From Causal Models to Sound Heuristic Inference

Ana Sofia Morais (morais@mpib-berlin.mpg.de)[1]
Lael J. Schooler (schooler@mpib-berlin.mpg.de)[1]
Henrik Olsson (h.olsson@warwick.ac.uk)[2]
Björn Meder (meder@mpib-berlin.mpg.de)[1]
[1]Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
[2]Department of Psychology, University of Warwick, Coventry CV4 7AL, UK

Abstract

We investigate whether people rely on their causal intuitions to determine the predictive value or importance of cues. Our real-world data set consists of one criterion variable (child mortality) and nine cues (e.g., GDP per capita). We elicited people's intuitive causal models about the domain. In a second task, we asked them to rank the cues according to their beliefs about the cues' predictive value. Alternative cue importance rankings were derived directly from their causal models using measures of causal centrality. The results show that people's judgments of cue importance corresponded more closely to the causal-based cue orders than to the statistical associations between the cues and the criterion. Using computer simulations, we show that people's causal-based cue orders form a sound basis for making inferences, even when information about the statistical structure of the environment is scarce or unavailable. Central to the simulations is take-the-best (TTB)—a simple decision strategy that makes inferences by considering cues sequentially. The simulations show that causal-based cue orders can be as accurate as individuals' judged orders. Causal-based cue orders allow TTB to perform as would be expected from estimating the weights of a linear model using about 35% of the available data. These findings suggest that people can rely on their causal intuitions to determine the importance of cues, thereby reducing the computational complexity involved in finding useful cue orders.

Keywords: Causal models; simple heuristics; take-the-best; cue orders; information search; inductive inference

Introduction

Simple heuristics that incorporate ideas of bounded rationality—such as considering very little information—can be surprisingly efficient and robust compared with models that rely on more information and complex computations (Todd, Gigerenzer, & The ABC Research Group, 2012). The take-the-best (TTB) heuristic, for example, looks for pieces of information (i.e., cues) in order of their predictive value (i.e., importance) and makes decisions based on the first cue that distinguishes between the options (Gigerenzer & Goldstein, 1996).

The question of how people find good cue orders is important to assess the psychological plausibility of simple heuristics like TTB. In fact, one criticism that has been leveled against TTB is that it is only seemingly simple, because it freeloads on the effort hidden in the computation of the cue order (e.g., Dougherty, Franco-Watkins, & Thomas, 2008). One response to this criticism is that natural selection and social learning can produce useful cue orders (Gigerenzer, Hoffrage, & Goldstein, 2008). Moreover, people's intuitions about the direction of the relation between the cues and the criterion can curtail the computational complexity involved in ordering cues, while maintaining good performance (Katsikopoulos, Schooler, & Hertwig, 2010).

We investigate another way of ordering cues by connecting two lines of research that are seldom considered together: causal reasoning and heuristic decision making. Our hypothesis is that people may rely on their intuitive causal models to order cues. We test this hypothesis by combining behavioral data and computer simulations. The behavioral study investigates how people's causal models relate to their judgments of cue importance, above and beyond the cues' predictive value in the environment. The simulations examine the usefulness of cue orders derived from people's causal models for making inferences with TTB using a real-world data set.

Intuitive causal models

Causal models can be represented as directed graphs, where the nodes denote the domain variables and the links represent the causal dependencies. Such representations mirror a characteristic property of our environment, namely, that some events, causes, can generate or prevent other events, their effects (Waldmann, Hagmayer, & Blaisdell, 2006; see Meder, Mayrhofer & Waldmann, in press, for a formal treatment of causal networks).

An intuitive causal model is a qualitative representation of a person's subjective beliefs about how the variables in a domain are causally related to each other. Intuitive causal models are not necessarily veridical. Rather, they mirror people's naïve assumptions about the causal structure of the environment, and how that structure gives rise to the observed data. One way of eliciting intuitive causal models is to present participants with a set of variables and ask them to indicate the presence and strength of the causal relations between the variables (e.g., Kim & Park, 2009; Sloman, Love, & Ahn, 1998). The elicited causal models can then be used to estimate a cue's relative importance as a function of its role in the network. Past research has shown that these measures account for people's judgments of cue importance in category-related judgments (Kim & Ahn, 2002; Rehder & Kim, 2006; Sloman, et al., 1998).

This research has recently been extended to investigate whether people's causal models may also determine their information search behavior. Using a categorization task, Morais, Olsson, and Schooler (2011) showed that participants queried structurally important features more frequently and earlier in search. This finding highlights the interplay between individuals' naïve assumptions about the world's causal structure and sequential search strategies.

Research questions

Previous research, however, did not consider the relation between the cues' actual predictive power according to some statistical measure (e.g., correlation with the criterion) and people's beliefs about the importance of cues. Another open question is how useful causal-based cue orders are when used with simple heuristics like TTB, compared to ordering cues according to some statistical measure. To address these questions, we use a real-world data set consisting of one continuous criterion (child mortality rate in different countries) and nine continuous cues (e.g., GDP per capita; Table 1). The data were taken from official statistics for 191 countries from around the world, provided by sources like the World Bank.

This data set provides the basis for investigating the following questions. First, can people's causal models of child mortality account for their judgments of cue importance, above and beyond the cues' predictive value in the environment? Second, what causal-based measure of cue importance best accounts for people's judgments of cue importance? Third, how does a simple heuristic like TTB perform when searching cues based on people's intuitions (either judged or based on causal models), compared to when ordering cues by their statistical association with the criterion?

From causal models to cue orders

How can cue orders be derived from causal models? We use measures of causal centrality that quantify the cues' importance as a function of their role in the causal model. The resulting centrality values can be used to order the cues by their relative importance. In addition to two measures from the categorization literature, we propose simplified accounts for determining causal importance. These measures are motivated by decision making research and have a heuristic-like flavor in that they use less information from the causal model and rely on simpler computations.

Number of Direct Relations People often consider a variable as more central or important to the extent that it is involved in a high number of direct causal relationships, regardless of the strength or the direction of the relations (Ahn, Kim, Lassaline, & Dennis, 2000; Rehder & Hastie, 2001). Based on this finding, the first centrality measure quantifies a variable's importance according to the total number of direct causes and effects that it has.

Weighted Number of Direct and Indirect Effects People also tend to judge a variable as being more important when it has a strong influence on many other variables in the network, via direct or indirect causal relations (Ahn et al., 2000; Sloman et al., 1998). In line with this finding, Sloman et al. (1998) proposed a measure of causal centrality that quantifies the importance of a variable according to the number of direct and indirect effects that it causes, weighted by the strength of these relations.

Unit-Weighted Number of Direct and Indirect Effects This measure simplifies the Sloman et al. (1998) model in that it considers the number of direct and indirect effects, but ignores the strength of those relations. This idea resembles unit-weight linear models in decision theory, which have been shown to yield surprisingly accurate predictions (Dawes, 1979). Based on the evidence that weights often do no help much, this measure quantifies the causal importance of a variable as a function of its number of direct and indirect effects, regardless of the strengths of the links with those effects.

Weighted Number of Direct Effects The second naïve measure discards indirect effects. Computationally simpler, this measure calculates causal importance as the sum of the strengths of the causal links that a variable has with its effects, divided by the total number of direct effects that the variable has. Thus, the variable's causal importance increases with the average causal influence that it has on its direct effects.

Unit-Weighted Number of Direct Effects The last measure ignores causal strengths and indirect effects. The causal importance of a variable is measured by the number of direct effects that it has: the higher the number of effects, the more important the variable. This measure corresponds to the concept of out-degree in network analysis.

Can people's causal models account for their judgments of cue importance?

We elicited people's intuitive causal models of child mortality before age five. The measures described above were then applied to participants' causal models to derive alternative predictions about participants' beliefs about cue importance. To explore the relation between the causal models and people's explicit judgments of cue importance, we asked them to indicate in what order they would query the cues, if they had to predict the child mortality rate (i.e., the probability of a child dying before age five) of an unknown country. To disentangle the effects of causal models and environmental statistics, the causal-based and judged cue orders were compared to the cue orders implied by the cue-criterion correlations.

Method

Participants Seventy participants (mean age 25 years, 33 female) participated in the study for 10 euros.

Child mortality data set We constructed a data set for the domain of child mortality, based on real-world statistics

Table 1: Correlations between cues and child mortality.

Cue	*r*
% of the population living below the poverty line	.80
% of births attended by skilled health personnel	-.76
% of children enrolled in primary school	-.75
% of infants with low birth weight	.58
% of the population living in cities	-.55
Gross Domestic Product (GDP) per capita	-.39
% of adults infected with HIV	.31
Population growth	.24
% of GDP spent on healthcare	-.10

Figure 1: A participant's causal model of child mortality.

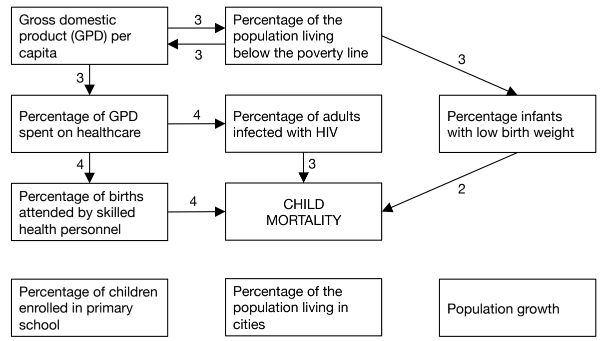

made available by the United Nations Children's Fund (UNICEF), the World Bank, the Joint United Nations Program on HIV/AIDS (UNAIDS), and the Central Intelligence Agency (CIA) World Factbook. The data set consists of ten continuous variables, including the criterion (child mortality before age five) and nine cues that can be used to predict child mortality rates in 191 countries. In a pilot study, we found that people were able to come up with a causal model of the domain, indicating that they have intuitions about how the variables are causally related. Table 1 shows the cues and the Pearson correlation coefficients for the relation between each cue and the criterion across the countries. The correlations were computed after replacing the missing values in the data by a method of multiple imputation (Honaker, King, & Blackwell, 2011). The imputation process preserved the strength of the cue-criterion correlations in the original data set.

Procedure The experiment consisted of two tasks: a causal model task and a cue-ranking task. The order of the tasks was counterbalanced across participants. The causal model task elicited participants' intuitive causal models using the software ConceptBuilder (Kim & Park, 2009). Participants were asked to draw a diagram of how the cues presented in Table 1 are causally related to child mortality, and how the cues are causally related to one another. Note that participants were not presented with any data through which they could learn the statistical structure of the environment.

First, participants were presented with the criterion variable and the cues on the screen, organized randomly in two rows. Then they read a one-page glossary with the definitions of the variables; the glossary was available throughout the course of the experiment. Subsequently, participants were instructed how to draw a causal model using the software. For every variable X that the participants considered to be causally related to a variable Y, they were asked to draw a directed link between the two variables pointing from cause to effect (e.g., $X \rightarrow Y$). A causal relationship was said to hold whenever a variable X causes or influences a variable Y. Participants were also informed that a variable can influence and be influenced by multiple variables, and that two variables can mutually influence each other. They were also told that if they considered a variable to be causally independent of all other variables, then it should not be connected (i.e., left without any incoming or outgoing links). After a link was drawn, participants were prompted to indicate the strength of the relation on a scale from 1 (= very weak) to 5 (= very strong). Figure 1 gives an example of a participant's causal model.

In the cue-ranking task, each participant was presented with the cues in Table 1, listed in randomized order. Their task was to order the cues according to their usefulness for predicting the child mortality rate of an unknown country. Specifically, participants were asked to indicate in what order they would ask the experimenter about each of the nine cues, one at a time, if they had to predict the child mortality in an unidentified country. The most important cue should be given a rank of 1; the second most important should get a rank of 2, and so on. In addition, for each cue participants were asked to indicate whether it is positively or negatively correlated with the criterion, or whether no relation exists between the two variables. The instructions provided participants with brief qualitative definitions, illustrated with examples, to clarify the concept of a correlation. No explicit relation was drawn between the causal model task and the ranking task. That is, participants were not prompted to judge the importance of cues based on the causal model and vice-versa.

Results and Discussion

Eight participants gave explicit judgments in which more than one cue was given the same rank, although the instructions asked them to rank the cues if they had to query them sequentially in order to make an inference. Three participants drew causal models in which the criterion was not influenced by any of the cues, suggesting a misunderstanding of the causal model task. These 11 participants were excluded from our analyses, resulting in a sample of 59 participants. We applied the alternative measures of causal centrality to each participant's causal model to derive the individual cue orders. Our analyses of the causal-based cue orders exclude the criterion variable.

How well do people's judgments and causal-based cue orders conform to the statistical cue order? We used the correlations in Table 1 to derive a cue order that reflects the

Table 2: Median correlations between cue orders.

	Cue order	
	Statistical	Judged
Judged cue orders	.23	
Causal-based cue orders		
Direct relations	.05	.42
Direct + indirect effects:		
weighted	.10	.30
unit weights	.13	.23
Direct effects only:		
weighted	-.03	.63
unit weights	.09	.44

Figure 2: Predictive accuracy of variants of TTB and linear regression.

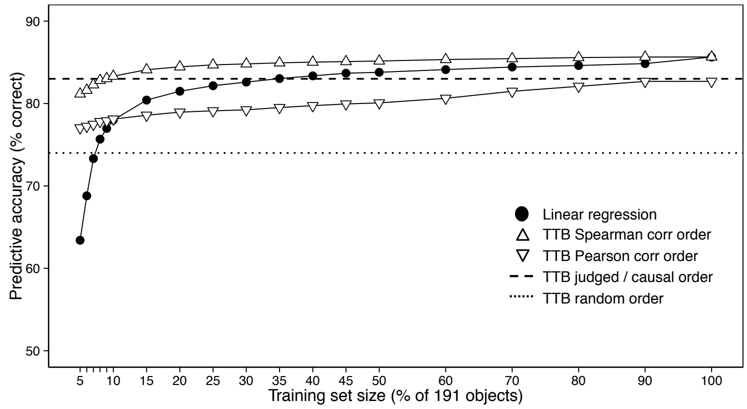

cues' predictive value in the child mortality data set. Table 2 (left column) shows the median Spearman rank correlations across participants between the statistical cue order and participants' judged cue orders, as well as between the statistical order and the alternative causal-based cue orders. The results show that participants' judged cue orders were moderately correlated with the cue order in the data set, indicating that participants have some intuitions about the predictive value of the cues. The orders derived from their causal models, however, did not conform to the cue order given by the cue-criterion correlations.

Do people's intuitive causal models account for their judgments of cue importance? Table 2 (right column) shows the median Spearman correlations across participants for the relation between the judged cue orders and the orders derived from the causal models. The direction of the correlations suggests that people's cue importance judgments varied as a function of the cues' role in the causal model: the more central a cue was in the network, the more useful for predicting child mortality people judged it to be. Among the causal centrality measures, the number of direct effects weighted by causal strength was the best predictor of people's judgments: cues that exert a strong direct influence on other cues were believed to be more important to predict child mortality. Thus, our simplified version of the Sloman et al. (1998) measure that ignores indirect effects provided the best account of people's judgments.

Summary Our results show that people's intuitive causal models accounted for their beliefs about the usefulness of cues, above and beyond the cues' statistical association with the criterion. A naïve centrality measure that only considers a cue's influence on its direct effects and discards indirect effects accounted best for people's intuitive judgments of cue importance.

Can Intuitive Causal Models Be the Basis for Sound Inference?

The finding that people's causal models account for their beliefs about cue importance raises the question of how useful those causal intuitions are for guiding cue search in heuristic inference. We address this question through computer simulations, using the child mortality data set and a paired comparison task in which the goal is to predict which of two unknown countries has the higher child mortality rate (the criterion), based on the nine cues shown in Table 1.

We began by assessing the predictive accuracy of the take-the-best (TTB) heuristic when searching cues in order of the cue-criterion correlations in the child mortality data set. Linear regression, with the cue weights estimated from the data set, served as benchmark. Next, we examined the performance of TTB when ordering cues according to participants' judged cue orders or causal-based orders.

Predictive Accuracy of TTB Using Statistical Cue Orders

TTB and linear regression differ substantially in how they process cues and make decisions. Regression estimates the weights of a linear model from the available data. To decide which country has the higher child mortality rate, the model combines the weighted cue values in a linear-additive fashion and selects the alternative with the higher estimate.

Unlike regression, TTB does not weigh and add cues. Instead, it searches cues in order and makes a decision based on the first cue that discriminates between the alternatives, ignoring all other cues (Gigerenzer & Goldstein, 1996). If the discriminating cue is positively related to child mortality, TTB predicts that the country with the higher cue value has the higher child mortality rate. If the relation is negative, the country with the lower cue value is considered to have the higher rate.

Procedure We investigated the predictive accuracy of TTB and linear regression through cross-validation, using the child mortality data set. In each simulation round, we split the data randomly into two parts: the training set and the test set. The training set was used to estimate the parameters of both models. For TTB, the cue order and the cue directions were derived from the cues' correlations with the criterion using the training data. We implemented two variants of TTB: one that searches cues according to their Pearson correlations with the criterion, and one that is based on the cue-criterion Spearman rank correlations. As our behavioral experiment was designed and analyzed with both Pearson and Spearman rank correlations, we use both measures in our simulations. For regression, we used the method of ordinary least squares to estimate the cues' weights from the training data. Model performance was evaluated by making pairwise inferences for all remaining objects, the test set.

Since more complex models are often less robust when making inferences based on small samples (Katsikopoulos et al., 2010), we varied the size of the training set, ranging from 5% to 95% of the 191 countries, in steps of 5%. Between 5% and 10%, we used an increase of 1% to assess the models' performance when making inferences from minute samples. We also evaluated the models' accuracy when the training set consisted of all objects in the data set, although this only indicates the models' capacity to fit the data, not their predictive accuracy. The simulation was repeated 5,000 times for each training set size.

Results Figure 2 shows the models' mean predictive accuracy, defined as the proportion of correct inferences in the test set. Across all training sets, TTB achieved a higher predictive accuracy when ordering cues by their Spearman rank correlation with the criterion than when ordering cues by the cue-criterion Pearson correlations. When based on Spearman correlations, TTB outperformed linear regression regardless of the size of the training set; yet regression reached the same level of performance when the full data set was used to both fit and test the models. When ordering cues according to the Pearson correlations, TTB outperformed regression for small training samples, but had lower predictive accuracy for larger samples. Overall, these results are consistent with previous work showing the high predictive accuracy of TTB when making inferences based on small samples (Katsikopoulos et al., 2010). Moreover, our findings indicate that TTB performs quite well in our data set, but that the heuristic's performance also varies as a function of the measure that it uses to order cues.

The usefulness of intuitive causal models

Do intuitive causal models yield useful cue orders for making inferences, when information about the statistical structure of the environment is scarce or not available? We tackled this question by examining how TTB performs when using people's judgments of cue importance or the cue orders derived from their causal models. In this case, the cue orders and the cue directions were not estimated from the training set, but fixed a priori based on the results of our behavioral study. Predictive accuracy was evaluated across 5,000 simulation runs, based on the same test sets used to compare the two TTB variants and linear regression.

We investigated four different implementations of TTB. All versions used the same cue directions, derived from the cue-ranking task in the behavioral study. We computed the aggregate cue directions by counting the number of times a cue was judged to be positively, negatively, or not related to the criterion. The direction that was judged by the majority of the participants for each cue provided the input to TTB. The aggregate subjective directions matched the cue directions in the data set (Table 1), except for the cue "percentage of children enrolled in primary school", which was judged to be uncorrelated with the criterion.

The four variants of TTB used the same cue directions, but different cue orders. The first variant used the judged cue orders derived from the cue-ranking task. We calculated each cue's average rank by taking the mean rank position across participants. The other implementations of TTB used cue orders derived from people's causal models. We considered three measures of causal importance: total number of direct relations, direct and indirect effects weighted by causal strengths, and weighted direct effects. For each measure, we used the cues' mean rank across participants. Finally, we tested TTB with random cue orders to evaluate the contribution of intuited cue directions to the model's performance. This variant of TTB provides a baseline for the performance that it achieves when using intuited cues directions, but searching cues in random order.

Results Since the order and direction of the cues was not estimated from the training sample, performance did not vary with the size of the training set. Therefore, we computed the grand mean across all test sets for each variant of TTB. Interestingly, the different causal-based cue orders achieved approximately the same predictive accuracy as the judged orders, around 83%. These models are represented in Figure 2 by the dashed line. The predictive accuracy of the alternative cue orders was as follows: judged cue orders 82.7%, direct relations 82.8%, direct and indirect effects weighted by causal strengths 82.8%, weighted direct effects 82.5%. Figure 2 further suggests that a similar level of accuracy can be expected from estimating the weights of a linear model based on 35% of the objects in the data set. In other words, cue orders derived from intuitive causal models can yield high accuracy without the benefit of a training set.

How useful are individuals' judged and causal-based cue orders relative to a random cue order? The dotted line in Figure 2 shows the predictive accuracy of TTB when searching cues in random order. Although the model only used people's intuitions about whether a cue is positively or negatively correlated with the criterion, it achieved a considerable accuracy of 73.9%. This is in line with past work showing that ordinary information about cue directions can yield good performance (Katsikopoulos et al., 2010). Even so, intuited cue orders (either judged explicitly or derived from causal models) led to a 9% improvement in predictive accuracy relative to the random cue order.

Why did TTB perform so well when using individuals' judged or causal-based cue orders? TTB makes a decision as soon as a cue is found which discriminates between the objects (i.e., a cue for which the two objects have different values). Since all cues in the child mortality data set are continuous, TTB often makes a decision based on the very first cue that is queried. The cue "percentage of the population living below the poverty line" (Table 1) was identified as the most important cue (at the aggregate level) by the three measures of causal importance, as well as by participants in the cue-ranking task. Similarly, TTB ranks this cue as the first or second most important when ordering cues by Pearson or Spearman correlation, respectively. Thus, people's intuition about the most important cue helped TTB achieve good performance.

Although people's intuition was that the poverty line cue was the best, there were slight differences in how the subsequent cues were ordered, which was reflected in predictive accuracy. In the cue-ranking task, participants identified the cue "percentage of births attended by skilled health personnel" as the second most important cue. Two causal importance measures (direct and indirect effects weighted by causal strengths and number of direct relations) selected "GDP per capita" as the second most important cue. Yet by the weighted number of direct effects, the second most important cue was the "percentage of GDP spent on health". TTB, in turn, examines this cue last when ordering cues by their Pearson or Spearman correlation with the criterion. This explains why cue orders based on the weighted number of directed effects performed a little less well than the other causal-based orders.

Note that the difference in accuracy between the judged cue orders and the orders based on the weighted number of directed effects is not at odds with the behavioral result that this causal importance measure accounted best for people's judgments (Table 2). While the relation between the two cue orders was evaluated by correlating the full orders, predictive accuracy was mostly determined by correctly identifying the first few most important cues.

Conclusion

One criticism leveled against simple heuristics like TTB is that they owe much of their simplicity and success to the complexity hidden in computing a useful cue order. Our results suggest that people may curtail complexity by relying on their causal intuitions to determine the predictive value of cues. First, we showed that people's intuitive causal models accounted for their beliefs about the usefulness of cues, above and beyond the statistical cue-criterion associations. Cues that participants indicated as exerting a strong direct influence on other cues were believed by the participants to be more important. Second, in an inference task, cue orders gleaned from intuitive causal models can be as accurate as people's judgments of cue importance. Causal-based cue orders allowed TTB to perform as would be expected from estimating the weights of a linear model using about 35% of the available data.

Countering the concern that TTB's success freeloads on the effort put into computing the cue order, we showed that causal intuitions may reduce effort without hurting accuracy. A simple heuristic can be robust even when the cue order is garnered from people's intuitive causal models.

Acknowledgments

This research was supported by a research fellowship from the MaxNetAging Research School to ASM, and grant ME 3717/2 to BM from the DFG as part of the priority program "New Frameworks of Rationality" (SPP 1516).

References

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361–416.

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

Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, *115*, 199–211.

Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas. *Psychological Review*, *115*, 230–239.

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

Honaker, J., King, G., & Blackwell, M. (2011). AMELIA II: A program for missing data. *Journal of Statistical Software*, *45*, 1–47.

Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information, *Psychological Review*, *117*, 1259–1266.

Kim, N. S., & Ahn, W. (2002). The influence of naive causal theories on lay concepts of mental illness. *American Journal of Psychology*, *115*, 33–65.

Kim, N. S., & Park, E. Y. (2009). ConceptBuilder: An open-source software tool for measuring, depicting, and quantifying causal models. *BRM*, *41*, 128–136.

Meder, B., Mayrhofer, R., & Waldmann, M. (in press). Structure induction in diagnostic causal reasoning. *Psychological Review*.

Morais, A. S., Olsson, H., & Schooler, L. J. (2011). Does the structure of causal models predict information search? In B. Kokinov, A. Karmiloff-Smith, & N. J. Nersessian (Eds.), *European Perspectives on Cognitive Science*. Sofia: New Bulgarian University Press.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *JEP: G*, *130*, 323–360.

Rehder, B. & Kim, S. (2006). How causal knowledge affects classification: A generative theory of categorization. *JEP: LMC*, *32*, 659–683.

Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*, 189–228.

Todd, P. M., Gigerenzer, G., & the ABC Research Group (2012). *Ecological rationality: Intelligence in the world.* New York: Oxford University Press.

Waldmann, M. R., Hagmayer, Y, & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, *15*, 307–311.