cognitive Science Society.

A revised and extended version of Chapter 5 has been published as follows:

- Marewski, J. N., Schooler, L. J., & Gigerenzer, G. (2010). Five principles for studying people’s use of heuristics. *Acta Psychologica Sinica*, *42*, 72-87. doi: 10.3724/SP.J.1041.2010.00072

CURRICULUM VITAE

Following the Free University's publishing rules, the curriculum vitae has been taken out of this dissertation.