Using Recognition in Multi–Attribute Decision Environments

Don van Ravenzwaaij (d.vanravenzwaaij@unsw.edu.au),
Ben R. Newell (ben.newell@unsw.edu.au), and **Chris P. Moore (christophermoore@gmail.com)**
School of Psychology,
University of New South Wales,
Sydney 2052, Australia

Michael D. Lee (mdlee@uci.edu)
Department of Cognitive Sciences,
University of California Irvine,
Irvine, CA, 92697–5100, USA

Abstract

An experiment examined the effect of 'pure' recognition — in the absence of concomitant evaluation — on inferences. In the first stage of the experiment, participants indicated whether they recognized a number of Italian and US cities. In the second stage, they decided which of two cities had the larger population. Crucially, names of the cities were not available in the second stage, but participants could find out whether they had recognized them (yes/no) in the first stage of the experiment (i.e., pure recognition). Additional predictive cues (e.g., presence/absence of a university) were also available. Participants used the recognition cue about 50% of the time, rarely examined it first, and used it differently as a function of whether recognition information was binary or continuous. Furthermore, participants used the recognition cue more often if they recognized more items, irrespective of its predictive validity. Implications for theoretical frameworks that view inference as driven by discrete heuristics or processes of evidence–accumulation are briefly discussed.

Keywords: Inference, heuristics, recognition, decision making.

Humans are decision makers. Throughout our lifes, we are constantly confronted with situations that force us to make a choice. Whether it is a preference decision, "Do I take the car or do I walk to work?", or a knowledge decision, "Which soccer team scored more goals last season, Borussia Dortmund or Bayern München?", we are evaluating alternatives.

Gigerenzer and colleagues have proposed a number of relatively simple *heuristics* that could help us making such decisions. In this paper, we will focus on one of the most prominent examples: the recognition heuristic (Goldstein & Gigerenzer, 1999, 2002; Gigerenzer & Goldstein, 2011).

In the original conceptualisation of the recognition heuristic, called take–the–best (TTB; Gigerenzer & Goldstein, 1996), the first step in deciding which of two response options to choose was to use recognition. So, if a decision maker knows the Bayern München soccer team, but has never heard of Borussia Dortmund, then respond that Bayern München scored more goals last season. When both teams are recognized (thus disabling the use of recognition) the heuristic consults relevant information, or *cues*, in memory that are indicative of the number of goals scored (e.g., "What was the team's final standing in the national competition?"). These cues should be consulted in descending order of informativeness, starting with the cue that will be most indicative of the criterion of interest (i.e., number of goals scored). Cue search stops when the decision maker examines a cue that points in one direction (i.e., Borussia Dortmund was first last season, Bayern München was second, so respond Borussia Dortmund).

This proposal for a simple mechanism based on recognition sparked a wide ranging debate about the plausibility, empirical validity, and generality of the recognition heuristic (for recent examples see the papers in the three special issues of the Journal of Judgment and Decision Making — Vol 6 (1) & (5), 2011; Vol 5 (4), 2010). Much of the debate revolves around some key assumptions about the nature and operation of recognition in inferential judgment.

In the paper that introduced the recognition heuristic as a stand–alone 'tool' (i.e. not just the first step in Take–the–Best), Goldstein and Gigerenzer (2002) assume, firstly, that recognition is binary. That is, we either recognize something, or we do not, and there is no room within the heuristic for the distinction between something being vaguely familiar and something being very familiar. Secondly, recognition is assumed to be noncompensatory. That is, when we recognize one option, but do not recognize the other, then we should always go with the recognized option, regardless of any additional information. Lastly, Goldstein and Gigerenzer (2002) make a distinction between familiarity and recognition: "The term familiarity is typically used in the literature to denote the degree of knowledge (or amount of experience) a person has of a task or object. The recognition heuristic, in contrast, treats recognition as a binary, all–or–none distinction; further knowledge is irrelevant." (pp. 77). Thus, according to a strict interpretation of the (2002 version of the) recognition heuristic, when deciding whether an Italian city you know has a larger population than an Italian city you do not know, it makes no difference whether the city you do know is Rome or Pisa.

All three of these assumptions have been roundly challenged in the literature on both empirical (e.g., Pohl, 2006; Newell & Shanks, 2004; Newell & Fernandez, 2006) and theoretical grounds (e.g., Hilbig, 2010; Newell, 2011). Responding to some of these critiques, Gigerenzer and Goldstein (2011) recast the adaptive use of the recognition heuristic as involving a two–step process: first recognition ("Do I

recognize one object but not the other?") and second, evaluation ("If so, is it reasonable to rely on the recognition heuristic in this situation?"). A view consistent with that is outlined in Newell and Shanks (2004).

While such a conceptualisation is undoubtedly more plausible, it makes the claims about the way recognition aids inference that much more difficult to define and test empirically. Perhaps the trickiest aspect of the problem is that recognition almost always entails further information about the recognized object. If you have heard of Pisa, it is highly likely that you know something else about it (e.g., that it has a leaning tower) which may or may not be relevant to the criterion of interest, in this case population (cf. Oppenheimer, 2003). In other words, it is difficult to isolate the influence of 'pure' recognition — how useful is just knowing that I recognize an object for drawing an inference?

Isolating this 'pure' recognition — recognition without concomitant evaluation — is important because it can shed new light on the distinction between recognition and familiarity, and the extent to which people will rely on recognition even when they cannot directly evaluate their reason(s) for recognising an object. In order to isolate pure recognition we introduced a novel element to the standard task in which participants decide which of two cities has the larger population (e.g., Goldstein & Gigerenzer, 2002). In our task we created the distinction between recognition and familiarity alluded to by Goldstein and Gigerenzer (2002) by first asking participants to provide recognition data about a pool of response options (city names). We then presented participants with a series of forced choice decisions between two cities about which different pieces of information could be obtained (e.g., presence/absence of a university), but for which the city names were unavailable. Although participants could not discover the names they could — crucially — discover whether or not they had recognized one, both or neither of the cities when they had been presented in the first stage of the experiment. This information was available in the same manner as all the other cues — that is via clicking on relevant buttons ("Did you recognize this city when you were shown its name?") to reveal a yes/no answer (see Figure 1).

A key question here is: How often and when do participants examine the 'pure' recognition cue when drawing an inference? Will recognition remain a primary driver of decisions (cf., Pachur & Hertwig, 2006) even in the absence of evaluation? In a sense, the use of recognition in this task allows us to gain insight into participants' meta–cognitions about the usefulness of recognition in different environments. For example, do decision makers use recognition more often as they recognize more items in a pool of response categories, irrespective of the informativeness of recognition? To facilitate examining these questions we presented each participant with two decision environments in which we assumed he or she would know a different proportion of the items, thereby allowing us to directly compare response strategies: a US cities environment and an Italian cities environment.

An additional feature of the experiment was that we offered participants (between–subjects) the opportunity to use recognition as a binary (yes/no) or a continuous (slider from 0 to 100) cue. If it is true that recognition operates in a binary fashion, participants should only use the endpoints of a scale when asked to give a continuous rating of their recognition. Similarly, the usage of the recognition cue should not differ between a condition in which it was indicated as binary and a condition in which it was indicated as continuous.

In the next section, we will describe the experiment and each of its conditions in greater detail. Then, we will present some results and conclude on both the tenability of the recognition heuristic, and the use of recognition as an aid to inference more generally.

Method

Participants

All participants were first year undergraduate students at the University of New South Wales who participated in return for course credit. A total of 100 participants (62 females, 38 males), aged 17 to 39 (mean = 19.5, SD = 2.9) took part in the experiment. They were randomly divided between four between–subject conditions ($n = 25$ each).

Material

The tasks we used for this experiment are based on the German cities task (Gigerenzer & Goldstein, 1996) in which over the course of consecutive trials, a participant has to decide which of two cities has the larger population. This decision can be made by extracting information on different cues in any order. An example of such a cue could be "Is this city the national capital?". Rather than using German cities, we administered an Italian cities environment and a US cities environment to each participant (see Lee & Zhang, 2012), expecting to see higher recognition percentages for the US cities than for the Italian cities.

Table 1: *The nine cues as used in the Italian and US cities environments. Env = Environment, Val = Cue Validity, Dis = Cue Discriminability.*

Env	Nr.	Cue	Val	Dis
Italy	1	Is the city the national capital?	1	0.04
	2	Does the city have a railway station?	0.92	0.36
	3	Is the city a regional capital?	0.84	0.38
	4	Does the city have a football team in the Serie A league?	0.81	0.36
	5	Does the city have a university?	0.80	0.55
	6	Does the city have an airport?	0.76	0.49
	7	Does the city have a football team in the Serie B league?	0.70	0.30
	8	Is the city in the Po Valley?	0.60	0.52
	9	Did you recognize this city when you were shown its name?	varies	varies
US	1	Does the city have an airport?	0.78	0.51
	2	Does the city have a sport team?	0.74	0.53
	3	Does the city have a metro?	0.74	0.23
	4	Does the city have an exposition site?	0.73	0.26
	5	Is the city the national capital?	0.67	0.03
	6	Does the city have a railway station?	0.66	0.35
	7	Is the city a state capital?	0.59	0.34
	8	Did you recognize this city when you were shown its name?	varies	varies

In the first stage of the experiment, participants indicate

whether or not they recognize each of the cities used in the subsequent stage of the experiment. In a between participants manipulation, recognition was either measured as a binary or as a continuous variable. Taking the Italian environment as an example, in the binary condition participants were asked "Do you recognize this city in Italy" for a total of 66 Italian cities with response options "yes" or "no". In the continuous condition, participants were asked "How well do you recognize this city in Italy". Answers were indicated on a slider going from 0 ("I am certain that I do not recognize this city") through 50 ("I am not sure whether or not I recognize this city") to 100 ("I am certain that I recognize this city").

In the second stage of the experiment, participants were asked "Which Italian city has a higher population?". The participant could choose between A and B, both of which represented Italian cities the participant had provided recognition data on in the first stage of the experiment. As noted in the introduction, our key focus is on investigating pure recognition without associated knowledge of the response options and thus we effectively disabled internal memory–based search by concealing the names of each city. In order to aid the decision making process, participants were presented with a number of cues on screen for which they can retrieve information.[1] Crucially, one of the cues the participant could access was the recognition cue of which data was provided in the first stage of the experiment. A screenshot of the second stage of the experiment is provided in Figure 1.

The final manipulation in our experiment consisted of the availability of information on two key aspects of each cue. These are each cues' *validity* and *discriminability*. The validity of a cue quantifies the number of times a cue points you to the right answer as a ratio of the times it discriminates between the two response options. For instance, in the Italian cities Environment, the cue "Is the city the national capital?" has a validity of 1, because whenever one alternative scores positive on this cue, that will be because that alternative is the city Rome and Rome is the largest Italian city. The discriminability of a cue quantifies the number of times a cue discriminates between two response alternatives as a ratio of all possible cue comparisons for each question. The national capital cues does not discriminate very often and therefore has a low discriminability, because this cue will only discriminate when one of the two response alternatives is Rome. In the "+info" condition, cue validity and discriminability was shown on screen, in the "-info" condition, this information was not available to the participant. Note that for the continuous recognition cue, the cue discriminates if both scores are different from eachother. Thus, if one cue scores 0, it makes no difference whether the other scores 1 or 100. This manipulation was included to examine whether provision of information about the usefulness of recognition, in particular, affected its use. The validity and discriminabilty information can be seen as an aid to answering the meta–cognitive ques-

tion facing the participant — i.e., how useful is knowing that I recognize an object for drawing an inference?

All cues and their validities and discriminabilities for both environments are shown in Table 1. These cue validity and cue discriminability rates were calculated for the subset of 100 comparisons the participants had to make in the task, rather than for the whole set of possible comparisons. The reason for this was to ensure that participants in the "+info" condition could relate the presented cue validity and cue discriminability rates as close as possible to their actual experience when performing the task. The presented information could be used by participants to base their search order on cue validity, cue discriminability or a combination of the two. After each trial, participants received feedback with respect to the accuracy of their response. The experiment was self–paced.
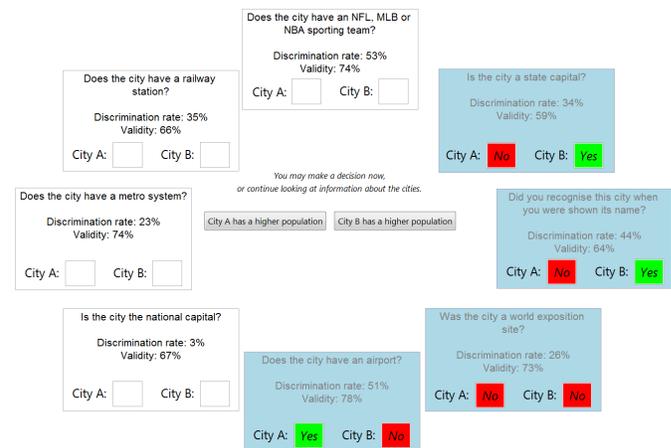

Figure 1: Screenshot of a trial of the binary version of the US cities task with cue information present ("+info"). See text for details.

Cues were presented in a circular array on the screen in random order. Participants examined cues by clicking on them. The order in which buttons were clicked was self–controlled. Deciding to stop examining additional cues was self–controlled, but conditional on having encountered at least one discriminating cue to dissuade guessing.

Procedure

Participants completed the Italian version and the US version of the task in random order. Participants were given instructions that they would have to indicate whether or not they recognized a number of cities, after which they performed stage 1 of the experiment, the recognition phase. Participants were subsequently instructed that they repeatedly had to make a choice between pairs of two alternatives. The concepts cue validity and cue discriminability were explained. Participants then performed the second stage of the experiment. After completing the experiment for the first environment, the second environment was administered.

---

[1]This is different from the original German cities task, in which cues had to be retrieved from memory and city names were revealed.

Design

Our experiment consists of eight conditions. The cities environment was a within–subject manipulation with two levels: Italian and US. The recognition mode was a between–subject manipulation with two levels: binary and continuous. Cue information was a between–subject manipulation with two levels: +info (info present on screen) and -info (no info on screen).

Results

For all statistical analyses, we report not only conventional p–values but also Bayes factors (e.g., Jeffreys, 1961). In contrast to p–values, Bayes factors allow researchers to quantify evidence in favor of the null hypothesis vis–a–vis the alternative hypothesis. For instance, when the Bayes factor $BF_{01} = 10$ the observed data are 10 times more likely to have occurred under H_0 than under H_1. When $BF_{01} = 1/5 = 0.20$ the observed data are 5 times more likely to have occurred under H_1 than under H_0. In the following, Bayes factors for analysis of variance are based on the BIC approximation (e.g., Wagenmakers, 2007; Masson, 2011), and Bayes factors for t–tests are based on the default Bayesian t–test proposed by Rouder, Speckman, Sun, Morey, and Iverson (2009).

We ran a 2x2x2 ANOVA with mode and cue information as between–subject independent variables and environment as a within–subject independent variable. Response accuracy was higher in the Italian environment (73.9%) than in the US environment (69.2%; $F(1,96) = 78.5$, $p < .05$, $BF_{01} = 1.1 \cdot 10^{-12}$). The following subsections report on the recognition proportion, the recognition validity and discriminabillity, and the recognition usage respectively.

Recognition Proportion

Figure 2 shows the proportion of cities that were recognized for each environment. As expected, recognition was higher for the US cities environment than for the Italian cities environment, as evidenced by a main effect for environment $(F(1,96) = 298.9$, $p < .05$, $BF_{01} = 2.0 \cdot 10^{-30}$; cf., Goldstein & Gigerenzer, 2002). Continuous recognition led to some parts of the scale being used besides the two extremes, suggesting that participants did not treat recognition as purely binary. However, the extremes were still the most popular.

Recognition Validity and Discriminability

Recall that the validity and discriminability of the recognition cue was calculated for each participant separately based on their answers in stage 1 of the experiment. Based on the recognition proportion for each environment, we expected to find that recognition was more valid, but less discriminating, for the Italian environment than for the US environment (cf., Goldstein & Gigerenzer, 2002). We were interested to see how cue mode would affect cue validity and discriminability.

Figure 3 shows recognition validity and discriminability for both environments. For recognition validity, there is a main effect for environment $(F(1,96) = 4.9$, $p < .05$, $BF_{01} = 0.82$; note that the Bayesian test indicates the evidence is
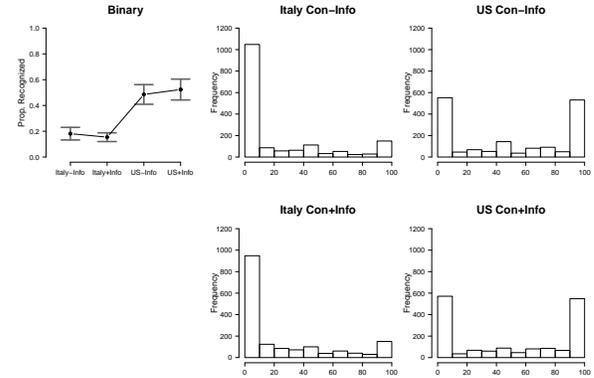

Figure 2: Proportion of cities recognized for the binary environment (top–left panel) and the continuous environment (other panels).

ambiguous); recognition validity may be higher in the Italian environment than in the US environment. There is also a tentative main effect for mode $(F(1,96) = 4.6$, $p < .05$, $BF_{01} = 0.99$; note that the Bayesian test indicates the evidence is ambiguous); validity for binary cues may be higher than for continuous cues.

For recognition discriminability, there is a main effect for environment $(F(1,96) = 57.8$, $p < .05$, $BF_{01} = 5.8 \cdot 10^{-10}$); recognition discriminability is lower in the Italian environment than in the US environment. There is also a main effect for mode $(F(1,96) = 70.4$, $p < .05$, $BF_{01} = 1.1 \cdot 10^{-11}$); continuous cues discriminate better than binary cues.
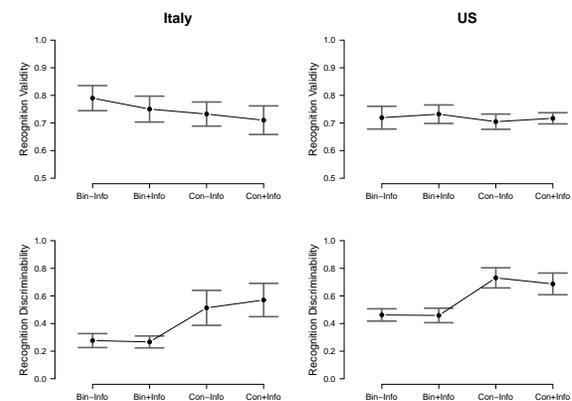

Figure 3: Recognition validity (top) and discriminability (bottom) for the Italian (left) and US (right) environments.

We have established there are differences in the validity and discriminability of the recognition cue that are a direct consequence of the cue being binary or continuous: recognition discriminates between response alternatives more often, but the extra information is, tentatively, less valid. Do decision makers use the recognition cue differently depending on the mode of the cue?

Recognition Usage

The top panels of Figure 4 show the proportion of trials the recognition cue was used for each environment. On average, participants did not use the recognition cue on all trials. On an individual basis, 9% of the participants used the recognition cue on all trials in the Italian environment and 8% of the participants used the recognition cue on all trials in the US environment.

For recognition use, there was a main effect of mode ($F(1,96) = 9.4$, $p < .05$, $BF_{01} = 0.09$); decision makers use continuous recognition more than binary recognition. There was also a main effect of environment ($F(1,96) = 22.8$, $p < .05$, $BF_{01} = 2.4 \cdot 10^{-4}$), decision makers use the recognition cue more in the US environment than in the Italian environment. Interestingly, there was a mode by environment interaction ($F(1,96) = 4.0$, $BF_{01} = 1.33$). In the Italian environment, recognition is used more often if it is continuous than if it is binary ($t(98) = -3.65$, $p < .05$, $BF_{01} = 0.02$). In the US environment, recognition usage does not depend on the mode of the cue ($t(98) = -1.52$, $p > .05$, $BF_{01} = 2.21$; note that the Bayesian test indicates the evidence is somewhat ambiguous). It is likely then, that the benefits of continuous recognition are highest when only a small portion of items are recognized. Finally, there is little evidence for a main effect of cue information ($F(1,96) = 3.3$, $p > .05$, $BF_{01} = 1.88$; but note that the Bayesian test indicates the evidence is somewhat ambiguous).

The bottom panels of Figure 4 show the average position in which the recognition cue was searched, given that it was examined for each environment. On average, participants did not search the recognition cue first. On an individual basis, for both environments, the lowest mean position of examination for the recognition cue was exactly 2, suggesting that not a single individual used the recognition heuristic in its most stringent form.

For recognition position, there was a main effect of mode ($F(1,95) = 4.5$, $p < .05$, $BF_{01} = 0.08$)[2]; recognition was used earlier when it was binary than when it was continuous. There was no main effect of environment ($F(1,94) = 2.1$, $p > .05$, $BF_{01} = 3.27$)

Conclusion

The goal of our experiment was to isolate 'pure' recognition and to examine participants' use of recognition information in the absence of concomitant evaluation. We argued that this would give us insight into participants' meta–cognition about the usefulness of recognition in different environments. What have we learned?

First we note that the accuracy of inferences about population size was higher for an environment about which participants, initially, knew less (Italian cities) than for one about

---

[2]Two participants never used the recognition cue and as such had no recognition position data. Recognition position was divided by the total number of cues for each environment to make both environments compatible.
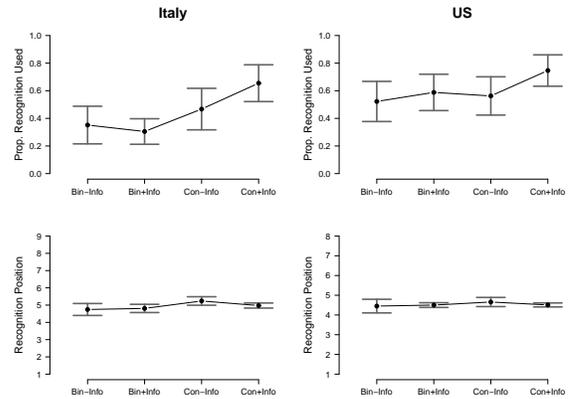

Figure 4: Proportion of trials the recognition cue was used (top) and position in search order (bottom) for the Italian (left) and US (right) environments.

which they knew more (US cities). But was the difference driven by adaptive use of recognition information? Under a strict interpretation, even in our novel task, recognition should be consulted on every trial — unless a participant has reason to believe that it will never discriminate (e.g., if they know they recognized either all or none of the cities in the environment). The results showed that was clearly not the case, with only a small proportion of the participants using the recognition cue on all trials. Moreover, even the participants that did examine recognition on every trial did not exclusively examine this cue first, challenging the idea that recognition information is somehow privileged in inference tasks (e.g., Pachur & Hertwig, 2006).

Additionally, we examined whether we could increase usage of recognition by measuring recognition on a continuous rather than a binary scale. We concluded that for the Italian environment where only a small proportion of the items were recognized, measuring recognition on a continuous scale led to recognition being used more often, despite the fact that recognition was less valid for the continuous scale than for the binary scale. No such effect was found for the US environment, in which on average about half of the cities were recognized. Though it is possible that the intermediate datapoints on the continuous recognition scale simply reflect perceived task demands by the participants, this alternative explanation does not seem to be in line with the fact that participants subsequently use continuous recognition more.

Our final question was whether participants would use recognition more often when their recognition cue was more valid. Surprisingly, we concluded the opposite: recognition was used more frequently in the US environment than in the Italian environment, despite the fact that recognition was more valid, on average, in the Italian environment. This finding suggests that meta–cognition about the usefulness of recognition is not particularly fine–tuned: adaptive use of recognition would predict greater reliance in environments where it is more useful (Gigerenzer & Goldstein, 2011).

Thus we conclude that 'pure' recognition can be compensated: knowing that we recognized one object and not another, but not knowing why is not enough for most participants to make a decision. Furthermore, it is not the first piece of information participants search for. In addition, our results show that recognition is more than a binary yes–or–no phenomenon. Allowing participants to indicate their recognition on a continuous scale led to an average increase in usage.

These results provide a novel and intriguing set of empirical regularities concerning the use of pure recognition information in a multi–attribute decision task. The next stage of this project will draw on the considerable advances that have been made in developing computational models of recognition–based judgments (e.g., Marewski & Mehlhorn, 2011; Marewski & Schooler, 2011) in an attempt to describe these data more fully. Our starting point will be to compare models that assume recognition information is used as evidence that can be accumulated much like any other cue to aid inference (e.g., van Ravenzwaaij, Moore, Lee, & Newell, 2013) with those that afford recognition an elevated status.

Acknowledgments

References

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

Gigerenzer, G., & Goldstein, D. G. (2011). The recognition heuristic: A decade of research. *Judgment and Decision Making*, *6*, 100–121.

Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *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.

Hilbig, B. E. (2010). Reconsidering "evidence" for fast–and–frugal heuristics. *Psychonomic Bulletin & Review*, *17*, 923–930.

Jeffreys, H. (1961). *Theory of probability*. Oxford, UK: Oxford University Press.

Lee, M. D., & Zhang, S. (2012). Evaluating the coherence of Take–the–best in structured environments. *Judgment and Decision Making*, *7*, 360–372.

Marewski, J. N., & Mehlhorn, K. (2011). Using the ACT–R architecture to specify 39 quantitative process models of decision making. *Judgment and Decision Making*, *6*, 439–519.

Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review*, *118*, 393–437.

Masson, M. E. J. (2011). A tutorial on a practical Bayesian alternative to null–hypothesis significance testing. *Behavior Research Methods*, *43*, 679–690.

Newell, B. R. (2011). Recognising the recognition heuristic for what it is (and what it's not). *Judgment and Decision Making*, *6*, 409–412.

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

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.

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

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t–tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.

van Ravenzwaaij, D., Moore, C., Lee, M., & Newell, B. R. (2013). How to collect information in uncertain decision environments? A hierarchical Bayesian modeling approach. Manuscript submitted for publication.

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p*-values. *Psychonomic Bulletin & Review*, *14*, 779–804.