The recognition heuristic: A decade of research

Gerd Gigerenzer[*]   Daniel G. Goldstein[†]

Abstract

The recognition heuristic exploits the basic psychological capacity for recognition in order to make inferences about unknown quantities in the world. In this article, we review and clarify issues that emerged from our initial work (Goldstein & Gigerenzer, 1999, 2002), including the distinction between a recognition and an evaluation process. There is now considerable evidence that (i) the recognition heuristic predicts the inferences of a substantial proportion of individuals consistently, even in the presence of one or more contradicting cues, (ii) people are adaptive decision makers in that accordance increases with larger recognition validity and decreases in situations when the validity is low or wholly indeterminable, and (iii) in the presence of contradicting cues, some individuals appear to select different strategies. Little is known about these individual differences, or how to precisely model the alternative strategies. Although some researchers have attributed judgments inconsistent with the use of the recognition heuristic to compensatory processing, little research on such compensatory models has been reported. We discuss extensions of the recognition model, open questions, unanticipated results, and the surprising predictive power of recognition in forecasting.

Keywords: recognition heuristic, ecological rationality, inference, heuristics, bounded rationality.

1   Introduction

With Herbert Simon's (1990) emphasis on recognition memory and limited search as a starting point, it was only a small logical step towards the recognition heuristic, which exploits the potential information in a lack of recognition. In accordance with Simon's emphasis on computational models, the recognition principle (as it was first called) was formulated as a building block of take-the-best and other heuristics, in order to model inferences from memory (Gigerenzer & Goldstein, 1996). Subsequently, it was realized that this initial building block could function as a stand-alone model for the same type of inferences, and it was named the recognition heuristic (Goldstein & Gigerenzer, 1999, 2002).

In reality, the recognition heuristic was not derived in such a logical manner. Serendipity, the luck of finding something one was not looking for, assisted its birth. Gigerenzer, Hoffrage, and Kleinbölting (1991, Prediction 4) had deduced from probabilistic mental models theory a situation in which the "hard-easy" effect would disappear. In his dissertation, Ulrich Hoffrage (1995; described in Hoffrage, 2011) set out to test this prediction, for which he needed two sets of questions, one hard, one easy. Hoffrage chose questions concerning the populations of American cities and German cities, which are respectively hard and easy for German students—or so everyone thought. Surprisingly, the students scored slightly higher when tested on a representative sample of American cities than on German ones. The result ruined the experiment. How could people score more correct answers in a domain in which they knew less? For days, our research group failed to think of a cognitive process that makes more out of less. Finally, Anton Kühberger pointed out that the explanation was tucked away in the Gigerenzer et al. (1991) article, which mentioned "familiarity" as a probabilistic cue. If a person has heard of one city but not the other, this lack of recognition can be informative, indicating that the recognized city probably has the larger population. For the German cities, the participants could not rely on the recognition heuristic—they knew too much. This serendipitous discovery also revealed a crucial condition for the successful reliance on recognition: a substantial correlation between recognition and population (the recognition validity), and a representative sampling of the cities. We return to this condition later.

One possible reason why it took us so long to find the answer was our training in classical statistical models. In a weighted linear model, adding a cue or predictor can never decrease its fit, such as unadjusted R^2, and the same is true for Bayes' rule (McGrath, 2008). This more-is-better principle holds for fitting parameters to known data, but not necessarily for predicting what one does not already know, as the German students had to do. A good cognitive heuristic, however, should excel in foresight as well as in hindsight.

The possibility that people could sometimes do better with less knowledge has generated much interest and controversy in the social sciences and in the media. In May 2009, the BBC, intrigued by the idea of less being more, decided to test the effect on their Radio 4 "More or less" program. Listeners in New York and London were asked whether Detroit or Milwaukee has the larger population. In exploratory studies for his dissertation, one of us (Goldstein, 1997) had found that about 60% of American students answered correctly ("Detroit"), compared to 90% of a corresponding group of German participants. The BBC is not known for running tightly controlled studies, and so we were somewhat uneasy about whether they could reproduce a less-is-more effect. But they did. In New York, 65% of the listeners got the answer right, whereas in London, 82% did so—as close as one can hope for an informal replication.

Our initial work on the recognition heuristic has stimulated dozens of articles comprising theoretical advancements, critique, and above all, progress. This work has contributed much to understanding the potential and limits of this simple model, but we also believe that its broad reception represents a larger shift in research practice. This change occurs in three directions:

1. *From labels to models of heuristics.* It is an interesting feature of the recent history of psychology that vague labels such as *availability* had remained largely unquestioned for three decades (for an exception, see Wallsten, 1983), whereas the precise model of the recognition heuristic immediately led to heated debates.

2. *From preferences to inferences.* While formal models such as elimination-by-aspects and lexicographic rules have been studied for preferences (e.g., Payne, Bettman, & Johnson, 1993; Tversky, 1972), their accuracy was typically measured against the gold standard of adding and weighting all information. In this earlier research, a heuristic could not—by definition—be more accurate, only more frugal. For inferences, in contrast, there exists a clear-cut criterion, and thus it was possible to show that cognition can actually achieve more accuracy with simple heuristics than with weighted additive rules. This leads to the third shift in research.

3. *From logical to ecological rationality.* For decades, human rationality was studied in psychology by proposing a logical or statistical rule (e.g., truth table logic; Bayes' rule) as normative in all situations, and then constructing an artificial problem in which this rule could be followed, such as the Wason Selection Task (Wason, 1971) or the Linda Problem (Kahneman & Tversky, 1982). In contrast, the question

of ecological rationality asks in which environment a given strategy (heuristic or otherwise) excels and in which it fails. No rule is known that is rational *per se*, or best in all tasks. Parts of the psychological research community have resisted the asking of questions about ecological as opposed to logical rationality.

We begin our review of the progress made in the last decade with the two key processes that govern the use of the recognition heuristic: recognition and evaluation, the latter of which corresponds to a judgment of its ecological rationality.

2　The recognition process

The recognition heuristic makes inferences about criteria that are not directly accessible to the decision maker. When the criterion is known or can be logically deduced, *inferential* heuristics like the recognition heuristic do not apply. Relying on the heuristic is ecologically rational in an environment R where the recognition of objects $a, b \in R$ is strong and positively correlates with their criterion values. For two objects, the heuristic is:

> If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion.

In our original work (Gigerenzer & Goldstein, 1996, pp. 651–652; Goldstein & Gigerenzer, 2002, pp. 76–78), we assumed that the recognition heuristic will model inferences when three conditions hold:

(i) there is a substantial recognition validity;

(ii) inferences are made from memory, rather than from tables of information ("inferences from givens"), meaning that cue values for unrecognized objects are missing values; and

(iii) recognition stems from a person's natural environment (i.e., before entering the laboratory), as opposed to experimentally induced recognition.

We return to these characteristics below. We would like to emphasize that such a definition of the domain is essential, just as in game theory, where rational strategies are defined (and restricted) to game features such as anonymity, a fixed number of repetitions, and no reputation effects. This is not to say that studies that test predictions outside the domain are useless; on the contrary, they help to map out the boundary conditions more clearly, as we ourselves and other researchers have done. For example, we conducted a long-run experiment that subtly induced recognition over several weeks to investigate the effect of exogenous recognition on choice (Goldstein & Gigerenzer, 2002, pp. 84–85).

The recognition heuristic stands on the shoulders of the core psychological capacity of recognition memory; without it, the heuristic could not do its job. However, the recognition heuristic is mute about the nature of the recognition process, just as Bayes' rule is mute about where the prior probabilities come from. Heuristics exploit core capacities in order to make fast and frugal judgments. Examples include recall memory (e.g., take-the-best), frequency monitoring (e.g., fast-and-frugal trees), mimicry (e.g., imitate-the-majority), and object tracking (e.g., gaze heuristic), with some heuristics taking advantage of several capacities (Gigerenzer & Brighton, 2009).

2.1 Connecting the recognition heuristic with the recognition process

In our original work, we did not investigate the link between the recognition heuristic and theories of the underlying recognition process. Since then, progress has been made towards this goal of theory integration, a topic that is of utmost importance in fields such as physics but is given little attention in psychology. As Walter Mischel (2006) put it, many psychologists still tend to treat theories like toothbrushes—no self-respecting person wants to use anyone else's. In one step towards theory integration, the recognition heuristic has been implemented based on the ACT-R model of memory (Anderson & Lebiere, 1998), which showed in some detail how forgetting—a process often seen as a nuisance and a handicap—can be functional in the context of inference (Schooler & Hertwig, 2005). In this same work, the fluency heuristic (Table 1) was formulated for the situation when both alternatives are recognized, that is, when the recognition heuristic cannot be applied. This work also integrated earlier work on fluency (e.g., Jacoby & Dallas, 1981) into the simple heuristics framework, defined the difference between the recognition and fluency heuristics, and thus contributed towards replacing verbal labels with computational models. Moreover, Schooler and Hertwig's analysis challenges the common belief that cognitive limits, such as forgetting or a limited working memory, inevitably pose liabilities for the human mind. Some cognitive limits foster specific cognitive processes, and at the same time some cognitive processes exploit specific cognitive limits—as may be the case in the interplay of forgetting and heuristic inference.

A second theoretical integration has combined a signal detection model of recognition memory with the recognition heuristic (Pleskac, 2007). In our original work, we had not separately analyzed how the recognition validity changes depending on what proportion of recognition judgments are correct. When recognizing an object, people can go wrong by erroneously recognizing something that they have never encountered before ("false alarms")

and by failing to recognize something that they have previously encountered ("misses"). Pleskac showed that, as the error rate of recognition increases, the accuracy of the recognition heuristic declines, and that the less-is-more effect is more likely when participants' sensitivity (d') is high, whereas low sensitivities lead to "more-is-more." Furthermore, when people are cognizant of their level of recognition knowledge, they can increase their inferential accuracy by adjusting their decision criterion accordingly. When the amount of knowledge is very low, it may be prudent to be conservative in judging something as previously encountered; with increasing knowledge, however, it is better to become more liberal and classify something as previously encountered, even if one is not absolutely certain.

2.2 Is recognition binary?

In our original work, we modeled the input for the recognition heuristic, the recognition judgment, as binary. For instance, the brand name *Sony* would be modeled as either recognized or not. This simplifying assumption has been criticized (e.g., Hilbig & Pohl, 2009). However, it is consistent with theories that distinguish between a continuous process of recognition (or familiarity) and a binary judgment, such as signal detection theory, where an underlying continuous sensory process is transformed by a decision criterion into a binary recognition judgment. Moreover, there is now evidence that not only the recognition judgment, but the recognition process itself may be binary or threshold-like in nature. Bröder and Schütz (2009) argued that the widespread critique of threshold models is largely invalid, because it is, for the most part, based on confidence ratings, which are nondiagnostic for rejecting threshold models. In a reanalysis of 59 published studies, they concluded that threshold models fit the data better in about half of the cases. Thus, our assumption of a binary input into the recognition heuristic is a simplification, but not an unreasonable one, as it is supported by evidence and theories of the nature of recognition. (But see Hoffrage, 2011, sec. 3.3.5, for some evidence against a simple threshold.) Note that a model that assumes binary recognition judgments does not imply that organisms are unable to assess the degree to which something is familiar or frequent in the environment (Malmberg, 2002). In fact, models such as the fluency heuristic exploit such information (Schooler & Hertwig, 2005).

2.3 Individual recognition memory constrains the selection of heuristics

No heuristic is applied indiscriminately to all situations (Payne, Bettman, & Johnson, 1993; Goldstein et al.,

Table 1: Four heuristics from the adaptive toolbox. Which to use for a given task? The content of individual memory determines whether an individual can apply the recognition heuristic (or other heuristics), and an evaluation process determines whether it should be applied.

Heuristic	Definition	Ecologically rational if:	Predictions
Recognition heuristic (Goldstein & Gigerenzer, 2002).	If one of two alternatives is recognized, infer that it has the higher value on the criterion.	Recognition validity >.5	Contradicting information about the recognized object is ignored; less-is-more effect; forgetting is beneficial.
Fluency heuristic (Schooler & Hertwig, 2005)	If one alternative is recognized faster than another, infer that it has the higher value on the criterion.	Fluency validity >.5	Less-is-more effect; forgetting is beneficial (Hertwig et al., 2008)
Take-the-best (Gigerenzer & Goldstein, 1996)	To infer which of two alternatives has the higher value: (1) search through cues in order of validity, (2) stop search as soon as a cue discriminates, (3) choose the alternative this cue favors.	Cue weights vary highly; moderate to high redundancy; scarce information (Hogarth & Karelaia, 2005, 2006; Katsikopoulos & Martignon, 2006; Martignon & Hoffrage, 1999, 2002).	Can predict as or more accurately as linear regression (Czerlinski et al. 1999), neural networks, exemplar models, and CARTs (Brighton, 2006).
Tallying / unit-weight linear model (Dawes, 1979)	To estimate a criterion, do not estimate weights but simply count the number of favoring cues.	Cue weights vary little; low redundancy (Hogarth & Karelaia, 2005, 2006).	Can predict as or more accurately than multiple regression (Czerlinski, Gigerenzer & Goldstein, 1999).

2001), and the recognition heuristic is no exception. How are heuristics selected from the adaptive toolbox? Marewski and Schooler (2011) have developed an ACT-R model of how memory can constrain the set of applicable heuristics. Consider the following set of strategies: the recognition heuristic, the fluency heuristic, take-the-best, and tallying (Table 1), in connection with the task of betting money on which tennis player, Andy Roddick or Tommy Robredo, will win against the other. Each of the four heuristics is *potentially* applicable for this task (the gaze heuristic, for instance, would be inapplicable). Whether a strategy is *actually* applicable for a given individual, however, depends on the state of individual memory. First, if an individual is ignorant about tennis and has heard of neither of the players, none of the heuristics can be applied and that person might simply guess. Second, if a person has heard of Roddick but not of Robredo, this state of memory restricts the choice set to the recognition heuristic; the bet would be on Roddick. As it turns out, Roddick and Robredo have played 25 sets against each other so far (by 2010) and Roddick has won 24 of them. The person who uses the recognition heuristic will, by definition, not be able to recall this fact from memory, having never heard of Robredo, but can nevertheless bet correctly. Third, consider an individual who has heard of both players, but recalls nothing else about them. This

state of memory excludes the recognition heuristic, as well as take-the-best and tallying, and limits the choice set to the fluency heuristic: If both players are recognized, but one was recognized more quickly than the other, predict that the more quickly recognized player will win the game.

Finally, consider an individual more knowledgeable about tennis who has heard of both players, and can also recall the values of both on relevant cues, such as their current ATP Champions Race ranking, their ATP Entry ranking, their seeding by the Wimbledon experts, and the results of their previous matches. This state of memory again excludes the recognition heuristic, but leaves the other three heuristics in the choice set. To choose between these, principles of ecological rationality come into play. For instance, if cues are moderately to highly redundant, take-the-best has an advantage over tallying, and participants in experiments tend to prefer take-the-best after simply observing the structure of the environment (such as the degree to which cues were intercorrelated): No feedback about accuracy appears to be necessary (Dieckmann & Rieskamp, 2007). When feedback is available, Strategy Selection Learning theory (SSL theory) provides a quantitative model of heuristic selection (Rieskamp & Otto, 2006). SSL theory makes predictions about the probability that a person selects one heuristic

within a defined set and shows how learning by feedback leads to adaptive strategy selection.

To summarize: In the decade after our initial work, we now have a sharp distinction between the recognition heuristic and the fluency heuristic. The effects of misses and false alarms on the accuracy of the recognition heuristic are better understood. Our grossly simplifying assumption of modeling recognition judgments as binary turned out to be largely consistent with a body of empirical evidence, although this issue is far from being settled. We postulate that individual recognition memory is the basis for the first of two steps by which an individual decides whether to rely on the recognition heuristic for solving a given task. The state of recognition memory determines whether it can be applied, while an evaluation process, our next topic, determines whether it should be applied.

3   The evaluation process

If the recognition heuristic satisfies the individual memory constraint (to recognize one of two objects), then an evaluation process is needed to determine whether relying on the recognition heuristic is ecologically rational for the particular inference being made. We titled our 2002 article "Models of *ecological* rationality: The recognition heuristic", emphasizing that the heuristic is not general-purpose, but selected in an adaptive way that depends on the environment (i.e., ecology). In our original work, we had specified one condition for the ecological rationality of the recognition heuristic (Goldstein & Gigerenzer, 2002, p. 87):

Substantial recognition validity. The recognition validity for a given criterion must be substantially higher than chance ($\alpha > .5$).

Evaluating the recognition validity requires the existence of reference class R of objects (Goldstein & Gigerenzer, 2002, p. 78; Gigerenzer & Goldstein, 1996, p. 654). We take this opportunity to clarify:

Precondition 1: Existence of a reference class. Without a reference class R (such as the class of all 128 contestants in Wimbledon Gentleman Singles), neither the experimenter nor the participant can estimate whether there is a substantial recognition validity. In other words, the more uncertain one is about the identity of the reference class, the less one can know about whether relying on the recognition heuristic is ecologically rational.

Precondition 2: Representative sampling. Assuming a substantial recognition validity, the successful use of the recognition heuristic for a specific pair $a, b \in R$ presupposes that it has been representatively sampled from R, rather than selectively sampled in a biased way (e.g., such that a high recognition validity in R is misleading for the

specific task). For instance, when we asked international audiences during talks we have given outside the United States whether Detroit or Milwaukee has the larger population, some answered "Milwaukee" despite never having heard of it; they explained that they thought it was a trick question, that is, one selectively sampled for being counterintuitive. This suspicion reflects the widespread view that psychologists routinely deceive their participants (Hertwig & Ortmann, 2001), but is not the only reason why people may suspect biased sampling and overrule the recognition heuristic. Yet, whereas biased sampling can be hard for a participant to judge, the absence of a meaningful reference class can easily be noticed. Thus, we assume that, in the presence of a substantial recognition validity, people will consider applying the recognition heuristic by default, that is, unless there is reason to assume biased sampling of objects. The issue of representative sampling of questions is described in detail in Gigerenzer et al. (1991) and Hoffrage (2011).

Table 2 (in the Appendix) includes all studies we know of that report correct predictions of judgments by the recognition heuristic ("accordance rates"). It reports the reference class, the criterion, and whether the three conditions that define the domain of the recognition heuristic were in place. It also reports two methodological features: whether the recognition heuristic was tested comparatively against alternative models, and whether individual analyses were performed, as opposed to means only (see below). The last column shows the recognition validity and the mean correct predictions of judgments by the recognition heuristic.

3.1   Does the strength of recognition validity relate to the predictive accuracy of the recognition heuristic?

In our 2002 article, we had not systematically investigated this question. The research of the last years suggests that the answer may be affirmative. For instance, name recognition of Swiss cities is a valid predictor of their population ($\alpha = .86$), but not for their distance from the center of Switzerland ($\alpha = .51$). Pohl (2006) reported that 89% of inferences accorded with the recognition heuristic model in judgments of population, compared to only 54% in judgments of the distance. Similarly, the first study on aging indicates that old and young people alike adjust their reliance on the recognition heuristic between environments with high versus low recognition validities, even though old people have worse recognition memory (Pachur, Mata, & Schooler, 2009).

In Figure 1, we plotted the recognition validities against the proportion of correct predictions by the recognition heuristic (accordance rates) for all study conditions from Table 2. Included are all studies that re-

Figure 1: Relationship between recognition validity and mean percentage of correct predictions of the recognition heuristic (accordance rate). Included are all 43 experiments or conditions in Table 2 where alpha and accordance rates were reported, inside and outside the domain of the recognition heuristic. Black symbols represent experiments/conditions with natural recognition and inferences from memory. Black triangles = 3 negative (contradicting) cues; black squares = 1 negative (contradicting) cue. White diamonds = repetition during the experiment rather than natural recognition (Bröder & Eichler, 2006); white diamonds with cross = repetition and inferences from givens (Newell & Shanks, 2004). Here, repetition validity is reported instead of recognition validity. Richter and Späth (2006, Exp. 3) reported a rank correlation instead of alpha, which we transformed into an estimate of alpha using Equation 2 in Martignon and Hoffrage (1999). Mixtures of positive and negative cues (Pachur, Bröder & Marewski, 2008, Exp. 1, all accordance rates >.96) are not included. The best fitting linear relation is shown; the Pearson correlation is $r = .57$.
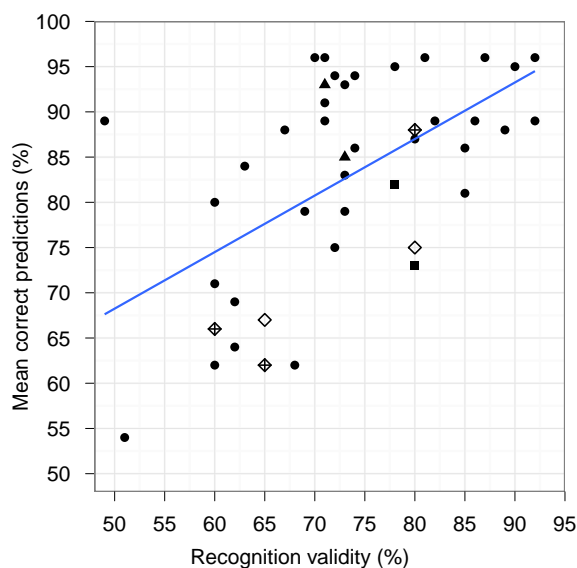
ported recognition validities and correct predictions, even when the objects were not representatively sampled, as well as studies that tested the recognition heuristic outside its domain. Figure 1 shows that when participants were provided with up to three negative cues (black triangles and squares), the results still fall into the general pattern, while studies that tested the recognition heuristic outside its domain appear to result in lower correct predictions. Studies that used inferences from givens or experimentally induced recognition are shown by white symbols. Across all study conditions, the correlation be-

tween recognition validity and proportion of correct predictions is $r = .57$.

Note that there are two interpretations of this correlation. One is that most individuals engage in probability matching, that is, they rely on the heuristic in a proportion of trials that corresponds to the recognition validity. However, the evidence does not support this hypothesis (Pachur, Bröder, and Marewski, 2008, Figure 5; Pachur & Hertwig, 2006). The second assumes differences among individuals and tasks, for instance that people tend to rely on the recognition heuristic consistently when the validity for them is high, but when the validity decreases, people increasingly suspend the default and follow some other strategy. Analyses of individual differences, such as in Figure 2 below, indicate that the latter may be the rule.

3.2 Is reliance on the recognition heuristic sensitive to the specification of a reference class?

Substantial recognition validities have been reported in various environments, including predicting the winners in the Wimbledon Gentlemen Singles matches in the class of final contestants (Scheibehenne & Bröder, 2007; Serwe & Frings, 2006), inferring the career points of professional hockey players (Snook & Cullen, 2006), and forecasting the election results of political parties and candidates (Marewski et al., 2010). In each of these studies, a large majority of participants consistently made inferences that were predicted by the recognition heuristic. In contrast, accordance appears to be low in studies where no reference class was specified or neither researchers nor participants could possibly estimate the recognition validity. For instance, Oppenheimer (2003, Experiment 2) asked students at Stanford University in a questionnaire to make paired comparison judgments of city population. There were six key pairs, each of which consisted of a nonexistent, fictional city, such as "Heingjing", and a real city, selected for specific reasons, such as Chernobyl (nuclear disaster), Cupertino (close proximity), and Nantucket (popular limerick). Because no real reference class exists, it is not possible to determine the recognition validity in this study. In the end, in less than half of the cases were the recognized cities judged to be larger (Table 2). The study concluded that it "found no evidence for the recognition heuristic despite using *the same task* as in the original studies" (p. B7, italics in the original). However, it was not the same task. In the original studies, participants knew the reference class, knew it was real, and were tested on its identity (Goldstein & Gigerenzer, 2002).

Let us make it clear. We do not assume that people follow the recognition heuristic unconditionally, for example independently of recognition validity, as a few researchers have implied. Sensitivity to the specifica-

tion of reference classes (or lack thereof) has been documented in research by ourselves and others, and is of general importance for understanding human judgment. For instance, single-event probabilities by definition do not specify a class, which results in confusion and large individual differences in interpretation, as in probabilistic weather forecasts (Gigerenzer et al., 2005) and clinical judgments of the chances of a patient harming someone (Slovic, Monahan, & MacGregor, 2000).

3.3 Are people sensitive to the presence or lack of representative sampling?

As far as we can see, in environments with substantial recognition validity, a well-defined reference class, and representative sampling, a substantial proportion of participants act in accordance with the recognition heuristic when making inferences from memory. Moreover, if there is a reference class with a substantial recognition validity, such as the height of the largest mountains, but the objects are selected so that the sample recognition validity is close to chance level, recognition heuristic accordance can still be quite high. The outlier in the left upper left corner of Figure 1 is such a case of selected sampling (Pohl 2006, Experiment 4; see Table 2). As mentioned above, although it is easy to detect whether there is a meaningful reference class, it is sometimes difficult to detect whether objects are randomly or selectively sampled from this class. Pohl's experiment with a selected sample of mountains suggests that people might assume random selection in the absence of any red flags. Except for Hoffrage (1995; 2011), we are not aware of any systematic study that varied representative and biased sampling in inferences from recognition; however, studies on Bayesian judgments suggest sensitivity to random versus biased sampling if the sampling is performed or witnessed by the participant (Gigerenzer, Hell, & Blank, 1988).

In summary, two processes, recognition and evaluation, do much to guide the adaptive selection of the recognition heuristic. They can be formulated as two questions: "Do I recognize one object but not the other?" "If so, is it reasonable to rely on the recognition heuristic in this situation?" The first constrains the set of applicable heuristics. The second process provides an evaluation check, judging the ecological rationality of the heuristic in a given task. Experimental results indicate that participants were sensitive to differences in recognition validity between tasks (Figure 1) and that they were sensitive to the existence of a meaningful reference class R (Table 2). Sensitivity to sampling from R needs to be studied. How this sensitivity and the associated evaluation process works is not yet well understood; however, the research in the following section provides some progress and hypotheses.

4 The neural basis of the recognition and evaluation processes

An fMRI study tested whether the two processes, recognition and evaluation, can be separated on a neural basis (Volz et al., 2006). Participants were given two kinds of tasks; the first involved only a recognition judgment ("Have you ever heard of Modena? of Milan?"), while the second involved an inference in which participants could rely on the recognition heuristic ("Which city has the larger population: Milan or Modena?"). For mere recognition judgments, activation in the precuneus, an area that is known from independent studies to respond to recognition confidence, was observed (Yonelinas, Otten, Shaw, & Rugg, 2005). In the inference task, precuneus activation was also observed, as expected, and in addition activation was detected in the anterior frontomedian cortex (aFMC), which has been linked in earlier studies to evaluative judgments and self-referential processing. These results indicate that the neural processes elicited by the two tasks of recognition and evaluation are not identical, as an automatic interpretation of the use of the heuristic would imply, but suggest a separate evaluation process that determines whether to select the recognition heuristic for a given task. The aFMC activation could represent the neural basis of this evaluation of ecological rationality.

The neural evidence furthermore suggests that the recognition heuristic may be relied upon by default, as opposed to being just one of many strategies. The default can be overthrown by information indicating that it is not ecologically rational to apply the heuristic in a particular task because recognition is not predictive of the criterion (Volz et al., 2006). The default interpretation is also supported by behavioral data. Response time data from Pachur and Hertwig (2006) as well as Volz et al. suggest that recognition judgments are made before other knowledge can be recalled. Consistent with this hypothesis, these authors show that response times were considerably faster when participants' inferences accorded with the recognition heuristic than when they did not. Similarly, participants' inferences accorded with the recognition heuristic more often when they were put under time pressure. Moreover, even though older people have slower reaction times, they also reacted faster when choosing the recognized object (Pachur et al., 2009). These findings are consistent with the recognition memory literature, indicating that a sense of recognition (often called *familiarity*) arrives in consciousness earlier than recollection (e.g., Ratcliff & McKoon, 1989). Recognition judgments are made very quickly, and the recognition heuristic appears to be a default strategy that can be overthrown by information contradicting its perceived ecological rationality.

5 Correcting misconceptions

We have seen three misconceptions about the nature of the recognition heuristic. The first was already briefly mentioned, the second concerns the meaning of a non-compensatory strategy, and the third the original domain of the heuristic.

Misunderstanding #1: All people rely indiscriminately on the recognition heuristic in all situations. Some researchers have ascribed to us the view that the recognition heuristic is "universally applied" or that people "rely on recognition blindly" (e.g., Richter & Späth, 2006, p. 160). Others used multinominal models to test the null hypothesis that people rely on the recognition heuristic 100% (or 96%) of the time, and found that only some people exhibit this level of consistency (see Brighton & Gigerenzer, 2011). We know of no model of judgment that predicts 96% correctly, and in all situations. In contrast, our view was and is that the recognition heuristic—like other heuristics—is likely to be applied when it is ecologically valid, not in all situations. This is implied by the very notion of the adaptive toolbox. Furthermore, different individuals select different heuristics, as we shall discuss.

Misunderstanding #2: *A noncompensatory strategy ignores all other information, not just other cue values.*

Consider an ordered set of M binary cues, $C_1, \ldots C_M$. These cues are noncompensatory for a given strategy if every cue C_j outweighs any possible combination of cues after C_j, that is, C_{j+1} to C_M. In the special case of a weighted linear model with a set of weights $W = \{w_1, \ldots, w_M\}$, a strategy is noncompensatory if (Martignon & Hoffrage, 1999):

$$w_j = \sum_{k>j} w_k \text{ for every } i \le j \le M \qquad (1)$$

In words, a linear model is noncompensatory if, for a given ordering of the weights, each weight is larger than the sum of all subsequent weights. A simple example is the set $\{1, 1/2, 1/4, 1/8, 1/16\}$. Noncompensatory models include lexicographic rules, conjunctive rules, disjunctive rules, and elimination-by-aspects (Hogarth, 1980; Tversky, 1972).

The definition shows that *noncompensatory* refers to a relationship between one cue and other cues, not a relationship between one cue and the criterion (e.g., knowing a city is small). We clarify this here, because in our original article we used the terms *further cue values* and *further information* interchangeably, assuming the technical meaning of *noncompensatory* to be known. But that was not always the case, and thus we may have contributed to the misunderstanding. For instance, people recognize the name *Madoff* as a money manager of the last decade but do not infer him to have been reputable because they have direct knowledge about him on this criterion. With certain knowledge about the criterion, no inference may be needed—one might simply make deductions using a local mental model (see Gigerenzer et al., 1991, Figure 2). If there is no inference to be made, the recognition heuristic does not apply, and the proportion of inferences consistent with the recognition heuristic is likely to be at or below chance level (Pachur & Hertwig, 2006; see also Hilbig, Pohl, & Bröder, 2009). Proper tests of noncompensatory processing introduce cues for the recognized object (but not the unrecognized object) that contradict the inference the recognition heuristic would make (see below). In sum, a noncompensatory process ignores cues, but not information in general, such as information concerning criterion values or the recognition validity.

Misunderstanding #3: The recognition heuristic is a model of inference in general, rather than of inference from memory The recognition heuristic, like take-the-best, was explicitly proposed as a model of inferences made from memory, that is, inferences in which each object's cue values are retrieved from memory, as opposed to inferences from givens, in which the cue values are provided by the experimenter (Gigerenzer & Goldstein, 1996, pp. 651–652). Inferences from memory are logically different from inferences based on external information. If one has not heard of an object, its cue values cannot be recalled from memory (although the name itself may, quite rarely, impart cue values, much like "80 proof whiskey" reveals its alcohol content). Thus, in inferences from memory, recognition is *not* like other cues. Rather, recognition can be seen as a prior condition for being able to recall further cue values from memory. In inferences from givens, in contrast, this logical relation does not hold, and recognition could, but need not, be treated as just another cue. Note that proponents of evidence accumulation models or parallel constraint satisfaction models tend to neglect this fundamental distinction when they consider "recognition as one cue among others" (Hilbig & Pohl, 2009, p. 1297; see also Glöckner & Bröder, 2011).

In addition to this logical difference, memory-based inferences are also psychologically different. Memory-based inferences require search in memory for cue values, whereas inferences from givens do not. Importantly, search in memory has been shown to elicit more noncompensatory processing (Bröder & Schiffer, 2006). Nevertheless, some tests of the recognition heuristic focused on inferences from given information, even about unrecognized objects (e.g., Newell & Shanks, 2004; Glöckner & Bröder, 2011). Moreover, in some studies, recognition

was induced experimentally by repetition within a session rather than arising naturally over time (e.g., Bröder & Eichler, 2006, Newell & Shanks, 2004). These studies went beyond the domain of the recognition heuristic and mostly show lower levels of correct predictions (Figure 1). We are very supportive of testing precise models for a variety of tasks, such as making inferences about an unrecognized product whose attributes values are listed on the box, but we emphasize that this task environment is outside that which is modeled by the recognition heuristic.

To summarize: No strategy is, or should be, universally applied. Assuming an automatic use of the recognition heuristic is the first misunderstanding. The proper question is: When do people rely on the heuristic, and when should they? The terms *noncompensatory* and *compensatory* refer to how cues are processed when making inferences about a criterion; they do not refer to ignoring or using information about the criterion or about the ecological rationality of a strategy. Finally, the recognition heuristic is a model of inference from memory, not from givens, with recognition and its validity learned in a person's natural environment.

# 6   Testing noncompensatory inferences

Although some models of heuristics are compensatory (for instance, unit weight models or tallying in Table 1, and the compensatory "recognition cue" models in Gigerenzer & Goldstein, 1996), many process information without making trade-offs, that is, in a noncompensatory way. For instance, the availability heuristic (Tversky & Kahneman, 1973) predicts that judgments of likelihood are based on the speed (or number, since definitions vary) with which instances come to mind. It appears to process speed (or number) in a noncompensatory way; no integration of other cues is mentioned. Similarly, the affect heuristic (Slovic, Finucane, Peters, & MacGregor, 2002) captures the notion that judgments are based on the affective tag associated with an object. Both heuristics appear to entail that people make judgments based on only a single piece of information—ease of retrieval and affect, respectively—and ignore further cue values. Being described in general terms rather than as an explicit computational model, however, the assumption of noncompensatory processing is not made explicit and may not even be intended by some authors. Perhaps this lack of clarity is one reason why the various apparent examples of one-reason decision making postulated by the heuristics-and-biases program have not sparked debate over noncompensatory processing.

Not so with the recognition heuristic. When we spelled out that the recognition heuristic is a model that relies on recognition and does not incorporate further probabilistic cues, this modeling assumption drew heavy fire. The intense reaction continues to puzzle us, given that noncompensatory processes have been frequently reported. Over 20 years ago, a classic review of 45 *process-tracing* (as opposed to outcome) studies of decision-making concluded, "the results firmly demonstrate that noncompensatory strategies were the dominant mode used by decision makers" (Ford et al., 1989, p. 75). Today, we know of several structures of environments in which not making trade-offs leads to faster, more accurate, and more robust inferences than one can achieve with compensatory processes, and vice versa (e.g., higher cue redundancy favors noncompensatory processing whereas higher independence between cues favors compensatory processing; see Gigerenzer & Brighton, 2009; Hogarth & Karelaia, 2006; Katsikopoulos & Martignon, 2006; Martignon & Hoffrage, 2002).

Ideally, research proceeds by first identifying environments in which a noncompensatory heuristic is ecologically rational, and then testing whether people rely on that heuristic in this environment or switch to compensatory strategies when the environment is changed accordingly (e.g., Dieckmann & Rieskamp, 2007; Rieskamp & Otto, 2006). Tests of whether and when people process recognition in a noncompensatory way fall mostly into two groups. One group did not test the recognition heuristic in its domain, that is, with substantial recognition validity, inferences from memory, and natural recognition (Table 2). Instead, tests were performed in situations where the recognition validity was unknown or could not be determined (e.g., Richter & Späth, 2006, Experiment 1; Oppenheimer, 2003, Experiments 1 and 2), in which recognition was not natural but induced by the experimenter (e.g., Bröder & Eichler, 2006; Newell & Shanks, 2004, Experiments 1 and 2), in which inferences were made from givens (e.g., Newell & Shanks, 2004, Experiments 1 and 2) or cue values were provided for unrecognized objects (Glöckner & Bröder, 2011). The second group tested noncompensatory processing of recognition in its proper domain. One of the first was an experiment by Richter and Späth (2006, Experiment 3), which we briefly review here given that it has been incorrectly presented as evidence against noncompensatory processing.

Richter and Späth asked whether the recognition heuristic would predict inferences in the presence of a strong, contradicting cue. German participants were taught whether certain recognized American cities have international airports or not. The airport cue was chosen as being the most valid (mean subjective validity = .82) among six cues tested in a pilot study. Moreover, the biserial rank correlation between population rank and air-

Figure 2: A reanalysis of Richter & Späth's (2006) Experiment 3, which tested the noncompensatory use of recognition in inferences from memory with substantial recognition validity. Each bar represents one participant, and its height the number of inferences (out of a total of 32) consistent with the recognition heuristic. The upper panel shows how often each participant judged a recognized city as larger than an unrecognized one when they were told that the recognized city had an international airport (positive cue). The middle panel shows the same when the participants had no information about whether the city had an international airport (no cue). The lower panel shows the critical test in which participants were told that the recognized city had no such airport (negative cue). Even in this critical test, the majority of participants made nearly every inference in accordance with the recognition heuristic. In contrast to this reanalysis, Richter and Späth (2006) did not report their individual data and concluded from the group means (98%, 95%, and 82% of the choices consistent with the recognition heuristic) that there is "no evidence was found in favor of a noncompensatory use of recognition" (see text).
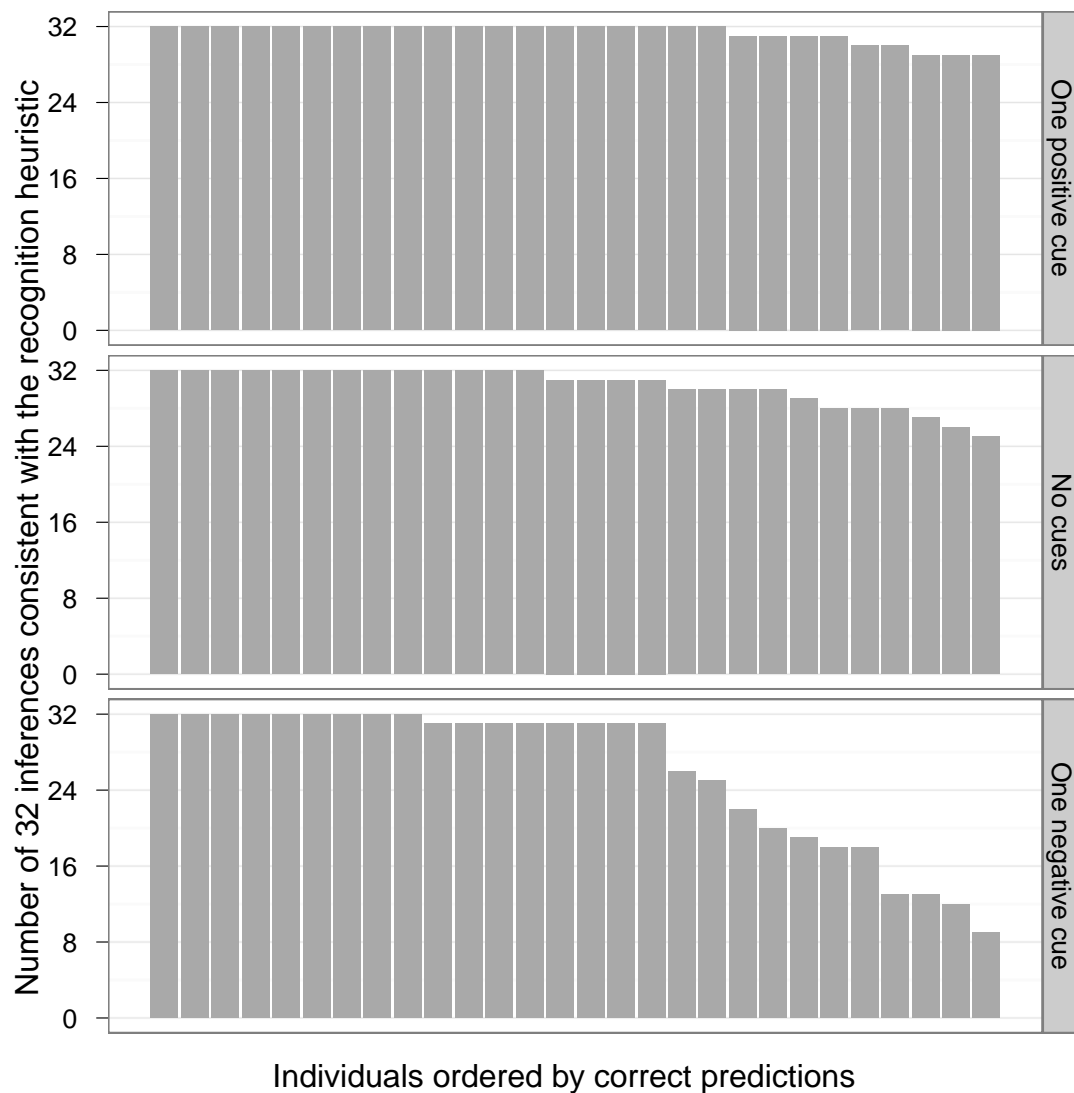
port was larger than that between population rank and recognition, $-.71$ versus $-.56$. There were three memory states for recognized cities: positive cue (with international airport), no cue (unknown), and negative cue (no international airport). Richter and Späth reported that in these three states, 98%, 95%, and 82% of the inferences were in accordance with the recognition heuristic, respectively, and they concluded that "no evidence was found in favor of a noncompensatory use of recognition" (p. 159). Puzzled by that conclusion, which was based on averages, we asked the authors for the individual data, which they cordially provided and which are shown in Figure 2. These data show that, in the presence of a strong contradicting cue (lower panel), the *majority of people* chose the recognized objects 97% to 100% of the time, as predicted by the recognition heuristic, while the others appeared to guess or follow some other strategy. This pattern was intra-individually highly consistent, with zero or one deviations out of 32 judgments per participant, a degree of consistency rarely obtained in judgment and decision-making research.

Pachur et al. (2008) reviewed the literature and found similar results to those in Figure 2. They concluded that, when the recognition validity is high, inferences from memory are frequently consistent with a noncompensatory use of recognition, even in the presence of conflicting cues. In the authors' own study, participants had knowledge of three conflicting (negative) cues indicating that the recognized object should have a small criterion value; nevertheless, about half of the participants chose the recognized object in every single trial. Individual differences were similar to those in Figure 2 (lower panel). Note that if it were true that most people consistently made trade-offs between recognition and opposing-valued cues or sets of cues *that have a higher validity than the recognition validity*, then in such situations most people should exhibit about 0% accordance. However, such individuals are not observed in Figure 2. Only a few were observed in Pachur et al.'s experiments (2008, p. 195) and in their reanalysis of Newell & Fernandez (2006).

Similarly, in studies on the role of name recognition in political forecasts, most voters always behaved in accordance with the recognition heuristic, whether or not there was a strong conflicting cue present (Marewski et al., 2010). As Table 2 shows, individual analyses reveal that a large proportion of participants consistently made inferences in accordance with the recognition heuristic, even with up to three conflicting cues.

How to model the people who deviate from the predictions of the recognition heuristic? A common proposal in the literature has been that these people integrate recognition information with other cues in a compensatory fashion. But, to our knowledge, in none of these articles was a compensatory model formulated and tested against the recognition heuristic. Testing such models is essential to theorizing for several reasons. First, if some individuals do not accord with the recognition heuristic, this does not logically imply that they rely on a compensatory process. They might simply guess, or rely on the best cue beyond recognition, as in a lexicographic rule, and thus adopt a different noncompensatory process. Second, because no model can explain all behavior, one needs to show that there are others that can explain more.

We know of only one study that has formulated compensatory models and tested them against the recognition heuristic (Marewski et al., 2010).[1] The five alternatives integrate recognition with further cues for the recognized object (Table 2). The alternative models had free parameters that allowed them to both mimic the recognition heuristic *and* predict the opposite pattern, depending on the parameter tuning. That is, they included the recognition heuristic as a special case. Because these alternatives use free parameters and the recognition heuristic uses none, it is important to test how well the models predict (rather than fit) judgments. None of the five compensatory models could predict judgments better than the recognition heuristic, which performed the best overall. The study showed that although the recognition heuristic cannot predict with 100% accuracy, particularly in the presence of contradicting cues, this by itself does not imply that compensatory models can actually predict better.

To summarize: The recognition heuristic is a simple, noncompensatory model of inference from memory. We personally have no doubts that recognition is sometimes dealt with in a compensatory way, especially when the ecology favors doing so. A number of studies have conducted critical tests in which recognized objects with negative cue values were compared with unknown objects. The typical results were that (i) the mean accordance rates decreased when one or more negative cue values were introduced, while (ii) a large proportion of participants' judgments nevertheless accorded consistently with the recognition heuristic's predictions. Result (i) has been interpreted as implying compensatory decision making, but no compensatory models were put forth to test this claim. In contrast, the first test of five compensatory models showed that in fact none could predict people's inferences as well as the noncompensatory use of recognition.

7 Methodological principles

The previous section suggests a methodology to be followed. We summarize here two relevant principles.

---

[1]As mentioned, Glöckner & Bröder (2011) work outside the domain of the recognition heuristic in which people know the cue values of unrecognized objects.

Principle 1: Test heuristics against competing models; do not claim victory for a model that was neither specified nor tested. This principle seems obvious, but it has been routinely neglected in the study of the recognition heuristic. Among the studies that claimed to have found evidence for compensatory models, we are not aware of a single one that has actually tested such a model. Hilbig and Pohl (2009) attempted to do so, and we applaud the direction they took. They used as alternatives two large model classes, evidence-accumulation models and neural nets, which they also treated as one. Since these model classes can mimic the outcomes of the recognition heuristic, multiple regression models, as well as virtually any inferential strategy ever proposed in cognitive psychology, it is not clear to us how they derived their specific predictions from such flexible models. We ourselves have proposed and tested linear "recognition cue" models that treat recognition in a compensatory way (Gigerenzer & Goldstein, 1996). We suspect that the origin of this methodological flaw is deeply rooted in generations of researchers who have been taught that hypotheses testing amounts to null hypothesis testing, that is, rejecting a precisely stated null hypothesis in favor of an unspecified alternative hypothesis. This biased procedure is not a swift route to scientific progress (Gigerenzer et al., 1989).

Principle 2: Analyze individual data; do not base conclusions on averages only. This principle is necessary because there are systematic individual differences in cognitive strategies. These differences have been reported across the entire life span, from children's arithmetical reasoning (e.g., Shrager & Siegler, 1998), judgments of area (Gigerenzer & Richter, 1990), and Bayesian inferences (Zhu & Gigerenzer, 2006) to decision making in old age (Mata, Schooler, & Rieskamp, 2007). If individual differences exist, analyses based only on means (across individuals) do not allow conclusions about underlying processes. One simple solution is to always analyze data on the individual level, as in Figure 2, which can reveal the existence of individual differences. A more theoretically guided approach would be to specify competing models, and test what proportion of participants can be predicted by each model (Marewski et al, 2010).

8 Results we had not anticipated a decade ago

Thanks to the researchers who set out to study the recognition heuristic by means of analysis, computer simulation, and experiment, we have more than once been taught lessons by unexpected results. We cannot list all here, but briefly mention three of the many surprises.

8.1 Less-is-more effects are theoretically stronger in group decision making

Reimer and Katsikopoulos (2004) extended the role of name recognition from individual to collective decision making. If one member in a group recognizes only one of two alternatives, but the others recognize all and have some further cue knowledge, should the group follow the most ignorant member who can rely on the recognition heuristic? The authors first deduced analytically that less-is-more effects can emerge in a group context and that these effects are stronger in magnitude than in individual decisions. The conditions are similar to those we had arrived at for individual decisions: If the recognition validity is higher than the knowledge validity, both are independent of the number n of objects recognized, and some further assumptions concerning the homogeneity of the groups hold, then the relationship between accuracy and n is inversely U-shaped. That is, there should exist groups whose members recognize fewer objects but reach a higher accuracy than do groups who recognize more objects. The authors reported less-is-more-effects in group decision making in an experiment, where also a fascinating new phenomenon emerged. Consider a group of three in which one member recognized only city a, while the other two members recognized both cities and individually chose b as the larger one. What would the group decide after consulting with one another? The majority rule predicts b, yet in most cases, the final group decision was a. This result suggests that a lack of recognition has a special status not only in individual decisions, as originally proposed, but in group decisions as well.

8.2 Less-is-more effects are stronger with > 2 alternatives and positive framing

In our original work, we relied on tasks with two alternatives to deduce the size of less-is-more effects analytically. Recently, the recognition heuristic has been generalized to more than two alternatives (McCloy, Beaman, & Smith, 2008; Marewski et al, 2010). Does the less-is-more effect also hold in choices involving multiple objects? It does. An experiment on inferring who is the richest citizen in a set demonstrated less-is-more effects irrespective of whether the task was to choose among two, three, or four alternatives (Frosch, Beaman, & McCloy, 2007). Moreover, one can analytically show that this generalization of the recognition heuristic implies that the size of the effect increases when more alternatives are involved. Surprisingly, for three and more alternatives, the model implies a framing effect. If the question is framed

positively, such as "Which of the people is richest?" the less-is-more effect is more pronounced than for the question "Which of the people is poorest?" (McCloy et al., 2008). This work illustrates the importance of analytically deriving predictions from the recognition heuristic, in order to see what the model implies and what it does not.

8.3 The power of laypersons' recognition for prediction

A widely entrenched view about heuristics is that yes, people rely on them because of their limited cognitive capacities, but no, they cannot often lead to good inferences. Skeptical of the power of the recognition heuristic to yield good decisions, Serwe and Frings (2006) set out to test it in a task in which they were confident that it would fail: predicting the winners of the 127 Gentleman Singles Wimbledon tennis matches. They were skeptical for good reasons. First, tennis heroes rise and fall quickly; by the time their names have finally found a place in collective recognition memory, their prowess may already be fading. Second, athletes are best known within their home country, even if they do not perform particularly well in the international arena. Recognition of an athlete should thus be a poor guide to predicting whether he or she will win an international match. To demonstrate these suspected Achilles' heels of the recognition heuristic, Serwe and Frings needed semi-ignorant people, ideally, those who recognized about half of the contestants. Among others, they contacted German amateur tennis players, who indeed recognized on average only about half of the contestants in the 2004 Wimbledon Gentlemen's Singles tennis tournament. Next, all Wimbledon players were ranked according to the number of participants who had heard of them. How well would this "collective recognition" predict the winners of the matches? Recognition turned out to be a better predictor (72% correct) than the ATP Entry Ranking (66%), the ATP Champions Race (68%), and the seeding of the Wimbledon experts (69%). These unexpected results took the authors by surprise. When they presented their results to the ABC Research Group, the surprise was on both sides. Could it have been a lucky strike, ripe for publication in the *Journal of Irreproducible Results*? Scheibehenne and Bröder (2007) set out to test whether the findings would replicate for Wimbledon 2005—and found basically the same result. In addition, when asked to predict the match winners, the amateur tennis players predicted in around 90% of the cases that the recognized player would win. Thus, there can be powerful wisdom in lay people's collective recognition.

Collective recognition has also been used for investment in the stock market, which is reviewed in Ort-mann, Gigerenzer, Borges, and Goldstein (2008), and for forecasting sport events, as reviewed in Goldstein and Gigerenzer (2009).

9 Open questions and future research

In this article, we have addressed a number of research directions that we think are important to pursue, such as integration with theories of recognition memory and deeper understanding of the evaluation process. We close with a selection of open questions and issues.

9.1 Is the noncompensatory process implemented in the stopping rule or in the decision rule?

The definition of the recognition heuristic allows both interpretations. The classic definition of compensatory and noncompensatory processes locates the difference in the decision rule: Cues for the recognized object may or may not come to mind; the relevant question is whether they are used when making the decision. In contrast, we had made a stronger modeling assumption, namely that search for cue information in memory is stopped if only one alternative is recognized, which locates the absence of trade-offs already in the stopping rule. It is not easy to decide between these alternatives. For instance, Hilbig and Pohl (2009) reported that mean decision times were shorter when participants had further knowledge about a city as opposed to when they had none, and interpreted this difference against our interpretation of the process and in favor of an unspecified "evidence-based model". Decision time predictions, however, cannot be derived from our simple model without making additional assumptions, and highly specific ones. We do not deal with decision times extensively in the limited space of this review, but elaborate on one point here. Decision time predictions as well as recognition time (fluency) are best derived from a model of the core cognitive capacities involved (Marewski & Schooler, 2011). To illustrate, it is not correct that our interpretation implies no difference in decision times; this prediction would, among others, require that the speed of recognition (fluency) be uncorrelated with the size of the object or, in general, the criterion value. Recognition, however, tends to be faster for larger objects (Hertwig, Herzog, Schooler, & Reimer, 2008; Schooler & Hertwig, 2005). Thus, if speed of recognition is correlated with the actual size and if objects that people know more about are larger, mean decision times are likely to be shorter when additional knowledge is available (Marewski et al., 2010; Hilbig & Pohl, 2009, Experiment 3, did address this issue). No cue integration

is needed to explain this result. The question whether the noncompensatory process is located in the stopping or in the decision rule is still an open issue. The answer to this question does not concern the outcome prediction of the recognition heuristic, only the process that leads to this outcome.

9.2 How do people adapt their use of recognition to changes in recognition validity?

Two striking observations have been reported. First, whereas accordance rates are correlated with the recognition validities across tasks (see Figure 1), the individual accordance rates within a task appear to be unrelated to the individual recognition validities (Pachur & Hertwig, 2006; Pohl, 2006). This result may be due to the low mean recognition validity (.60) or the low variability in individual recognition validities (most were between .55 and .65) in Pachur and Hertwig's study, or to the use of selected rather than representative samples in some of Pohl's sets (mountains, rivers, and islands). Whatever the reasons, this observation deserves a closer investigation. Although it suggests limits to the adaptive use of recognition, a second observation suggests an enhanced adaptive use: Pohl and Hilbig (Pohl, 2006; Hilbig & Pohl, 2008) reported that the recognition heuristic fits the data better when the heuristic would lead to a correct inference than when it would lead to an incorrect one. For instance, in Hilbig and Pohl's (2008) experiment, 76% chose the recognized city in pairs when it was incorrect and 82% when it was correct. The authors interpret this slight difference in means as indicative of additional knowledge being relied on in some cases, but what this knowledge is remains unclear. It could be related to criterion knowledge, an issue that Pachur & Hertwig (2006), and Hilbig et al. (2009) have taken different sides on. What is needed is a model that can predict when this effect occurs. A clarification of these two observations will hopefully contribute to better theories about the evaluation process.

9.3 Recognition plus additional knowledge

Pohl (2006) reported that the recognition heuristic predicted inferences better on R+U pairs (comparison between a recognized object about which a person has additional knowledge [R+] and an unrecognized object [U]) than on RU pairs (R = mere recognition). The question is how to explain this difference. Pohl (2006) concluded from this result that some people use a compensatory strategy, but without specifying and testing any such a strategy. Yet this is not the only interpretation. Another is that the difference follows from systematic variations in the strength of the recognition signal and the recognition

validity (Marewski et al., 2010; Marewski & Schooler, 2011).

9.4 Recognition and preference formation

Although we formulated the recognition heuristic as a model for inferences, it can also serve as a model for preferences. Consider consumer choice, in which the classical model of brand preference is a formalization of Fishbein's (1967) work on beliefs and attitudes:

$$A_b = \sum_{i=1}^{N} W_i B_{ib} \qquad (2)$$

where A_b = the attitude toward brand b, W_i = the weight of the ith product attribute, B_{ib} = the consumer's belief about brand b where attribute i is concerned, and N = the number of attributes deemed important for choosing a brand.

The resemblance to the weighted linear models studied in judgment and decision-making research is clear. With weights and beliefs that are typically elicited from the decision maker, such models do a good job in fitting consumers' brand choices for orange juice, lipstick, and the like (Bass & Talarzyk, 1972). However, *what* people choose is different from *how* people choose, as those studying decision processes have noticed. We illustrate this here with *noncompensatory screening* and *halo effects*.

Before choosing products, consumers often reduce a large number of possible alternatives to a smaller set, which they inspect more closely. Such "consideration sets" turn out to be excellent predictors of what is ultimately chosen (Shocker, Ben-Akiva, Boccara, & Negungadi, 1991; Hauser, 1978). Although these considerations sets can in theory be created by compensatory multiattribute procedures that integrate all available information (Roberts & Lattin, 1991), studies suggest that products are filtered into a consideration set by means of noncompensatory heuristics (Gilbride & Allenby, 2004; Laroche, Kim, & Matsui, 2003; Payne, 1976; Bettman & Park, 1980). The generalization of the recognition heuristic to the domain of preferences and multi-alternative choice enables its use as a building block in consideration set formation (Marewski et al., 2010). Recognition-based consideration sets facilitate decisions when the initial choice set is large. Sometimes recognition itself is the desirable attribute, as when students choose universities.

Further deviations from the classical linear model of brand preference have been suggested by the presence of halo effects, that is, the tendency for people who favor a brand to evaluate it positively on all attributes and those who dislike it to do the opposite (Beckwith & Lehmann, 1975). Such behavior suggests that the expressed beliefs about attributes may themselves be inferences, as

opposed to the result of recall from memory. Beliefs about the attributes of unrecognized brands cannot be stored in memory and must be constructed on the fly. Extending beyond the original domain of the recognition heuristic, one exciting possibility is that the effect of recognition on attribute beliefs can be even stronger than that of direct experience. Experimental studies on food choice indicate not only that people buy the products they recognize but that brand recognition often dominates other cues to a degree that can change the perception of the product. For instance, in a blind taste test, most people preferred a jar of high-quality peanut butter over two alternative jars with low-quality peanut butter. Yet when one familiar and two unfamiliar brand labels were randomly assigned to the jars, preferences changed. When the high-quality product was in a jar with an unknown brand name, it was preferred only 20% of the time, whereas the low-quality product in the jar with the familiar brand name was chosen 73% of the time. When the exact same peanut butter was put into three jars, one showing a familiar brand name and two showing unfamiliar brand names, the (faux) familiar product won the taste test 75% of the time (Hoyer & Brown, 1990; see also Macdonald & Sharp, 2000). One way to interpret this result is that, for the majority of consumers, brand name recognition dominates the taste cues present in the blind test. But there is an interesting alternative to this noncompensatory processing hypothesis. The taste cues themselves might be changed by name recognition—people "taste" the brand name. Such a process could be modeled in a similar way as the change of perceived cue values in the RAFT model of hindsight bias (Hoffrage, Hertwig, & Gigerenzer, 2000). This interpretation—like the halo effect—suggests a model in which recognition imputes or changes attribute values themselves. Such a process is likely to occur when cue values are direct subjective experiences such as tastes, which are neither presented as propositions nor retrieved from memory. It would provide an alternative hypothesis for the processing of brand name recognition, one that is not based on the distinction between noncompensatory and compensatory processes.

9.5 Do humans and other animals share common heuristics?

Behavioral biologists have documented in detail the rules of thumb (their term for heuristics) that animals use for choosing food sites, nest sites, or mates. For instance, Stevens & King (in press) discuss how animals use simple heuristics for recognizing kin to facilitate cooperation, and Shaffer, Krauchunas, Eddy, & McBeath (2004) report that dogs, hoverflies, teleost fish, sailors, and baseball players rely on the same heuristics for intercepting prey, avoiding collisions, and catching balls. Biology of-

fers numerous, specific examples, but no systematic theory of heuristics, such as in terms of rules for search, stopping, and decision (Hutchinson & Gigerenzer, 2005). If we can find signs of the same rules across species, this might provide converging evidence for specific models of heuristics. The recognition heuristic seems to be a good candidate. For instance, rats and mice prefer foods they recognize from having tasted or from having smelled on the breath of fellow rats, a tendency known as *neophobia*. They may also rely on recognition to infer which of several foods made them sick. In one experiment, Norway rats were fed two foods. Both were relatively novel, but one was familiar from the breath of a fellow rat. After these rats were given a nauseant, they subsequently avoided the food they did not recognize from the neighbor's breath (Galef, 1987). As in the experiments with humans, one can test whether recognition is overruled by a powerful cue. Consider a similar situation, where one food is recognized from the breath of a fellow rat, but now the fellow rat is also (experimentally made to appear) sick at the time its breath is smelled. Surprisingly, observer rats still chose the recognized food from the breath of the sick neighbor (Galef, McQuoid, & Whiskin, 1990). As in humans, accordance rates were not 100%, but around 80%. Thus, recognition appears to overrule the sickness cue that advises against selecting the recognized food.

The question of which heuristics humans and other animals share has been recently discussed in a target article in *Behavioural Processes* (Hutchinson & Gigerenzer, 2005) and commentaries (e.g., Cross & Jackson, 2005; Shettleworth, 2005), although we know of no systematic research on the topic. Some comparative research has focused on common biases, but few researchers have tested models of common heuristics. The recent dialogue with psychologists studying the adaptive toolbox has led biologists to revisit the fundamental question of how to model behavior. Models of heuristics are not simply approximations to optimizing models; rather, they allow scientists to study behavior in uncertain, complex worlds as opposed to the certain, small worlds required for the ideal of optimization. "Although behavioral ecologists have built complex models of optimal behaviour in simple environments, we argue that they need to focus on simple mechanisms that perform well in complex environments" (McNamara & Houston, 2009, p. 670).

10 Conclusion

The recognition heuristic is a simple model that can be put to many purposes: describing and predicting inferences and preferences, and forecasting such diverse events as the outcomes of sporting events and elections. The research on the recognition heuristic has promoted

the use of testable models of heuristics (instead of vague labels), and of simple models in which each parameter can be directly measured rather than fitted. With such precise models, one can easily observe when a heuristic makes correct predictions and when it fails. But the emerging science of heuristics also caused unease in some research communities, breaking with cherished ideals such as general-purpose models of cognition, the assumption of a general accuracy-effort tradeoff, and the conviction that heuristics are always inferior to complex strategies. This may well explain why every critique we know of the recognition heuristic claims that minds add and weigh cues; none has proposed and tested a different, perhaps simpler model. As the last decade has shown, however, there is clear evidence that this simple model consistently predicts the judgments of a substantial proportion of individuals, even in the presence of contradicting cues. Moreover, research now suggests that people may use heuristics in an adaptive way, as witnessed in the substantial correlation between recognition validities and accordance rates. We thank all fellow researchers and critics for carving out the details of the adaptive toolbox, and thus contributing, in the words of Herbert Simon (1999), "to this revolution in cognitive science, striking a great blow for sanity in the approach to human rationality."

References

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

Bass, F. M., & Talarzyk, W. W. (1972). An attitude model for the study of brand preference. *Journal of Marketing Research, 9,* 93–96.

Beckwith, N. E., & Lehmann, D. R. (1975). The importance of halo effects in multi-attribute attitude models. *Journal of Marketing Research, 12*, 265–275.

Bettman, J. R., & Park, C. W. (1980). Effects of prior knowledge and experience and phase of the choice process on consumer decision analysis: A protocol analysis. *Journal of Consumer Behavior, 7,* 234–248.

Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & R. 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.

Brighton, H., & Gigerenzer, G. (2011). Towards competitive instead of biased testing of heuristics: A reply to Hilbig & Richter (2011). *Topics in Cognitive Science, 3*, 197–205.

Bröder, A., & Eichler, A. (2006). The use of recognition information and additional cues in inferences from

memory. *Acta Psychologica, 121,* 275–284.

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

Bröder, A., & Schütz, J. (2009). Recognition ROCs are curvilinear - or are they? On premature arguments against the two-high-threshold model of recognition. *Journal of Experimental Psychology: Learning, Memory & Cognition, 35,* 587–606.

Cross, F. R., & Jackson, R. J. (2005). Spider heuristics. *Behavioural Processes, 69,* 125–127.

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.

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

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition, 35,* 1801–1813.

Fishbein, M. (1967). A consideration of beliefs and their role in attitude measurement. In M. Fishbein (Ed.), *Readings in attitude theory and measurement* (pp. 389–400). New York: Wiley.

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.

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

Galef, B. G., Jr. (1987). Social influences on the identification of toxic food by Norway rats. *Animal Learning & Behavior, 15,* 327–332.

Galef, B. G., Jr., McQuoid, L. M., & Whiskin, E. E. (1990). Further evidence that Norway rats do not socially transmit learned aversions to toxic baits. *Animal Learning and Behavior, 18,* 199–205.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science, 1,* 107–143.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103,* 650–669.

Gigerenzer, G., Hell, W., & Blank, H. (1988). Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 513–525.

Gigerenzer, G., Hertwig, R., van den Broek, E., Fasolo, B., & Katsikopoulos, K. (2005). "A 30% chance of rain

tomorrow:" How does the public understand probabilistic weather forecasts? *Risk Analysis, 25,* 623–629.

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

Gigerenzer, G., & Richter, H. R. (1990). Context effects and their interaction with development: Area judgments. *Cognitive Development, 5,* 235–264.

Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Krüger, L. (1989). *The empire of chance. How probability changed science and everyday life.* Cambridge, United Kingdom: Cambridge University Press.

Gilbride, T. J., & Allenby, G. M. (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science, 23,* 391–406.

Glöckner, A. & Bröder, A.. (2011). Processing of recognition information and additional cues: A model-based analysis of choice, confidence, and response time. *Judgment and Decision Making, 6,* 23–42.

Goldstein, D. G. (1997). *Models of bounded rationality for inference.* Doctoral thesis, The University of Chicago. Dissertation Abstracts International, 58(01), 435B. (University Microfilms No. AAT 9720040).

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. 37–58). New York: Oxford University Press.

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

Goldstein, D. G., & Gigerenzer, G. (2009). Fast and frugal forecasting. *International Journal of Forecasting, 25,* 760–772.

Goldstein, D. G., Gigerenzer, G., Hogarth, R. M., Kacelnik, A., Kareev, Y., Klein, G., et al. (2001). Group report: Why and when do simple heuristics work? In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 173–190). Cambridge, MA: MIT Press.

Hauser, J. R. (1978). Testing the accuracy, usefulness, and significance of probabilistic models: An information-theoretic approach, *Operations Research, 26,* 406–421.

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., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences, 24,* 383–451.

Hilbig, B. E., & Pohl, R. F. (2008). Recognizing users of the recognition heuristic. *Experimental Psychology, 55,* 394–401.

Hilbig, B. E., & Pohl, R. F. (2009). Ignorance- vs. evidence-based decision making: A decision time analysis of the recognition heuristic. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 1296–1305.

Hilbig, B. E., Pohl, R. F., & Bröder, A. (2009). Criterion knowledge: A moderator of using the recognition heuristic? *Journal of Behavioral Decision Making, 22,* 510–522.

Hoffrage, U. (1995). *Zur Angemessenheit subjektiver Sicherheits-Urteile: Eine Exploration der Theorie der probabilistischen mentalen Modelle* [The adequacy of subjective confidence judgments: Studies concerning the theory of probabilistic mental models]. Unpublished doctoral dissertation, University of Salzburg, Austria.

Hoffrage, U. (2011). Recognition judgments and the performance of the recognition heuristic depend on the size of the reference class. *Judgment and Decision Making, 6,* 43–57.

Hoffrage, U., Hertwig, R., & Gigerenzer, G. (2000). Hindsight bias: A by-product of knowledge updating? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 566–581.

Hogarth, R. M. (1980). *Judgement and choice: The psychology of decision.* Chichester, United Kingdom: John Wiley & Sons.

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

Hogarth, R. M., & Karelaia, N. (2006). "Take-the-best" and other simple strategies: Why and when they work "well" with binary cues. *Theory and Decision, 61,* 205–249.

Hoyer, W. D., & Brown, S. P. (1990). Effects of brand awareness on choice for a common, repeat purchase product. *Journal of Consumer Research 17,* 141–148.

Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes, 69,* 97–124.

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

Kahneman, D., & Tversky, A. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty:*

Heuristics and biases (pp. 84–98). New York: Cambridge University Press.

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

Laroche, M., Kim, C., & Matsui, T. (2003). Which decision heuristics are used in consideration set formation? *Journal of Consumer Marketing, 3*, 192–209.

Macdonald, E., & Sharp, B. (2000). Brand awareness effects on consumer decision making for a common, repeat purchase product: A replication. *Journal of Business Research, 48*, 5–15.

Malmberg, K. J. (2002). On the form of ROCs constructed from confidence ratings. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 380–387.

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.

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.

Marewski, J. N., & Schooler, L. J. (2011). *How memory aids strategy selection.* Unpublished manuscript.

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.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Lexicographic heuristics for paired comparison. *Theory and Decision, 52,* 29–71.

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.

McCloy, R., Beaman, C. P., & Smith, P. T. (2008). The relative success of recognition-based inference in multichoice decisions. *Cognitive Science, 32,* 1037–1048.

McGrath, R. E. (2008). Predictor combination in binary decision-making situations. *Psychological Assessment, 20,* 195–205.

McNamara, J. M., & Houston, A. I. (2009). Integrating function and mechanism. *Trends in Ecology and Evolution, 24,* 670–675.

Mischel, W. (2006). Bridges toward a cumulative psychological science. In Paul A. M. Van Lange (Ed.) *Bridging social psychology: Benefits of transdisciplinary approaches* (p. 437–446). Mahwah, NJ: Lawrence Erlbaum.

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

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

Ortmann, A., Gigerenzer, G., Borges, B., & Goldstein, D. G. (2008). The recognition heuristic: A fast and frugal way to investment choice? In C. R. Plott & V. L. Smith (Eds.), *Handbook of experimental economics results: Vol. 1 (Handbooks in Economics No. 28)* (pp. 993–1003). Amsterdam: North-Holland.

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., Mata, R., & Schooler, L. J. (2009). Cognitive aging and the adaptive use of recognition in decision making. *Psychology and Aging, 24*, 901–915.

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

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

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

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

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the

time course of retrieval. *Cognitive Psychology, 21,* 139–155.

Reimer, T., & Katsikopoulos, K. 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,*1501–1562.

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

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

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

Serwe, S., & Frings, C. (2006). Who will win Wimbledon? The recognition heuristic in predicting sports events. *Journal of Behavioral Decision Making, 19*, 321–322.

Shaffer, D. M., Krauchunas, S. M., Eddy, M., & McBeath, M. K. (2004). How dogs navigate to catch Frisbees. *Psychological Science, 15,* 437–441.

Shettleworth, S. J. (2005). Taking the best for learning. *Behavioural Processes, 69,* 147–149.

Shocker, A. D., Ben-Akiva, M., Boccara, B., & Nedungadi, P. (1991). Consideration set influences on consumer decision making and choice: Issues, models, and suggestions, *Marketing Letters, 2,* 181–197

Shrager, J., & Siegler, R. S. (1998). SCADS: A model of strategy choice and strategy discovery. *Psychological Science, 9,* 405–410.

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

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

Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (Vol. 2, pp. 397–420). New York: Cambridge University Press.

Slovic, P., Monahan, J., & MacGregor, D. G. (2000). Violence risk assessment and risk communication: The effects of using actual cases, providing instruction, and employing probability versus frequency formats. *Law and Human Behavior, 24,* 271–296.

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

Stevens, J. & King, A. (in press). The life of others: Social rationality in animals. In Hertwig, R., Hoffrage, U., the ABC Research Group. *Simple heuristics in a social world.* New York: Oxford University Press.

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.

Volz, K. G., Schubotz, R. I., Raab, M., Schooler, L. J., 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.

Wallsten, T. S. (1983). The theoretical status of judgmental heuristics. In R. W. Scholz (Ed.), *Decision making under uncertainty* (pp. 21–39). Amsterdam: Elsevier.

Wason, P. C. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology, 23,* 63–71.

Yonelinas, A. P., Otten, L. J., Shaw, K. N., & Rugg, M. D. (2005). Separating the brain regions involved in recollection and familiarity in recognition memory. *Journal of Neuroscience, 25,* 3002–3008.

Zhu, L. & Gigerenzer, G. (2006). Children can solve Bayesian problems: The role of representation in mental computation. *Cognition, 98,* 287–308.

Appendix, Table 2: An overview of experimental studies on the recognition heuristic (RH) reporting mean correct predictions (accordance rates). Three plusses in Columns 4–6 mean that the domain was one for which the recognition heuristic was proposed as a model: α = substantial recognition validity; Mem = Inferences from memory (as opposed to inferences from givens); Nat = natural recognition (as opposed to experimentally induced). Studies that satisfy these three conditions are listed first; others follow. (RU) = comparison between a recognized object (R) and an unrecognized object; (R$^+$U) = comparison between a recognized object about which a person has additional knowledge (R$^+$) and an unrecognized object.

Article/Exp.	Reference class	Criterion	Task in the domain of RH?			Alter-native model?	Individual analysis?	Mean correct predictions of judgments by RH (and recognition validity α)
			α	Mem	Nat			
Frosch et al. 2007	Richest individuals in UK	Richest Poorest	+ +	+ +	+ +	no no	no no	(r=.73): 79% (r=.73): 83%
Marewski et al. 2009	Political parties	Election result	+	+	+	yes	yes	87%-89% For experimentally induced recognition, RH predicts the forecasts of 49 voters, fluency heuristic of 12, and both equally of 4 voters.
Marewski et al. 2010/1	Political candidates	Election result	+	+	+	yes	Most voters always followed RH, independent of whether there was a conflicting cue	(α=.80): 87% 1 negative cue: 73% RH predicted better than compensatory models tested.
Marewski et al. 2010/2	Political parties	Election result	+	+	+	yes	no	(α=.92): 89% (R$^+$U) (α=.68): 62% (RU) RH predicted better than compensatory models tested.
Marewski et al. 2010/3	Largest cities	Population	+	+	+	yes	Most voters always followed RH, independent of whether there was additional knowledge (R$^+$)	(α=.81): 96% (R$^+$U) (α=.74): 86% (RU) RH predicted better than compensatory models tested.

Article/Exp.	Reference class	Criterion	Task in the domain of RH? α Mem Nat			Alter-native model?	Individual analysis?	Mean correct predictions of judgments by RH (and recognition validity α)
Marewski et al. 2010/4	Political parties	Election result	+	+	+	yes	For 25 out of 26 participants, RH predicted better than 2 compensatory models	RH predicted better than compensatory models tested.
Pachur & Biele 2007	2004 European Soccer Championship	Winner	+	+	+	yes	no	(α=.71): 91% RH predicted better than 4 alternative models.
Pachur et al. 2008/1	British cities	Population	+	+	+	no	More negative than positive cues: 60% always (16 out of 16 trials) followed RH	(α=.70): 96% No effect of up to 3 negative/positive cues on RH accordance
Pachur et al. 2008/2	British cities	Population	+	+	+	no	3 negative cues: 46% always (18 out of 18 times) followed RH	(α=.72): 94% 3 negative cues: 85%
Pachur et al. 2008/3	British cities	Population	+	+	+	no	3 negative cues: 48% always (10 out of 10 times) followed RH	(α=.71): 96% 3 negative cues: 93%
Pachur et al. 2009	Infectious diseases; U.S. cities	Frequency Population	+ +	+ +	+ +	no no	no no	Young adults: (α=.90): 95% (α=.62): 64% Elderly: (α=.92): 96% (α=.60): 71%
Pohl 2006/2	Largest Swiss cities mixed with ski resorts	Population	+	+	+	no	no	(α=.72): 75%
Pohl 2006/3	Italian cities Belgian cities	Population	+ +	+ +	+ +	no no	no no	(α=.82): 89% (α=.89): 88%

Article/Exp.	Reference class	Criterion	Task in the domain of RH? α Mem Nat			Alter-native model?	Individual analysis?	Mean correct predictions of judgments by RH (and recognition validity α)
Richter & Späth 2006/3	Largest U.S. cities	Population	+	+	+	no	no Reanalysis (see Figure 1): Majority consistently followed RH	(rank corr.= −.56) 1 positive cue: 98% no cue: 95% 1 negative cue: 82%
Scheibehenne & Bröder 2007	Wimbledon Gentlemen Singles 2005	Winner	+	+	+	no	no	Amateurs (α=.71): 89% Laypeople (α=.69): 79%
Serwe & Frings 2006	Wimbledon Gentlemen Singles 2003	Winner	+	+	+	no	no	Amateurs (α=.73): 93% Laypeople (α=.67): 88%
Snook & Cullen 2006	NHL hockey players	Career point	+	+	+	no	no	(α=.87): 96%
Volz et al. 2006	Cities from 7 countries	Population	+[1]	+	+	no	no	(α=.63): 84%
Pohl 2006/4	Mountains Rivers Islands	Height Length Area	− + +	+ + +	+ + +	no no no	no no no	(α=.49): 89% (α=.74): 94% (α=.85): 81%
Oppenheimer 2003/1	Undefined (real and fictional cities)	Population	−	+/−	+/−	no	no	α cannot be determined Adherence: < chance
Oppenheimer 2003/2	Undefined (real and fictional cities)	Population	−	+/−	+/−	no	no	α cannot be determined Adherence: < chance
Richter & Späth 2006/2	Largest airlines in the world	Safety	−[2]	+	+	no	no	α unknown. 3 positive cues: 98% 2 positive/1neg cue: 88% 2 negative/1pos cue: 81% 3 negative cues: 67%
Bröder & Eichler 2006	Unknown small towns	Population	−	+	−	no	no	α cannot be determined; reported is validity of repetition (v). (v=.80): 75% (v=.65): 67%
Newell & Shanks 2004/1	Undefined (fictional company names)	Investment decision	−	−	−	no	yes	α cannot be determined; reported is validity of repetition (v). (v=.80): 88% (v=.65): 62%
Newell & Shanks 2004/2	Undefined (fictional company names)	Investment decision	−	−	−	no	yes	α cannot be determined; reported is validity of repetition (v). (v=.60): 66%

[1] low recognition validities (α) introduced to study decisions against RH.

[2] recognition validity not reported; reported was a partial $r = -.29$ between recognition and number of fatalities when year established was controlled for.