

Racing for the City: The Recognition Heuristic and Compensatory Alternatives

Katja Mehlhorn (s.k.mehlhorn@rug.nl)

Dept. of Artificial Intelligence and Dept. of Psychology, University of Groningen, the Netherlands

Julian N. Marewski (marewski@mpib-berlin.mpg.de)

Max Planck Institute for Human Development, Berlin, Germany, and MIT Sloan School of Management, Cambridge, MA

Abstract

In the last decade a debate in the decision making literature has centered on the question whether decisions can be better described by simple non-compensatory heuristics or by more complex compensatory strategies. We argue that this debate should be led at a higher level of precision. Theories about decision strategies are implemented at different levels of description and they often only make verbal, qualitative predictions. This makes it difficult to compare between them and to test them against quantitative process data. A way to make theories comparable and improve the precision of their predictions is to model them within one computational framework. Using the example of the recognition heuristic, we show how simplifying dichotomies such as the one between non-compensatory and compensatory decision strategies can dissolve when using detailed quantitative models.

Keywords: recognition heuristic, compensatory strategies, cognitive model, ACT-R

Introduction

Imagine you are asked which of two cities is larger, York or Stockport. You do not know the correct answer but remember that you recently read an article about York in the newspaper. You recall that it had some mentionable industry, but no international airport and also no premier league soccer team. Of the city of Stockport you have never heard before. Which city will you answer to be the larger one? To respond, you could employ the *recognition heuristic* (Goldstein & Gigerenzer, 2002). According to it, if you recognize one of the alternatives, but not the other, you may infer the recognized one to be larger. Your answer would be York. As an alternative to the recognition heuristic, you may rely on a strategy that uses your knowledge about the city's attributes as cues. Following corresponding *compensatory* models of decision-making, (e.g., unit-weight linear strategy), you might conclude that the absence of an airport and a premier league soccer team speak against York being a large city. Consequently, you might infer Stockport to be larger.

The example illustrates a debate that has received much attention in the decision-making literature (for an overview see Marewski, Pohl, & Vitouch, 2010). Are decisions better to be described by simple non-compensatory heuristics, or by complex compensatory decision strategies? The recognition heuristic is a non-compensatory model for memory-based decisions: Even if further knowledge beyond recognizing an alternative is retrieved, this knowledge is ignored when the heuristic is used. In contrast to this assumption, many other decision models posit that people

evaluate alternatives by using knowledge about their attributes as cues. The common idea behind such compensatory models is that an alternative's value on one cue can be traded off against its value on another cue.

A large amount of evidence has been gathered for as well as against both positions - for support of the recognition heuristic see for example: Gigerenzer, Hoffrage, and Goldstein, 2008; Pachur, 2010; and Volz, et al., 2006; for challenges of the heuristic see for example: Beaman, Smith, Frosch, and McCloy, 2010; Dougherty, Franco-Watkins, and Thomas, 2008; and Oppenheimer, 2003. However, non-compensatory and compensatory decision strategies themselves are broad categories that subsume a number of different models. For instance, compensatory strategies propose that knowledge about the alternatives is used in some way; however they do not agree on how this is done. Constraint satisfaction models, for example, assume that all available information is integrated at once, in a parallel, automatic fashion (Glöckner & Betsch, 2008). Evidence accumulation models, in contrast, assume that evidence for the alternatives is accumulated sequentially until a decision boundary is reached (e.g., Lee & Cummins, 2004).

In testing different decision strategies against each other, research has encountered various problems. First, theories are often specified at varying levels of detail, making it difficult to directly compare them. Second, many theories have been formulated at a verbal qualitative level and are therefore underspecified relative to the empirical data against which they are tested. Consider the city size example again. Based on different theories, one might generate predictions about decision times; the time participants need to decide which of the two cities is larger. Figure 1 illustrates the paradigm that is typically used to assess these times. Presented with the names of two alternatives (like the cities York and Stockport), the participant is asked to infer which of the two has a larger value on a criterion (e.g., which of the two cities is larger). All information the person wants to use for this decision has to be retrieved from memory. Using this paradigm for testing different strategies, one could for example assume that decision time increases with the amount of knowledge that is used in the decision making process (Bröder & Gaissmaier, 2007). However, participants' decision times will not only depend on the decision strategy itself, but also on other factors, like the time it takes to read the names of the cities, to retrieve information from memory, and to enter a response. Consequently, the contribution of the decision strategies themselves might be drowned out by these

additional factors (cf. Hertwig, Herzog, Schooler, & Reimer, 2008).

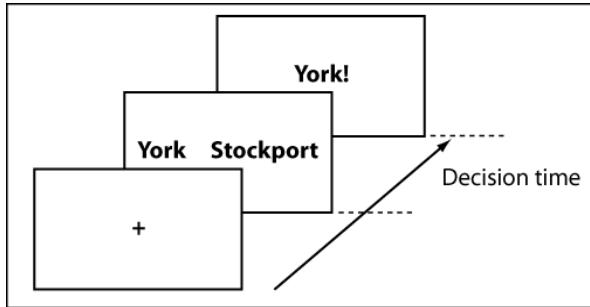


Figure 1. Sample trial of the memory paradigm as it is usually used to assess participants' decisions and decision times.

In the current project we try to tackle both of the issues mentioned above. First, we implement different decision strategies into one cognitive modeling framework. This results in directly comparable quantitative predictions of the strategies. Second, by using a cognitive architecture for this implementation, we take into account the interaction with additional components of cognition, like reading, memory retrieval and giving a motor response. This allows us for assessing the contribution of different decision strategies at a higher level of precision and for directly comparing them against empirical data.

Methods

Empirical Data

The empirical data we used to test different strategies was gained by reanalyzing the data of Experiment 1 from Pachur, Bröder, and Marewski (2008), which has been argued to provide evidence for both the recognition heuristic and compensatory strategies (Gigerenzer, Hertwig, & Pachur, 2010). In the decision phase of the experiment, Pachur et al. presented their participants with choices between cities, as in the introductory example: a recognized city with three associated cues and an unrecognized city about which nothing was known. The cues were industry, airport, and soccer. Each cue could be either positive (speaking for a city being large) or negative (speaking against a city being large). The cities varied in the pattern of associated cues, with two, one, or zero of the three cues being negative (Table 1). For each pair of cities the decision and the decision time were assessed as shown in Figure 1. However, in their paper, Pachur et al. only reported the decision data. Before the decision phase of the experiment, participants learned the names and cue patterns of the six to-be-recognized cities. After the decision phase, recognition memory for all cities and cue memory for the to-be-recognized cities were assessed. Only pairs in which participants responded to have recognized one, but not the other city were analyzed and used in the model tests.

Table 1. Positive (+) and negative (-) cues associated to the six to-be-recognized cities in Experiment 1 of Pachur et al. (2008)

Cue	City					
	Aber-deen	Bristol	Notting-ham	Sheffield	Bright-on	York
Industry	+	+	+	+	+	+
Airport	+	+	-	-	-	-
Soccer	+	+	+	+	-	-

Models

To test between different strategies we generated quantitative models that were implemented in the cognitive architecture ACT-R (Anderson, et al., 2004). From all possible modeling accounts we chose ACT-R because it takes into account both sub-symbolic and symbolic components of cognition as well as perceptual and motor processes. This allows for modeling the task as it was solved by the participants and to directly compare the modeling results to the empirical data. Like the participants in the experiments, the models read the city names from a screen, make a decision about which of the cities is larger and indicate their decision by pressing one of two keys. Models' decisions and decision times are assessed. Below we describe the details of the decision procedure for each model.

Assessing recognition. All models start with assessing recognition of the cities (see Pachur & Hertwig, 2006, for evidence suggesting that when asked to make a decision between alternatives, people will first assess their recognition). In modeling recognition, we follow Anderson et al. (1998) and Schooler and Hertwig (2005) in assuming that a city is recognized, if it can be retrieved from memory. In ACT-R, the probability and the time required for retrieving a city from memory depend on the city's level of activation. The activation A_i of a city i in memory is determined by three components as shown in Equation 1:

$$A_i = B_i + S_i + \varepsilon. \quad (1)$$

The first component is the city's base-level activation B_i which reflects the frequency of encounters with the city in the past. The second component is the spreading activation the city receives from the current context S_i , reflecting its usefulness in the current context. The third term ε is a random noise component that is calculated from a logistic distribution. Given a city's activation A_i , the probability that it will be retrieved is calculated by,

$$p = \frac{1}{1 + e^{-s}}, \quad (2)$$

where τ describes the threshold that has to be crossed for a retrieval and s describes a noise component. The time required for a successful retrieval decreases with increasing activation A_i of a city i , as shown in Equation 3,

$$\text{retrieval time} = Fe^{-A_i}, \quad (3)$$

where F describes the latency of the retrieval. If the model cannot retrieve a city, the time it takes to notice a retrieval

failure is calculated using the retrieval latency F and the retrieval threshold τ :

$$\text{retrieval failure time} = Fe^{-\tau}. \quad (4)$$

Assessing cue knowledge. After the initial assessment of recognition, a subset of the models additionally retrieves knowledge about cues from memory. We assume that this retrieval is performed using the same kind of retrieval processes as are described above for the retrieval of cities (see Equation 1-4). Analyzing the cue memory task that followed the decision phase in Pachur et al. (2008), we found that participants remembered negative cues slower than positive ones. To reflect this fact, we let the models retrieve positive cues faster than negative ones by giving them different activation levels.

Decision strategies. Whereas all models assess recognition as a first step, they differ in the strategies that lead to a decision.

Non-compensatory strategies. We implemented four models to test variations of the non-compensatory recognition heuristic. These models always decide for the recognized city. However, they differ in the amount of knowledge retrieved from memory before this decision is made. As memory retrieval takes time, depending on the amount of knowledge that is retrieved before the decision, the models produce different decision time predictions.

Model-1. Implementing the simplest version of the recognition heuristic (Goldstein & Gigerenzer, 2002), this model directly uses the outcome of the recognition assessment and responds with the recognized city.

Model-2. Implementing a newer proposal for the recognition heuristic (cf. Pachur, et al., 2008), this model retrieves knowledge about the three cues of the recognized city from memory. After the cues are retrieved, the model responds with the recognized city, without using the retrieved cue knowledge in the decision.

Model-1&2. This model presents a combination of Model-1 and Model-2, in assuming a race between their strategies. After recognition is assessed, the strategies to directly decide for the recognized city and to retrieve a cue race against each other¹. This race is repeated until the decision is made.

Model-1&2-F. This model is identical to Model-1&2, but it additionally assumes that retrieved cues will at times be forgotten. The intuition behind this assumption is that processing a cue can detract from previously retrieved cues

(cf. Mensink & Raaijmakers, 1988 for such interference accounts of forgetting). Forgetting is implemented by an additional race between retrieve-a-cue, respond-with-recognized and forgetting that starts as soon as at least two cues have been retrieved from memory.

Compensatory strategies. The remaining models were implemented to test different versions of compensatory decision strategies. Depending on the cue knowledge associated to a city, these models can decide for and against the recognized city. They differ in how the cue knowledge is used in this decision and they produce different decision time predictions.

Model-3. This model implements a strategy that assumes that cue knowledge is used implicitly, by memory activation processes (Glöckner & Betsch, 2008). After assessing recognition, it retrieves knowledge about the three cues of the recognized city from memory. After all cues are retrieved, the model tries to form an impression about the recognized city's size. It does this by attempting to retrieve information that indicates whether the city is large. The probability that this information can be retrieved depends on memory activation spreading from positive cues (Equation 2). The more positive cues are associated to a city, the more activation is spread and the higher the chance that the city is assessed as large. If the model cannot assess the city as large, it will enter the unrecognized city.

Model-1&3. In assuming a race between the strategies of Model-1 and 3, this model implements a combination of the non-compensatory recognition heuristic and a compensatory decision strategy. After recognition is assessed, the strategies to directly decide that the recognized city is larger and to retrieve a cue race against each other. This race is repeated until the decision is made or all cues are retrieved. If all cues are retrieved and no decision has been made yet, the model can additionally try to form an impression about whether the city is large by using memory activation as implemented in Model-3.

Model-1&3-F. This model is identical to Model-1&3, but it additionally assumes that retrieved cues will at times be forgotten as in Model-1&2-F.

Model-4. This model uses cue knowledge explicitly by means of a decision criterion as suggested by evidence accumulation models (Lee & Cummins, 2004). After assessing recognition, it retrieves knowledge about the cues for the recognized city. As soon as enough positive or negative cues are retrieved to meet the model's decision criterion, it responds with the recognized city (in case of positive cues) or the unrecognized city (in case of negative cues). To reflect different possible decision criteria, the model is implemented in different versions. **Model-4.1** responds as soon as one positive or negative cue is retrieved. **Model-4.2** needs two positive or negative cues for a decision, and **Model-4.3** needs all 3 cues to be positive or negative to reach its criterion. If the model cannot retrieve enough cues to reach its criterion, it uses recognition as its best guess.

¹ In the literature, the terms "race" or "race model" are sometimes used in similar ways as the terms "evidence accumulation" or "sequential sampling models". For instance, Gold and Shadlen (2007) define race models as process where "evidence supporting the various alternatives is accumulated independently to fixed thresholds" (p. 541) and as soon as one of the alternatives reaches the threshold, it is chosen. Applying the race to ACT-R's production rules, we implemented a simplified version of that mechanism, where competing production rules have equal *utilities* (Anderson, et al., 2004) and are therefore chosen at random. Put in Golden and Shadlen's terms, the production rules have equal chances of reaching the threshold.

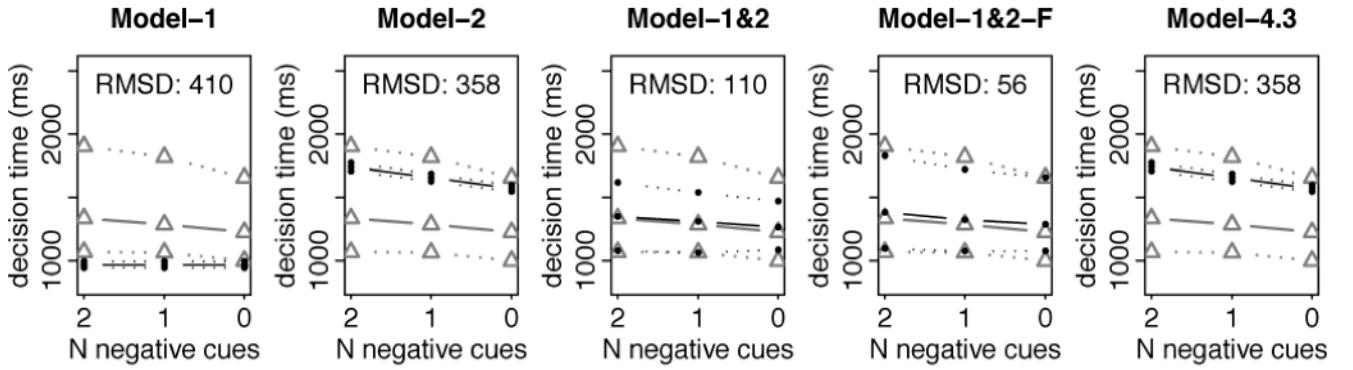


Figure 2. Decisions times (median and quartiles) for participants (grey) and models (black) that always chose the recognized city. RMSDs were calculated separately for median and quartiles and then averaged.

Results

Pachur et al. (2008) found that some participants answered always in accordance with the recognition heuristic ($N=25$), whereas others seemed to sometimes use their cue knowledge to decide against the recognized city ($N=15$). To investigate this difference further, we analyzed the data of these two groups of participants separately (subsequently referring to them as *recognition* and *cue group*). To investigate the effect of cue knowledge, we analyzed decisions (% of choices for recognized city) and decision time distributions (medians and quartiles) separately for two, one, and zero negative cues.

In fitting the models to the human data, we used a stepwise procedure to constrain the parameter space. Specifically, we first fit the parameters associated with recognition and cue retrieval on data of the recognition and cue-memory tasks. To do so, we implemented separate ACT-R models of recognition and cue retrieval, which were fit to the recognition and cue memory data. With these parameters fixed, we then estimated the remaining parameters from participants' decision times in the decision task². To allow maximum comparability between the models, all parameters were kept constant between the models. Each model was run 40 times on the trials of each participant.

Recognition Group. As it was to be expected, the non-compensatory models (Model-1, Model-2, Model-1&2, and Model-1&2-F) always decided for the recognized city, replicating the decisions of the recognition group. Also the compensatory Model-4.3 showed this decision behavior, because it could never reach its decision criterion of three negative cues that would have been necessary to decide against the recognized city. In producing 100% decisions for the recognized city, by definition of the groups, all these models reached a RMSD of 0 to the decision data of the

recognition group. The models largely varied in their decision time patterns (see Figure 2). In the empirical data, decision times had a large spread and increased in a linear fashion with the amount of negative cues associated to a city. As was to be expected, Model-1 produced fast decision times that did not vary as a function of cue knowledge and did not vary between trials. Retrieving cue knowledge before the decision, Model-2, Model-1&2, and Model-1&2-F produce a linear increase of decision times with the number of negative cues. Only the models that implement a race between different strategies are able to reach a spread comparable to the empirical data. The best fit is reached by Model-1&2-F which assumes that cues are at times forgotten.

Cue Group. Most compensatory models (Model-3, Model-1&3, Model-1&3-F, Model-4.1, and Model-4.2) decide for the recognized city in part of the cases. The exact proportion of these choices depends on the amount of negative cues associated to the cities and differs between the models (Figure 3). In the empirical data, the proportion of choices for the recognized city was overall high but decreased with the number of negative cues associated to a city. Using cue-knowledge implicitly, Model-3, Model-1&3, and Model-1&3-F reflect the decreasing proportion of choices for the recognized city. In doing so, Models-1&3 and 1&3-F fit the decisions of participants well, whereas Model-3 underestimates the overall proportion of choices for the recognized city. The decision patterns of Model-4.1 and Model-4.2 deviate substantially from the human data. Using a strict decision criterion, these models show a sudden drop in choices for the recognized city when reaching their decision criterion of one or two negative cues. As in the recognition group, the empirical decision times show a small linear increase with the number of negative cues and have a large spread. Whereas all models that fit the cue groups decisions produce linearly increasing decision times, only Model-1&3-F is able to fit the large spread.

² ACT-R's latency factor (F) was set to .1, the retrieval threshold (τ) to -.3, and activation noise (s) to .2. The base levels (B_i) were set to 4.1 for cities and positive cues and to .18 for negative cues. Spreading activation between cities and negative cues was set to 0. Maximum associative strength (S) was set to 3. Visual attention latency was set to .035 and imaginal delay to .1. All other parameters were kept at the default values of ACT-R 6.0 (Anderson, 2007).

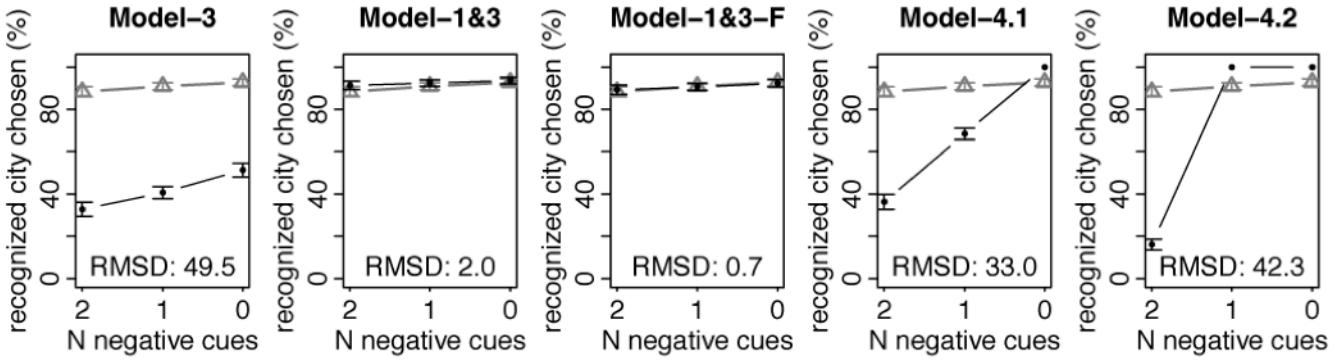


Figure 3. Proportion of choices for the recognized city for participants (grey) and models (black) that chose the unrecognized city in part of the trials.

Conclusions

A number of strategies have been proposed for how people make memory-based decisions between alternatives. In the current article we explore how such strategies can be evaluated against each other by using the precision of a cognitive architecture. By implementing a number of decision models that have originally been defined at different levels of description into *one* architectural modeling framework, we make these models directly comparable to each other. By not only modeling decision processes, but also the interplay of these processes with perceptual, memory, and motor processes, we produce quantitative predictions that can be directly compared to the empirical data.

For participants that always responded with the recognized city (recognition group), as well as for those that sometimes decided against the recognized city (cue group), models that implemented a race between different strategies performed best. These models were not only able to fit the participants' decisions, but also the distribution of their decision times. The success of these models is interesting, because, even though they were not identical, they were very similar. Both types of models retrieve and encode cues in some of the trials, whereas in others they merely assess recognition. The models only differ in one respect, namely, that the best-fitting model of the recognition group (Model-1&2-F) exclusively relies on recognition, whereas the best

fitting model of the cue-group (Model-1&3-F) sometimes additionally acts on implicit knowledge, from which it gains intuition about the recognized city's size. While this difference between the two winning models is thus a subtle one, it does seem to reflect an important psychological variable that is consistent with the literature on intuitive versus deliberate modes of processing (e.g., Evans, 2008).

Our results show how dichotomies like the one between non-compensatory and compensatory processes can dissolve when implementing decision strategies at a higher level of detail and precision in a cognitive architecture. In addition, the good fit of the race models to the human decision times highlight the possibility that even people who always responded with recognized cities most likely retrieved and encoded cues in at least some of the trials. People who sometimes responded with unrecognized cities, in turn, most likely based their decisions on cues in some of the trials but ignored these cues and relied entirely on recognition in others.

Before concluding, we want to discuss some possible limitations of our approach. First, the empirical data against which we tested our models was gathered in an artificial setting where participants were explicitly taught the cue values of the to-be-recognized cities before the experiment. This allowed for high control of participants' knowledge and simplified the models. For example, in modeling this experiment it was reasonable to assume that the different

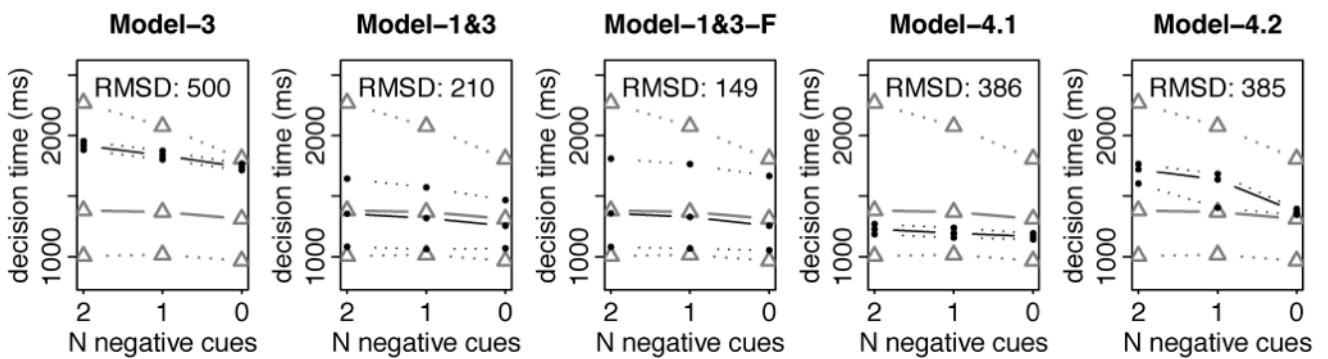


Figure 4. Decisions times (median and quartiles) for participants (grey) and models (black) that chose the unrecognized city in part of the trials. RMSDs were calculated separately for median and quartiles and then averaged.

cues (soccer, industry, and airport) are represented in memory with equal strength and that therefore, retrieval times and probability did not vary between the cues. However, in real-life, where knowledge is acquired naturally, the situation becomes more complex. In naturalistic settings, the activations of different pieces of information will vary as a function of the environment (Anderson & Schooler, 1991), resulting in different probabilities and speed of retrieval for different pieces of information. Future research will have to show if those models that won the model comparison here, will also be able to generalize best to such more naturalistic settings.

Second, in implementing decision strategies that differ in their level of description and that are often underspecified in aspects important for the implementation, we had to make a number of additional assumptions. All assumptions are grounded in the decision, memory, and ACT-R literatures. Often, however, these literatures offer more than one plausible solution. Future evaluation of the different strategies and their implementation will be necessary to test the extent to which our results are due to core features of the modeled strategies and to which extend they were caused by additional assumptions we had to make for implementing the strategies.

Summarizing, our results suggest that models, which implement a race between competing decision strategies, best predict people's decisions and decision time distributions. This demonstrates how simplifying dichotomies that are so often used in psychological research can dissolve when using quantitative models that specify the interplay of underlying cognitive processes.

References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036-1060.

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

Anderson, J. R., & Schooler, L. (1991). Reflections of the environment in memory. *Psychological Science*, 2(6).

Beaman, C., Smith, P., Frosch, C., & McCloy, R. (2010). Less-is-more effects without the recognition heuristic. *Judgment and Decision Making*, 5(4), 258-271.

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

Dougherty, M., Franco-Watkins, A., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, 115(1), 199.

Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.

Gigerenzer, G., Hertwig, R., & Pachur, T. (2010). *Heuristics: The Foundations of Adaptive Behavior*. Oxford University Press.

Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas (2008). *Psychological Review*, 115(1), 230.

Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberate decision making. *Judgement and Decision Making*, 3(3), 215-228.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535-574.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75-90.

Hertwig, R., Herzog, S., Schooler, L., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Learning, Memory*, 34(5), 1191-1206.

Lee, M., & Cummins, T. (2004). Evidence accumulation in decision making: Unifying the 'take the best' and the 'rational' models. *Psychonomic Bulletin & Review*, 11(2), 343.

Marewski, J. N., Pohl, R. F., & Vitouch, O. (2010). Recognition-based judgments and decisions: Introduction to the special issue. *Judgment and Decision Making*, 4(4), 207-215.

Mensink, G., & Raaijmakers, J. (1988). A model for interference and forgetting. *Psychological Review*, 95(4), 434-455.

Oppenheimer, D. (2003). Not so fast!(and not so frugal!): Rethinking the recognition heuristic. *Cognition*, 90(1), B1-B9.

Pachur, T. (2010). Recognition-based inference: When is less more in the real world? *Psychonomic Bulletin & Review*, 17(4), 589.

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(2), 183.

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(5), 983-1002.

Schooler, L., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, 112(3), 610-628.

Volz, K. G., Schooler, L. J., Schubotz, R. I., Raab, M., Gigerenzer, G., & Von Cramon, D. Y. (2006). Why you think Milan is larger than Modena: Neural correlates of the recognition heuristic. *Journal of Cognitive Neuroscience*, 18(11), 1924-1936.