Constraining ACT-R Models of Decision Strategies: An Experimental Paradigm

Cvetomir M. Dimov (cvetomir.dimov@unil.ch) and Julian N. Marewski (julian.marewski@unil.ch)
Department of Organizational Behavior, Université de Lausanne,
Quartier UNIL-Dorigny, Bâtiment Internef, 1015 Lausanne, Switzerland

Lael J. Schooler (schooler@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Lentzeallee 94
14195 Berlin, Germany

Abstract

It has been repeatedly debated which strategies people rely on in inference. These debates have been difficult to resolve, partially because hypotheses about the decision processes assumed by these strategies have typically been formulated qualitatively, making it hard to test precise quantitative predictions about response times and other behavioral data. One way to increase the precision of strategies is to implement them in cognitive architectures such as ACT-R. Often, however, a given strategy can be implemented in several ways, with each implementation yielding different behavioral predictions. We present and report a study with an experimental paradigm that can help to identify the correct implementations of classic compensatory and non-compensatory strategies such as the take-the-best and tallying heuristics, and the weighted-linear model.

Keywords: Take-the-best, tallying, weighted-linear model, process models, ACT-R

Introduction

One important characteristic of well-developed scientific theories is precision. In psychology, theoretical precision can be achieved by complementing verbally formulated theories with formal models. Typically, formal models are specified in terms of mathematical equations or computer code. The goals, level of detail, and level of description of such models vary as a function of the psychological subdiscipline, research questions being asked, or the available technology, to name only a few factors. Computational models have become both increasingly popular and powerful, and have aided cognitive scientists in their endeavor to shed light into the behaviorist's black box. Computer models allow one to specify, on an algorithmic level, the cognitive processes psychological mechanisms are assumed to draw on.

Such process models predict not only what decision a person will make, but also how the information used to make the decision will be processed. The predictions made by these models can thus be tested not only on *outcome data* (e.g., what item is chosen) but also on *process data*, including on patterns of information search, response times, or neural activation. Such predictions can eventually differentiate among competing theories that make identical outcome predictions. In particular in the cognitive and decision sciences, describing cognitive processes represents a central goal of theorizing on its own. In fact, the past decades have seen repeated calls to develop process models.

Yet, surprisingly there are relatively few theories of decision making that yield detailed quantitative predictions about process data. Instead, typically qualitative predictions about response times and other process data are tested in experiments. This theoretical and methodological weakness contributes to fuelling important scholarly debates about which decisional processes describe behavior best: simple *non-compensatory* ones, for which decisions based on some predictors cannot be overturned by others, or complex *compensatory* integration processes, for which various predictors can neutralize each-other's influence (cf. Bröder & Schiffer, 2003; Glöckner & Betsch, 2008; Marewski et al., 2010).

One way to increase the precision of theories of decision making is to implement them in detailed cognitive architectures such as the *ACT-R theory of cognition* (e.g., Anderson, 2007). ACT-R is a quantitative framework that applies to a broad array of behaviors and tasks, formally integrating theories of memory, perception, action, and other aspects of cognition. ACT-R also allows modeling decision processes. When models of decision making are implemented in ACT-R, quantitative predictions about response time distributions at the millisecond level and other process data can be made and compared to experimental studies. Marewski and Mehlhorn (2011), for instance, implemented several compensatory and non-compensatory decision strategies in ACT-R. In doing so, they modeled for each of the strategies how decisional processes interplay with memory, perceptual, and motor processes, which, in turn, allowed them to quantitatively predict the response time distributions associated with using each strategy in a simple two-alternative forced choice decision task.

While the architectural approach can thus help remedying the aforementioned theoretical and methodological weakness, this approach does not come without its complications. Specifically, often a given strategy can be implemented in numerous different ways in ACT-R (or other cognitive architectures), with each implementation yielding different response time and other process predictions. Part of the problem is that many decision strategies are—in the worst case—only formulated verbally or—in the best case—specified mathematically or algorithmically, without spelling out the strategies' assumptions about lower-level cognitive processes. This *specification problem* (see Lewandowsky, 1993), namely how to translate an underspecified theory or strategy into a detailed cognitive model, poses a paramount modeling challenge to the researcher who sets out to find out which implementation is the most adequate one. To illustrate this point, Marewski and Mehlhorn (2011) actually ended up implementing over thirty ACT-R models of similar decision strategies without being able to make strong conclusions about which model most likely represented the correct one.

In this paper, we present and report a study with an experimental paradigm that can help to build and identify the correct implementations of decision strategies. In what we call the *train-to-constrain-paradigm*, participants are instructed in a detailed step-by-step procedure how to apply specific strategies in a decision task. Since the experimenter thus knows which strategies participants have relied on in the experiment, the resulting response times lend themselves to constraining ACT-R implementations of these strategies. Specifically, as an initial step, here we focus on a variant of that paradigm in which participants are instructed to apply three classic compensatory and non-compensatory strategies, namely the *take-the-best* (henceforth: *TTB*) and *tallying* heuristics, and the *weighted-linear model* (henceforth: *WLM*).

The remainder of this paper is structured as follows. First, we will explain in more detail the three decision strategies. Second, we will present the train-to-constrain-paradigm and, in doing so, report a study that we ran using that paradigm. Third, we will report the results of this study, and, fourth, briefly illustrate how these results can be used to build and constrain ACT-R implementations of the three strategies.

Decision Strategies

Tallying and WLM have been formulated in different ways (and at times also been given different names); here we use Gigerenzer and Goldstein's (1996) definitions as well as their TTB heuristic. Gigerenzer and Goldstein specified these strategies as models of inductive *inference* about unknown quantities or future events in simple two-alternative forced choice tasks. In such tasks, a person has to infer which of two alternatives (e.g., cities) has a larger value on a given criterion (e.g., population). One variant of this task that has received considerable attention during the past years is the *memory-based decision task* illustrated in Figure 1. In this task, a person has to make inferences by relying exclusively on the contents of their memory. The experimental paradigm for identifying correct ACT-R implementations of TTB, tallying, and the WLM that we propose here extends this memory-based task.

Take-the-best. The simple TTB heuristic stands in the tradition of Tversky's (1972) classic *elimination by aspects* model. TTB bases inferences on the attributes of the alternatives (e.g., whether a city has an airport), which it uses as *cues*. A cue can have a *positive* (e.g., a city has an airport, coded as "1"), *negative* (has no airport, coded as "-1"), or an *unknown* (coded as "0") value. The vector of cue values that define a person's knowledge about a specific alternative is called the alternative's *cue profile*. TTB bases inferences on just one good cue. Specifically, TTB orders the cues *i* unconditionally according to their validity v_i, with $v_i = c_i/(c_i + w_i)$, c_i being the number of correct inferences based on cue *i* given that it *discriminates* between two alternatives (i.e., cue values are 1 & 0, respectively, or 1 & -1, respectively), and w_i the number of incorrect inferences. TTB's rules for searching cues, stopping search, and making a decision can be summarized as follows:

Search: Search through cues in the order of their validity.
Stopping: Stop as soon as a cue is found that discriminates between the alternatives.
Decision: Infer that the alternative with the positive cue value has the higher value on the criterion of interest.

As can be seen, TTB is a non-compensatory strategy, which uses solely the first discriminating cue. Translated into a process prediction this implies, for example, that the time it takes to make decisions with TTB should depend on how many cues have been considered before a discriminating cue is found.

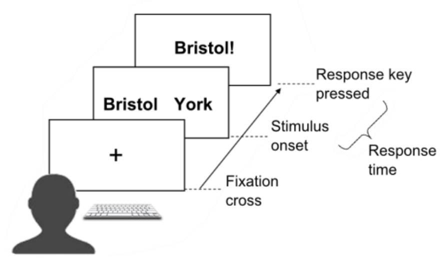

Figure 1: Illustration of the memory-based decision task

Tallying. In contrast to TTB and other non-compensatory strategies, many decision models posit that people evaluate alternatives by integrating knowledge about multiple cues. One such heuristic is tallying. This representative of classic unit-weight linear integration models (e.g., Dawes, 1979) simplifies decisions by treating all cues equally. For each alternative, tallying simply counts the cues with positive values and infers that alternatives with the larger number of positive cue values score higher on the criterion of interest. As a consequence, the various cues can neutralize each other's influence on the final decision, thus making tallying a compensatory model. Tallying's search, stopping, and decision rules read as follows:

Search: Search through cues in any order.
Stopping: Stop search after *m* out of a total of *M* cues (with $1 < m < M$) have been accessed.
Decision: Decide for the alternative that is favored by more positive cue values. If the number of positive cue values is the same for both alternatives, guess.

Weighted-linear model. The WLM is similar to tallying in that it integrates all the information available when choosing an alternative. In the WLM, cue values are coded like in TTB. As suggested by its name however, it integrates all cue information by multiplying the cue values by their validities and summing them over for each city, thus computing the weighted sum of the cues for each city. The WLM's rules can be summarized as follows:

Search: Search through cues in any order.

Stopping: Stop search after m out of a total of M cues (with $1 < m < M$) have been accessed. Multiply each cue value with its validity and compute the weighted sum of cues for each alternative.

Decision: Decide for the alternative that is favored by the larger weighted sum. If the weighted sum is the same for both alternatives, guess.

The WLM has a long tradition in the cognitive and decision sciences and beyond. For instance, variants of this model have been viewed as optimal rules for preferential choice and are often considered to define rational behavior (cf. Payne, Bettman, & Johnson, 1993).

Experimental Paradigm

The train-to-constrain-paradigm builds on several earlier studies on TTB, tallying, and the WLM (e.g., Bröder & Gaissmaier, 2007; Bröder & Schiffer, 2003; Mata, Schooler & Rieskamp, 2007) and on approaches that teach subjects to rely on specific decision strategies (e.g., Khader et al., 2011; Marewski & Schooler, 2011).

In our study, we implemented the training portion of our paradigm in a computerized experiment, in which subjects were told that they would participate in a quiz show. In that show, they first learned fictitious facts about how British cities would look like in the future, namely whether these cities would have an international airport, a train station, a university, and/or a premier league soccer team in the year 2100 (such facts are typically judged as useful for inferring city size; cf. Pachur, Bröder, & Marewski, 2008). In a second step, subjects learned how to employ a strategy that uses these facts as cues to make decisions. During the actual quiz show, they then saw pairs of cities on the computer screen and were instructed to always use the strategy to infer which of the two cities would be larger in the year 2100. Subjects were paid according to the degree to which their decisions agreed with predictions of the respective decision strategy.

Subjects and design. A total of 141 subjects participated in the experiment (89 male, $M_{age} = 25.3$), of which 120 finished it successfully. Subjects were randomly assigned to one of three between-subjects conditions. The conditions differed in terms of the strategy participants learned to use. In the first condition subjects learned TTB, in the other two conditions they learned tallying and the WLM, respectively.

Materials. Sixteen well-known British cities were used as alternatives. These cities correspond to those that most subjects in Pachur et al.'s (2008) pre-study 1 recognized. A pre-study suggested that subjects' familiarity with these cities' names aids them to learn a large number of facts about these cities. Since the degree of familiarity was roughly the same for all cities in both Pachur et al.'s pre-studies, no interference effects of familiarity were expected, and, indeed, also none found. These 16 cities were combined with 8 cue profiles, illustrated in Table 1. In doing so, each of the 8 cue profiles was used twice—albeit with different city names.

Table 1: Cue profiles used

	City1	City2	City3	City4	City5	City6	City7	City8
Airport	+	+	-	-	+	+	-	-
Soccer team	-	-	-	-	-	-	+	+
University	-	-	+	+	+	+	+	+
Train station	+	-	+	-	+	-	+	-

Learning task. The experiment started with a *learning task* (cf. Bröder & Schiffer, 2003), in which subjects were taught the 4 cues about the 16 British cities, corresponding to a total of $4 \times 16 = 64$ facts. Specifically, during learning, cities and cues were presented repeatedly in a random order until subjects correctly recalled at least 14 of the 16 cities' cue profiles perfectly. Cue profiles were assigned at random to specific cities.

Strategy learning task. After having learned all cues, in each of the three between-subjects conditions, subjects were trained how to use one of three decision strategies. The strategy learning procedure required subjects to go through a stepwise explanation of the decision process assumed by each strategy as well as to apply that strategy correctly on several practice trials that mimic the actual decision task. During practice, subjects received feedback about whether they had applied the strategy correctly, and the strategy was practiced until subjects' decisions concurred to 100% with the strategy's predictions. During the strategy learning task, subjects also memorized additional information that is necessary for applying the strategy, such as the cue validities in the case of TTB and WLM. The instructions on how to use each strategy were crafted such that they reflect the strategy descriptions from the literature.

Repetition of learning task. To make sure participants still remembered the 64 facts correctly, one round of the learning task was repeated upon completion of the strategy learning task.

Decision task. In a *decision task*, 72 pairs of the previously learned British cities were presented (one city on the left side of the computer screen, the other one on the right; see Figure 1). To avoid inducing frequency effects, the pairs were constructed such that each city name appears equally often. Subjects were instructed to always apply the strategy to decide which of the cities will be larger in the year 2100. For each correct application of the strategy, subjects received a bonus payment of 0.5 Euros (0.68 US$). Each decision inconsistent with the strategy's prediction resulted in a penalty of 0.5 Euros (no feedback was given).

Cue-memory task. In a *cue-memory* task, subjects had to reproduce the cue values they learned for the cities. The purpose of this task was to collect data about how well subjects remembered the cue values they were taught. This data will be used in future projects to populate the declarative memory of the ACT-R models.

Experimental Results

Figure 2 shows the mean of the 25th, 50th and 75th response time percentiles for the three experimental conditions as a function of the number of cues that have to be retrieved from memory prior to finding the most valid discriminating one (henceforth: *most valid discriminating cue*). Several important observations can be made. First, tallying participants made the fastest decisions. Their response time varied from under 3s for the 25th percentile to almost 6s for the 75th percentile. This is much faster than previous decision making experiments have reported. For example, Bröder and Gaissmaier (2007) reported mean response times between 6.5s and 8s in their first, and between 11s and 15s in their second experiment. It should be noted that those experiments did not instruct subjects to rely on specific strategies, but that instead used participants' decisions to infer, post hoc, by means of strategy classification procedures which strategies subjects have used.

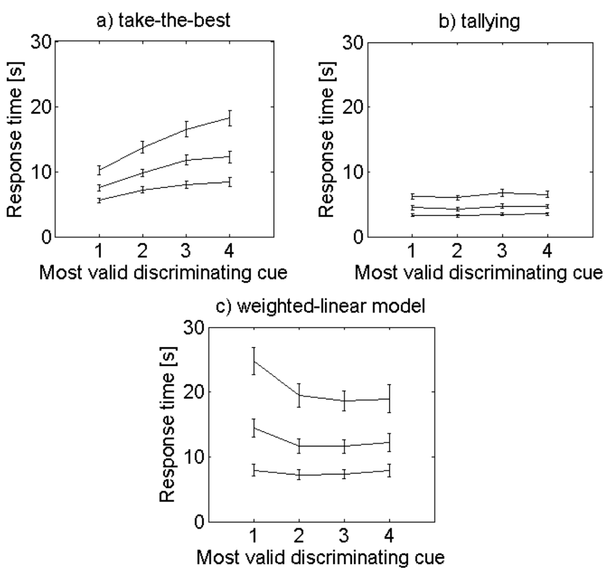

Figure 2: Participants' aggregate response time percentiles as a function of most valid discriminating cue. Error bars are standard errors of the mean computed across all participants in the respective experimental condition.

Second, the response times of TTB participants fall in the response time range of those reported in these previous experiments. However, this resulted in participants in the TTB condition being slower than tallying participants, which also is a finding that stands in contrast to previous studies, in which post hoc strategy classification procedures were used (e.g., Bröder & Gaismaier, 2007).

Third, WLM participants are the slowest, which is a result that is consistent with Bröder and Gaissmaier's (2007) earlier studies. Bröder and Gaissmaier reported mean response times between 10s and 11s in their first and between 15s and 23s in their second experiment, which fall close to the time range of our participants.

Fourth, as can be seen in Figure 2a, TTB participants' response times increase as a function of most valid discriminating cue. In contrast, Figures 2b and 2c show that for tallying and the WLM the response times do not exhibit such an increase when they are analyzed in the same way as for TTB participants. This result is to a large extent consistent with earlier work: in Bröder and Gaissmaier's (2007) experiments, participants who were inferred to have relied on TTB exhibited strong increases in mean response times as a function of the most valid discriminating cue, while those who were classified as likely users of tallying or the WLM did not exhibit increases that were as strong.

Implementing Strategies in ACT-R

In the *constraining* portion of our paradigm, the observed response times will be used to build and constrain ACT-R implementations of the three decision strategies. Specifically, each individual participant's responses in the memory task can be used to model the contents of that subject's declarative memory after having gone through the training phase. These declarative memory contents can then be used to model the retrieval processes associated with using each of the three decision strategies (cf. Marewski & Mehlhorn 2011, for this approach). Together with perceptual, motor, and other cognitive processes—all of which can be modeled in ACT-R—these retrieval processes will contribute to the response times predicted by the corresponding ACT-R models of the decision task.

Overview of ACT-R

ACT-R describes cognition as a set of modules that interact through a production system. The production system consists of *production rules* (i.e., if-then rules) whose conditions (i.e., the "if" parts) are matched against the contents of the modules. If a rule's conditions are met, then the production rule can fire and the specified action is carried out. Each module implements different cognitive processes. The *declarative module*, for instance, enables information storage in and retrieval from memory, the *intentional module* keeps track of a person's goals, while the *imaginal module* holds information necessary to perform the current task. A *visual module* for visual perception and a *manual module* for motor actions (e.g., typing on a keyboard) simulate interactions with the world. In coordinating the modules, the production rules can only act on information that is available in *buffers*, which can be thought of as processing bottlenecks, linking the modules' contents to the production rules. For instance, the production rules cannot access all contents of the declarative module, but only these that are currently available in the retrieval buffer. ACT-R distinguishes between a symbolic

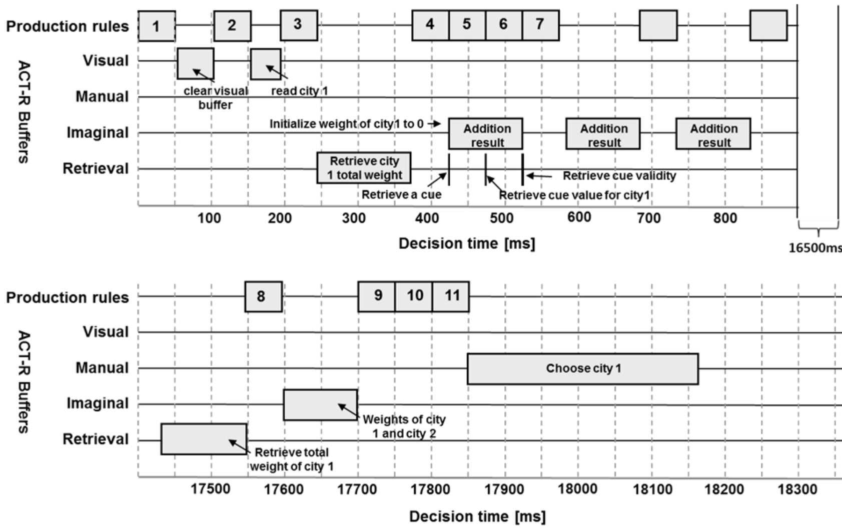

1. **If** there is visual information on the left side of the screen that has not been attended before
   Then move attention to this information
2. **If** the information on the screen (i.e. city1) is attended
   Then encode it in the visual buffer
3. **If** city1 is encoded in the visual buffer
   Then try to retrieve its total weighted sum
4. **If** city1's total weighted sum cannot be retrieved
   Then initialize its total weighted sum to 0 in the imaginal buffer and retrieve a cue
5. **If** a cue has been retrieved
   Then try to retrieve its value for city1
6. **If** the cue value has been retrieved and it is positive
   Then retrieve its validity
7. **If** the validity of the cue has been retrieved
   Then encode its validity in the imaginal buffer, and initiate adding its validity as an increment to the current total weighted sum of city1
8. **If** city1's total weighted sum has been retrieved and city2's total weighted sum has been stored in the imaginal buffer
   Then encoded city1's total weighted sum into the imaginal buffer
9. **If** both cities' total weighted sums are encoded into the imaginal buffer
   Then compare the hundreds of the two total weighted sums
10. **If** the hundreds of the two total weighted sums are the same
    Then compare the tens of the two total weighted sums
11. **If** one of the tens of the two total weighted sums is bigger
    Then choose the city with the higher tens

Figure 3: Processing stream of the weighted-linear model for the first and last seconds of the decision process. Production rules on the right hand side are stylized representations of the actual ACT-R productions for this model. Note that the model's decision time predictions can vary across different decision trials, for instance, as a function of perceptual and motor processes, or cue activation. Also note that the same production rules fire more than once during the process.

and a subsymbolic system. The symbolic system is composed of the productions rules as well as of the modules and buffers. Access to the information stored in the modules and buffers is determined by the subsymbolic system. This system is cast as a set of equations and determines, for instance, the timing of memory retrieval.

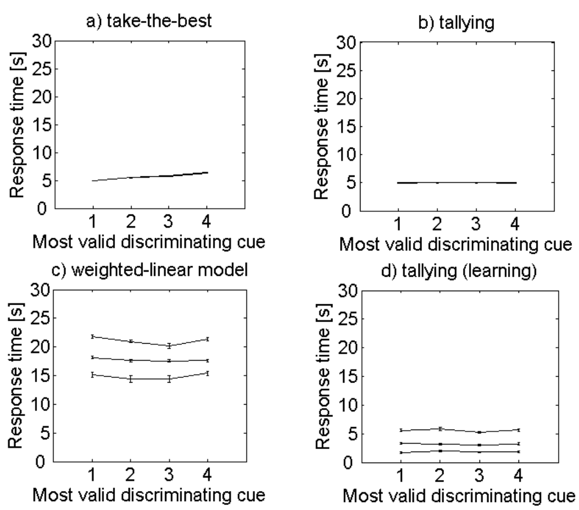

Figure 4: ACT-R predictions of response time percentiles of a tallying and weighted-linear model implementation. Error bars are standard errors of the mean, computed across 30 simulation runs of the ACT-R model.

Illustrating our ACT-R models

Figures 4a, 4b and 4c present our preliminary ACT-R models, developed prior to running the experiment as a source of rough predictions of participants' eventual behavior. All of these three models are, perhaps, the most naïve implementations which follow the above mentioned strategy definitions and experimental instructions. In developing these models, no parameters were fitted, but those from Marewski and Mehlhorn (2011) were used.

All models perform the same task as our experimental subjects: The models "read" two city names off a computer screen, process them, decide for one of them, and enter a response by "pressing" a key. To illustrate this, Figure 3 shows the first and last seconds of an 18-seconds-long processing stream of our preliminary ACT-R implementation of the WLM. The various decisional, memorial, perceptual and motor processes assumed by the model are coordinated by production rules.

Specifically, by first "reading" the names of both cities, the model tries to retrieve a memory trace of the city names called a *chunk*. Chunks are facts like "York is a city" or "York has an airport" which model people's familiarity with city names and their cue knowledge about these cities, respectively. For each cue, the model retrieves its validity. If the cue value is positive, the model adds the validity of this cue to the weighted sum of the city, initiating a summation procedure. If the cue value is negative, the model subtracts the validity of the corresponding cue from the weighted sum of that city, initiating a subtraction procedure. Finally, the model compares the total weighted sums of the two cities and chooses the one with the larger total weighted sum by pressing a key. As Figure 4c shows, the predicted response time percentiles of 30 simulation runs of this WLM ACT-R implementation lie close to the 75th percentile range observed in participants' data (Figure 2c), suggesting that this implementation is not an implausible model, but also that other processes which boost participants' response times, such as memorizing the weighted sum, are present in participants. Our preliminary tallying model (Figure 2b) predicts response times within experimental data, while the TTB model (Figure 2a) is faster. These three models have to be adapted to successfully capture participants' behavior, a more successful example of which is the tallying model presented on Figure 4d, which was built after the experiment. While the former tallying model did not include memorization of the number of positive cue values of already seen cities, the latter model did, which produced a response time distribution close to participants' response times. Exact modeling of each participant's cue knowledge is the next modeling step to be made. Naturally, after identifying the most promising implementations of all strategies, all models would then have to be tested in new experiments, this way ensuring that they can also account for behavior in tasks for which they were not developed.

Discussion and Conclusion

While it goes beyond the scope of this short proceedings paper to present more ACT-R implementations—that is part of a larger research paper—one legitimate question one may raise is what the methodological advantages of our approach over earlier experimental work is. As mentioned above, in earlier studies including Marewski and Mehlhorn's (2011) ACT-R modeling efforts and Bröder and Gaissmaier's (2007) response time analyses for TTB and other heuristics, participants' decisions had to be used to infer, post hoc, by means of strategy-classification and/or other model selection procedures which strategies participants relied upon in an experiment. As a result, the conclusions that could be drawn from analyses of response times crucially hinged on the accuracy of the strategy classification and/or model selection procedure. Our train-to-constrain approach, in contrast, allows identifying the response time patterns associated with a strategy without the need to use potentially inaccurate strategy classification. To illustrate this point, the deviations observed between Bröder and Gaissmaier's and our findings could, besides being a product of differences in the stimuli and materials used, also be a result from the strategy classification method used by these authors. More studies with our paradigm, including experiments that make use of Bröder and Gaissmaier's stimuli and materials, are warranted to decide between these and other competing explanations.

To conclude, response times such as the ones observed in our experimental paradigm can be used to find out which ACT-R implementation best mirrors classic decision strategies used by trained subjects. Once identified, these implementations can, hopefully, be used to model behavior both in previously published studies as well as in new studies in which subjects' decision strategies are unconstrained by training.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 107–129.

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

Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, *132*, 277–293.

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

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

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

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

Khader, P.H., Pachur, T., Meier, S., Bien, S., Jost, K., & Rösler, F. (2011). Memory-based decision-making with heuristics: Evidence for a controlled activation of memory representations. *Journal of Cognitive Neuroscience*, *23*, 3540–3554.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D., & Gigerenzer, G. (2010). From recognition to decisions: Extending and testing recognition-based models for multialternative inference. *Psychonomic Bulletin & Review, 3*, 287-309.

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.

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.

Newell, B. R., & Bröder, A. (2008). Cognitive processes, models and metaphors in decision research. *Judgment and Decision Making*, *3*, 195–204.

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.

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

Schulte-Mecklenbeck, M., Kühberger, A. & Ranyard, R. (Eds.). (2011). *A handbook of process tracing methods for decision research: A critical review and user's guide*. New York: Taylor & Francis.

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