

The Robustness of the Take The Best Configural Heuristic in Linearly and Nonlinearly Separable Environments

Ulrich Hoffrage (Ulrich.hoffrage@unil.ch)

Ecole des Hautes Etudes Commerciales (HEC), Université de Lausanne, BFSH1, CH-1015 Lausanne, Switzerland

Rocio Garcia-Retamero (rretamer@mpib-berlin.mpg.de)

Center for Adaptive Behavior and Cognition, Max Plank Institute for Human Development, Lentzeallee 94, D-14195 Berlin

Uwe Czienskowski (sciencec@mpib-berlin.mpg.de)

Center for Adaptive Behavior and Cognition, Max Plank Institute for Human Development, Lentzeallee 94, D-14195 Berlin

Abstract

Take The Best (TTB) is a fast and frugal heuristic for paired comparison that has been proposed as a model of bounded rationality. This heuristic does not take compound cues into account to predict the criterion. However, causal knowledge about the relationship between a criterion and several cues may suggest that processing cues as configurations could be adaptive in a certain environment. In a series of simulations, we show that an extension of TTB, namely TTB-Configural, outperforms both the more frugal TTB and a more demanding benchmark. Moreover, we review empirical evidence that people process cues as configurations when equipped with the corresponding knowledge about the causal structure.

One-Reason Decision Making: The Approach of Fast and Frugal Heuristics

When we are faced with a decision it is often impossible to consider all the possible alternatives, their attributes, and their potential consequences. In fact, in everyday life we often make fast decisions based on little information. Recently, the ABC Research Group at the Max Planck Institute for Human Development has suggested that we use *fast and frugal heuristics* in these situations, that is, simple but nevertheless fairly accurate rules in the mind's adaptive toolbox¹ for making decisions with a minimum of information (Gigerenzer, Todd, & The ABC Research Group, 1999).

The fast and frugal heuristic that has received the most attention to date is *Take The Best* (TTB; Gigerenzer & Goldstein, 1996). This heuristic is designed to infer which of two alternatives, described on several dichotomous cues, has a higher value on a quantitative criterion. TTB is a process model that consists of *building blocks*, which are precise steps of information gathering and processing involved in generating a decision. More specifically, this heuristic has a *search rule*, which describes the order in which to search for information (TTB looks up cues in the order of their *validity*, i.e., the probability that a cue will make the correct decision given that it discriminates between the alternatives); a

stopping rule, describing when information search is to be stopped (TTB stops after the first discriminating cue); and a *decision rule*, which describes how to use the available information to make a decision (TTB chooses the alternative favoured by the first discriminating cue).

When TTB was used to make predictions in real-world environments—for instance, to predict which of two cities has a higher homelessness rate, which of two professors earns more money, or which of two persons has a higher percentage of body fat—it turns out that it could compete well with more savvy strategies such as multiple regression, particularly in cross-validation (Czerlinski, Gigerenzer, & Goldstein, 1999). Cross-validation comparisons were made by training on half of the items in each data set and testing on the other half of the data. By being simple and focusing only on some, but relevant, information, TTB avoided being too closely matched to any particular environment and outperformed regression in accuracy when generalizing to the test set.

Compound Cue Processing within the Fast and Frugal Heuristics Approach

The fast and frugal research program has, however, been criticized (e.g., by Garcia-Retamero, Dieckmann, Hoffrage, & Ramos, 2005a) for considering only individual cues when setting up the cue search order, as is the case, for instance, in TTB. Consequently, fast and frugal heuristics do not benefit from the fact that the correct decisions in certain environments might depend on the interactions between individual cues. For instance, according to the diathesis–stress model (Walker & Diforio, 1997), hereditary predisposition and a current stressor (e.g., a traumatic experience) are assumed to be necessary for the onset of schizophrenia. Therefore, in a medical diagnostic context, a doctor might predict that a patient with a high predisposition to this disease who suffered a recent traumatic experience would be in worse shape than other patients who present either only one or neither of these factors. Note that in this example, the combination of two cues produces a different value on the criterion than either one cue or neither cue.

The dilemma of the selection of potential combinations of cues (i.e., compound cues) for decision making becomes even more critical if we consider the evidence in several research

¹ The adaptive toolbox is the term coined to describe the collection of specialized cognitive mechanisms in the human mind for inference and reasoning.

areas that suggests that people indeed can and do process compound cues as configurations² in certain environments (e.g., Shanks, Charles, Darby, & Azmi, 1998).

One issue in the literature on configural strategy use that has generated special interest is whether linearly separable environments are easier to learn than nonlinearly separable ones (see Smith, Murray, & Minda, 1997), which indeed seems to be the case. In a nonlinearly separable environment, such as the eXclusive-OR (or XOR, for short) logical structure, the optimal response depends on the relationship between the components of a compound cue (Shepard, Hovland, & Jenkins, 1961). In an environment where two cues are amalgamated into a compound that obeys the XOR logical rule, an object for which one of the cues that form the critical compound is present and the other is not has a higher criterion value than an object for which both are present or absent. A configural representation is needed to solve this problem because there is no set of weights for the individual cues that could make the system generate a proper response. On the other hand, a problem is linearly separable (e.g., the AND logical structure) when the optimal response can be generated as a linear function of the weighted cues (Minsky & Papert, 1969). In an environment where two cues are amalgamated into a compound that obeys the AND logical rule, an object for which both cues that form the critical compound are present has a higher criterion value than an object for which only one or neither cue is present (for an example, see the above-mentioned diathesis–stress model for the onset of schizophrenia).

Note that individual cues in the environment can be combined in many possible pairs, triples or even higher orders to form compound cues in many different ways, thus leading to a combinatorial explosion. Therefore, it is extremely difficult to select from this array of compounds those that are highly valid and meaningful (i.e., stable). A strategy that processes all possible combinations of cues as configurations by default would thus be too computationally demanding (Kehoe & Graham, 1988), that is, it would be neither fast nor frugal. One possible solution is to assume that decision makers do not spontaneously process all possible compound cues in the environment as configurations and keep track of their validities but instead focus on only the subset of those that seem “plausible”. The interesting question is then how this process might work.

The Adaptive Value of Causal Knowledge for Compound Cue Processing

When it is said that a cause brings about an effect, it implies that there is a stable causal link between the cause and the effect, and an underlying causal mechanism that is an essential property of the link between the events involved in this relationship (Ahn & Kalish, 2000). We hypothesize that

² In our terminology, a configuration is constructed from several individual cues to create a new stimulus that differs from its components. We differentiate between the internal representation of a combination of cues (i.e., a configuration) and the observed compound cue in the environment.

people might use their knowledge about the causal structure of the environment to interpret the empirical input when they make decisions (see also Waldmann & Holyoak, 1992; Waldmann & Hagmayer, in press).

Along these lines, we submit that causal knowledge might allow decision makers to infer whether an elemental or a configural strategy would be adaptive in a certain environment. Particularly, we hypothesize that decision makers would assume that certain cues in the environment might interact with each other when they are perceived to act through the same underlying causal mechanism in bringing about the decision criterion (e.g., drinking alcohol and taking a prescription drug in affecting the stomach, possibly causing nausea). In this case, the effect of one cue would frequently be expected to be modulated by other cues, and they might be represented together as a configuration (see also Waldmann, 1996). A more complex decision-making strategy that processes such configural cues would then be used in that environment, especially when the cues that are perceived to act through a common causal mechanism indeed form a highly valid compound cue in the physical world.

We further hypothesize that cues might be represented as independent elements when they are perceived to act through different causal mechanisms to bring about the decision criterion (e.g., tailwind and physical strength in affecting the speed of a cyclist). An elemental strategy would then be more adaptive and less demanding in terms of time and cognitive resources than a configural strategy. In this way, causal knowledge about the cues’ causal mechanisms might allow decision makers to deal adaptively with the countless number of combinations of cues that appear in a particular environment, by directing them to those that might be relevant.

TTB-Configural: A Fast and Frugal Heuristic that Processes Compound Cues as Configurations

Building both on the literature on simple heuristics on the one hand and causal processing on the other, Garcia-Retamero et al. (2005a) and Garcia-Retamero, Hoffrage, Dieckmann, and Ramos (2005b), recently proposed a configural version of TTB, namely the *TTB Configural* heuristic. This heuristic, like TTB, is precisely defined in terms of its building blocks of search, stopping search, and decision rules and also assumes that cues are ordered in a hierarchy according to their validities. The contrast to TTB, which only considers individual cues, is that TTB-Configural processes compound cues, which are represented as configurations and included as such in the cue hierarchy.

To be fast and frugal, TTB-Configural processes a combination of several cues as a configuration when they are perceived to act through a common causal mechanism to bring about the decision criterion. However, TTB-Configural processes these cues as several individual elements when they are perceived to act through different causal mechanisms. Consequently, TTB-Configural includes both configural and individual cues in the cue hierarchy that is ordered according to the cues’ validities and begins searching with the cue that

has the highest validity—be it an individual or a compound cue. If the critical piece of information is a compound cue, *TTB-Configural* looks up the individual cues that jointly constitute that compound (for a similar treatment of compound cues in the context of multilevel set-covering models, see Baumeister & Seipel, 2002).

In the following, we first report the results of a simulation study in which we compared the performance of three heuristics that differed with respect to their treatment of compound cues. Subsequently, we review two series of experiments that have demonstrated that *TTB-Configural* also performs well when used as a behavioural model.

Simulation Study

In the simulation study, we cross-validated the performance of three decision-making strategies in several environments (generation of environments and evaluation of strategies were programmed in Visual C# .NET (Microsoft .NET Framework 1.1, Version 1.1.4322 SP1)). These strategies differ in whether they benefit from causal knowledge about the cues in the environment to process compound cues as configurations:

- *TTB* processes individual cues as independent elements and does not consider compound cues.
- In contrast, *TTB-Configural* processes compound cues as configurations when the decision maker’s causal knowledge about the compounds’ components induces the heuristic to do so. Therefore, this strategy only processes some of the possible compound cues in the environment.
- Finally, *TTB-All* processes all possible combinations of individual cues in the environment as configurations.

The task was to infer which of two alternatives, described on several dichotomous cues, had a higher value on a quantitative criterion. The performance of the strategies was determined across a complete-paired comparison of all objects of a given environment. The environments we used were artificially created, with the following procedure:

An environment consisted of 50 (or 100) objects that were described on 6 (or 8) elemental cues. These cues formed 3 (or 4) relevant compound cues, that is, cues for which we assumed that the decision maker would have causal knowledge suggesting these cues should be treated as compound cues. We consider this assumption to be psychologically plausible and reasonable since we are not concerned here about how causal knowledge has been acquired (cf., Lagnado & Sloman, 2002; Novick & Cheng, 2004), but only about how it can be used to aid the selection of some compound cues from among the wide range of possibilities in the environment through a top-down process. For simplicity, we only considered compounds of two cues (although we see no reason why our conclusions should not generalize to compounds of higher order).

We created environments in which the compound cues had high (or low) validities. In addition, as a control condition, we also used environments in which no such relevant compound cues existed. The three factors – number of objects (50 vs. 100), number of relevant compound cues (3 vs. 4), and

validity of the relevant compound cues (high vs. low vs. chance) were fully crossed, yielding 12 types of environments.

When creating an environment, we first generated the values of the relevant compound cues. In environments with 3 relevant compound cues, the values were randomly generated with the constraint that the resulting validities equaled .95, .9, and .85 in the high validity condition and .9, .8, and .7 in the low validity condition. For environments with 4 relevant compound cues, their validities were .95, .9, .85, and .8 in the high validity condition and .9, .8, .7, and .6 in the low validity condition. In the chance condition, all validities equaled .5.

Next, we generated elemental cues by splitting up these compound cues. This was achieved by taking the values of a given compound cue as input and by determining—according to the probabilities in Figure 1—the values of the two resulting elemental cues as outputs. Thereby, we used either the logical AND rule or the logical XOR rule. Seen the other way around, that is, from output to input, it was either the AND or the XOR rule that had to be used for the amalgamation of the two elemental cues to result in the compound cue. The logical rule thus is the fourth factor that was varied when creating an environment, yielding altogether 24 types of environments.

AND Logical Rule			XOR Logical Rule			
Compound	Elemental Cues		Compound	Elemental Cues		
	Cue 1	Cue 2		Cue 1	Cue 2	
1	$\xrightarrow{p=1.0}$	1	1	$\xrightarrow{p=.5}$	0	1
	$\xrightarrow{p=.33}$	1	0	$\xrightarrow{p=.5}$	1	0
0	$\xrightarrow{p=.33}$	0	1	$\xrightarrow{p=.5}$	1	1
	$\xrightarrow{p=.33}$	0	0	$\xrightarrow{p=.5}$	0	0

Figure 1: Scheme according to which a compound cue with a particular value was split into two elemental cues. Note that the probabilities are conditioned on the value of the compound cue.

In a third step, we combined all the elemental cues to generate all possible compounds of two components, using both the AND and the XOR logical rule. Note that in this step the values of the 6 (or 8) elemental cues were the inputs and the values of the compound cues were the outputs. Because 6 (8) cues can be combined into 15 (28) pairs, the application of the AND rule and the XOR rule yielded 30 (56) compounds (which, of course, include the relevant 3 (4) compounds). Together with the 6 (8) elemental cues, an environment thus had 36 (64) cues, elemental or compound.

The three steps explained above—generation of relevant compound cues, generation of elemental cues from those compounds, and generation of all possible compounds from those elemental cues—were repeated 100 times for each of the 24 types of environments. Finally, we split each environment into two sets: the learning set and the test set,

with 50% of the objects in each set. We repeated the assignment of the objects to the learning and test set 100 times (yielding an output file with 24 * 100 * 100 lines). How do the three strategies process these cues? TTB only includes the 6 (8) elemental cues in its hierarchy. TTB-Configural includes the 6 (8) elemental cues and—directed by causal knowledge—in addition the 3 (4) relevant compound cues. TTB-All includes all 36 (64) cues in its hierarchy. Evident the spirit of simple heuristics is no longer present with TTB-All, as this strategy does not try to avoid a computation explosion.

What was our rationale for generating the environments according to the three steps explained above? The 3 relevant compound cues were created so that they had some validity in the whole set, and it can thus be assumed that this was also true for both the training and the test set. This, however, should only occasionally be the case for the other compound cues that were generated from the elemental cues in the third step. Here we expected that some of these would actually have a high validity in the training set, simply because of random fluctuations and because the set of possible compounds is quite large. But these compound cues that have a high validity in the training set will, due to statistical regression towards the mean, have much lower validities in the test set (as those with low validities in the training set can be expected to have higher validities in the test set).

As a consequence, TTB-Configural should be more robust than TTB-All. Specifically, in the training set TTB-All should outperform TTB-Configural, because the set out of which valid cues—compound or elemental—could be selected was larger for TTB-All. However, as most of the compounds that were highly valid in the training set would be useless in the test set, TTB-All should lose the advantage of selecting from more cues when evaluated in the test set.

Results

Data were first aggregated across the 100 runs per environment, and then these environment-specific means were aggregated across the 100 environments for a given type. The standard errors for the mean performance across the 100 environments for a given type ranged between .1 and .4 percentage points, and thus all factors that we manipulated turned out to be significant in a four-way analysis of variance. Still, the differences between the levels of the factors number of objects (50 vs. 100) and the factor number of elemental cues (6 vs. 8) were relatively small and uninteresting from the present perspective; therefore, all analyses reported below were performed across these two factors.

Figure 2 shows the performance of the three strategies in the AND environments. First, there was almost no difference between the environments where the compounds had a high validity and those with a low validity, but there was obviously a big difference between these environments and the chance condition, which did not contain any relevant compound cues. Second, in the learning set, TTB-All performed best, but in the test set, TTB-Configural outperformed both TTB and

TTB-All. Accordingly, TTB-Configural was most robust (the differences between performance in the training set minus that in the test set for the AND environments were 6.6, 6.1, and 8.9 percentage points for TTB, TTB-Configural, and TTB-All, respectively, averaged across all other factors; $F(2,1198)=2351.9, p<.001$).

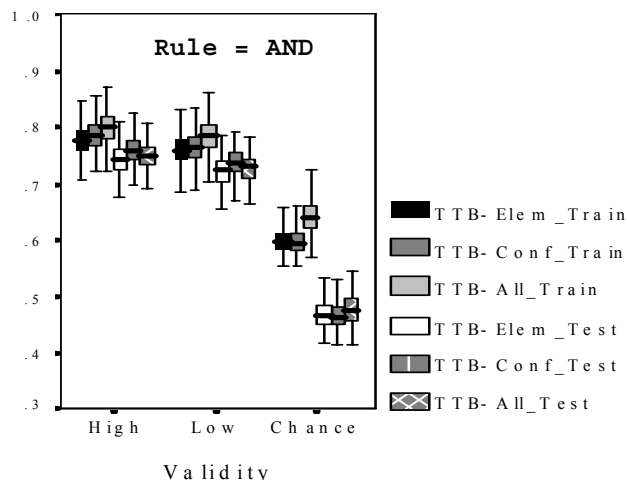


Figure 2: Proportion of correct inferences of TTB (in the legend denoted as TTB-Elem), TTB-Configural (in the legend denoted as TTB-Conf), and TTB-All in the three validity conditions (high vs. low vs. chance) of the AND-environments. The three leftmost box-plots within a validity condition depict performance of the three strategies in the training set, the three rightmost that in the test set.

Figure 3 shows the performance of the three strategies in the XOR environments. Again, there was almost no difference between the validity conditions (except for the contrast between the high and low validity conditions on the one hand and chance on the other). However, unlike in the AND environments, the simple TTB, which did not process compound cues, was markedly outperformed by TTB-Configural and TTB-All. This is hardly surprising, as the XOR environment is nonlinearly separable and thus one-reason decision making is not adaptive here (unless this one reason is a compound cue). Note that in the test set, TTB did not even perform better than chance. Third, TTB-All outperformed TTB-Configural in the training set, but the pattern reversed in the test set. This was the case because, as we had predicted, TTB-Configural was the most robust strategy (differences between performance in the training set and test set were 13.6, 2.4, and 4.0 percentage points for TTB, TTB-Configural, and TTB-All, respectively, $F(2,1198)=2234.4, p<.001$). Thus, TTB-Configural was more robust than TTB-All, even though TTB-All contained the relevant compounds as well. However, TTB-All had other compound cues even higher in its cue hierarchy that were not robust and had a substantially lower validity in the test set.

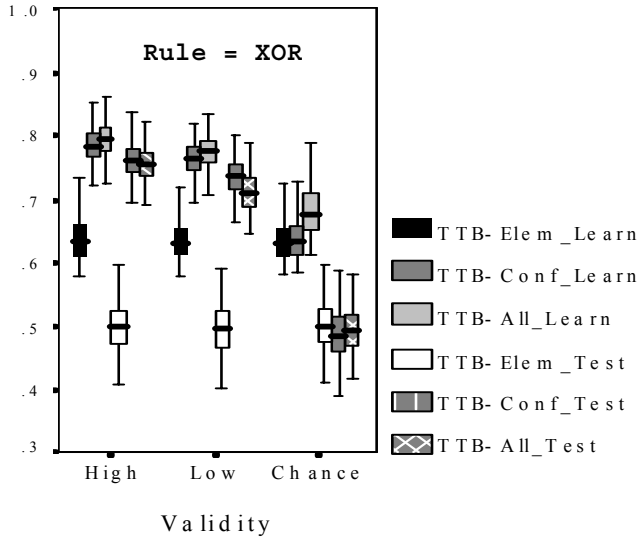


Figure 3: Performance of the three strategies in the XOR environments (for more explanations, see Fig. 2)

In sum, in linearly separable environments, where the relevant compound cues are constituted by amalgamating two elemental cues with the logical AND rule, (1) TTB-All was best when fitting known data, however, (2) TTB-Configural and TTB did not fall too far behind. Moreover, (3) when generalizing to new data, TTB-Configural was (slightly) better than TTB and TTB-All. In nonlinearly separable environments, the same pattern could be observed with respect to the comparison between TTB-Configural and TTB-All. However, TTB failed considerably in these environments, both when fitted to known data and in particular when used to predict new data.

Series of Experiments

In two previous series of experiments, Garcia-Retamero et al. (2005a, 2005b) explored whether decision makers will use TTB-Configural if a highly valid compound cue exists in the environment and if their causal knowledge about how the cues affect the criterion suggests that those elemental cues may interact with each other. In these experiments, both linearly separable and nonlinearly separable environments were used.

The task was framed as a medical diagnostic task. Specifically, participants were presented with information about two patients and had to choose the one who would show a higher body temperature. The information provided about the patients was whether they had ingested three different substances (A, B, and C, respectively). In a first block of trials, participants were provided with information about all three cues in a given comparison at no cost. Thereafter, they had to make a decision. Outcome feedback about the correct option was given to enable participants to learn the cues' validities. Subsequently, participants went through a decision-making phase in which the cue information was no longer available; instead they had to look it up sequentially. Again, outcome feedback was provided.

We generated complex environments where two of the three cues (i.e., A and B) were amalgamated into a highly valid compound cue. For the sake of brevity and simplicity, we will report results only for those environments where the compound was constituted by applying the logical XOR and the AND rule. The validity of the critical compound, AB, in those environments was 1.00, but the validity of each of its component cues, A and B, was .50. The validity of the third cue, C, which was not included in the highly valid compound, was .75. In a third control environment, the validity of the compound AB was .50, as was the validity of each of its component cues. Just as in the complex environments, the validity of cue C was .75; that is, this was the only individual highly valid cue.

We also manipulated participants' causal knowledge about the cues via the experimental instructions. Particularly, we differentiated between a configural, an elemental, and a neutral causal model condition. In the configural causal model condition, the instructions emphasized that cues act through the same causal mechanism and at the same point in time in bringing about the effect (i.e., the criterion value). In the elemental causal model condition, the instructions emphasized that the cues acted through different causal mechanisms and at different points in time in bringing about the effect. Finally, in the neutral causal model, participants did not receive any information about the possible causal mechanisms through which the cues acted. In short, having in mind a configural, an elemental, or a neutral causal model, participants received an XOR environment, an AND environment, or a control environment.

To classify participants according to a particular strategy in our experiments, we used the Bayesian method for multiple-attribute decision making proposed by Bröder and Schiffer (2003). Since TTB and TTB-Configural showed the highest fit for most of our participants, we will focus on only these two strategies. Table 1 shows the percentage of participants for whom one of these strategies achieved the highest fit. In accordance with our hypothesis, a high percentage of participants decided according to the highly valid compound cue, using the TTB-Configural heuristic, when they had a configural causal mental model of the environment, that is, when causal knowledge suggested that the component cues acted through a common causal mechanism. Furthermore, even if there was a highly valid compound in the environment, a high percentage of participants represented its component cues as independent elements and used the elemental TTB heuristic when these cues acted through different causal mechanisms (in the elemental causal model), but also spontaneously (in the neutral causal model). Interestingly, these results were found regardless of whether the component cues were amalgamated into a compound by applying the XOR or the AND logical rule. Finally, when there was no highly valid compound cue in the environment, that is, in the control condition, the elemental TTB was also the most frequently used strategy.

Table 1: Percentage of participants for whom TTB-Configural and TTB achieved the highest fit in a given causal mental model condition.

	TTB-Configural Participants			TTB Participants		
	Config. Model	Elem. Model	Neutral Model	Config. Model	Elem. Model	Neutral Model
XOR	50	8.3	8.3	16.6	58.3	66.7
AND	66.7	8.3	16.7	25	50	41.7
Control	0	0	0	66.7	100	83.3

In sum, results in these series of experiments suggest that TTB-Configural is used when the information structure in the environment and in the mind fit together, that is, when the causal knowledge about the cues induced participants to search for highly valid compound cues in the environment and a real highly valid compound cue existed in the physical world. When this was not the case, the elemental TTB was the best behavioural model in our experiments.

Conclusions

The results from the simulation and from the experiments complement each other: When there are valid compound cues in the environment, TTB-Configural, an extension of TTB that includes such cues in its cue hierarchy, outperformed—in cross-validation—both TTB and even another extension of TTB that included all possible compounds. Moreover, TTB-Configural performed well as a behavioural model: If, and only if, there was a valid compound cue in the environment and participants received a hint that cues may interact with each other, then they selected this heuristic from their adaptive toolbox.

References

Ahn, W., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 199-225). Cambridge, MA: MIT Press.

Baumeister, J., & Seipel, D. (2002). Diagnostic Reasoning with Multilevel Set-Covering Models. Proceedings 13th. International Workshop on *Principles of Diagnosis (DX-02)*, Semmering, Austria

Bröder, A., & Schiffer, S. (2003). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, 16, 193-213.

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

Garcia-Retamero, R., Dieckmann, A., Hoffrage, U., & Ramos, M. (2005a). *Compound cue processing within the fast and frugal heuristic approach in non-linearly separable environments*. Submitted for publication.

Garcia-Retamero, R., Hoffrage, U., Dieckmann, A., & Ramos, M. (2005b). *Do people use fast and frugal heuristics to process compound cues in linearly separable environments?* Submitted for publication.

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

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

Kehoe, E. J., & Graham, P. (1988). Summation and configuration: Stimulus compounding and negative patterning in the rabbit. *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 320-333.

Lagnado, D. A. and Sloman, S. A. (2002). Learning causal structure. In W. Gray & C. D. Schunn (Eds.) *Proceedings of the twenty-fourth annual conference of the cognitive science society*. Erlbaum.

Minsky, M., & Papert, S. (1969). *Perceptrons: An introduction to computational geometry*. Cambridge, MA: MIT Press.

Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, 111, 455-485.

Shanks, D. R., Charles, D., Darby, R. J., & Azmi, A. (1998). Configural processes in human associative learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 24, 1353-1378.

Shepard, R. N., Hovland, C. L., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, 75 (13), Whole no. 517.

Smith, J. D., Murray, M. J., & Minda, J. P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23, 659-680.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak, & D. L. Medin (Eds.), *The Psychology of learning and motivation* (Vol. 34, pp. 47-88). San Diego, CA: Academic Press.

Waldmann, M. R., & Hagmayer, Y. (in press). Seeing versus Doing: Two models of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, 121, 222-236.

Walker, E. F., & Diforio, D. (1997). Schizophrenia: A neural diathesis-stress model. *Psychological Review*, 104, 667-685.