

# Structured Correlation from the Causal Background

Ralf Mayrhofer<sup>1</sup>, Noah D. Goodman<sup>2</sup>, Michael R. Waldmann<sup>1</sup>, and Joshua B. Tenenbaum<sup>2</sup>

<sup>1</sup> {rmayrho, mwaldma}@uni-goettingen.de; Department of Psychology, University of Goettingen, Germany.

<sup>2</sup> {ndg, jbt}@mit.edu; Department of Brain and Cognitive Sciences, MIT.

## Abstract

Previous research has cast doubt on whether the Markov condition is a default assumption of human causal reasoning—as causal Bayes net approaches suggest. Human subjects often seem to violate the Markov condition in common-cause reasoning tasks. While this might be treated as evidence that humans are inefficient causal reasoners, we propose that the underlying human intuitions reflect abstract causal knowledge that is sensitive to a great deal of contextual information—knowledge of the “causal background”. In this paper, we introduce a hierarchical Bayesian model of causal background knowledge which explains Markov violations and makes additional, more fine-grained predictions on the basis of causally relevant category membership. We confirm these predictions using an experimental paradigm which extends that used in previous studies of “Markov violation.”

**Keywords:** causal reasoning; hierarchical Bayesian models; categorization; Bayes nets; Markov condition

## Introduction

Causal Bayes nets are an increasingly popular approach to human causal learning and reasoning (Waldmann & Martignon, 1998; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Griffiths & Tenenbaum, 2005; Sloman, 2005). However, recent research (see Rehder & Burnett, 2005; Waldmann, Cheng, Hagmayer, & Blaisdell, 2008; Walsh & Sloman, 2007) has cast doubt on the Markov condition, a central assumption of this approach. This condition states that a variable, conditioned on its causes, is independent of all variables except its effects (Pearl, 2000; Spirtes, Glymour, & Scheines, 2000). A cause thus “screens off” its effects: for example, in a simple common-cause network (Fig. 1a) inferences about the state of an effect variable,  $E_3$ , depend only on the state of its cause  $C$  and not on the states of the other effects ( $E_1$  and  $E_2$ ). However, it is not difficult to create scenarios in which our intuitions seem to violate this condition: Suppose one day while in your favorite cafe you badly need to check your email, but your laptop is unable to connect to the wireless network. From the past you know that e-mail does not always work, but there is a fairly high probability that it does when the transmitter is on. You see (from the indicator lights) that the transmitter is on, so you ask your neighbor if his computer is working, but it also isn’t connecting. You ask four more people and get the same answer. Will you try another person? You probably expect that the seventh computer will be no different than the first six—also not working. If we take the causal structure of this situation to be the overt one—the transmitter independently causes each computer to connect to the internet—then the Markov condition implies that the failure of any number of computers should not affect the probability that a final computer can connect. Thus, your intuition that six instances of failure indicate that a failure in the sev-

enth case is likely is a case of a “Markov violation”<sup>1</sup>. Such “Markov violations” in common-cause scenarios were experimentally documented by Waldmann, Mayrhofer, and Hagmayer (2007). Earlier work by Rehder and Burnett (2005) explored analogous “Markov violations” in reasoning about causally related properties of object categories.

One interpretation of this finding is that the competency of human causal reasoning is well described by causal Bayes nets, but that humans are simply not efficient causal reasoners. Another interpretation is that the causal Bayes net formalism is not able to capture human causal reasoning, and should be abandoned. We offer a third interpretation: that causal Bayes nets can be used to describe human causal reasoning, but that the overt causal structure must be augmented to account for knowledge of the *causal background structure*. The causal background is the sum of all of the unobservable causal influences potentially affecting a specific causal system. This goes beyond overt causal structure that is explicitly known in a given situation, summarizing the additional subtleties of the real system. We do not propose that people have specific knowledge of causal background factors (though this may sometimes be the case) but that they have abstract knowledge about classes of background factors that may influence causal relations between various classes of events. This knowledge allows people to account for many potential background influences implicitly, in aggregate form, without representing them explicitly or individually. This is related to the notion of the alternative cause in situations of elemental causal induction (Cheng, 1997), however it extends the alternative cause to non-elemental situations in which the alternatives may affect multiple variables.

If this causal background knowledge is a coherent and useful form of abstract knowledge, rather than a mere reflection of inefficiencies in causal reasoning, then we should be able to reveal finer structure within the intuitions of “Markov violation”. Imagine that, in the example above, the seventh computer is of a completely different type than the other six (say, a Mac instead of a PC). In this case you may be much more likely to try once more to check your email, using this new computer. Motivated by this intuition, we consider the effect of categorical knowledge on “Markov violation”. It is well known that categories can influence causal inferences and vice versa (Lien & Cheng, 2000; Waldmann & Hagmayer, 2006; Ahn, 1998; Rehder, 2003). The notion that objects have abstract causal types, which reflect similarities in causal powers, has been formalized in the causal-schemata model (Kemp, Goodman, & Tenenbaum, 2007) within a hierarchi-

<sup>1</sup>We use the quoted term “Markov violation” to indicate simply the empirical phenomenon.

cal Bayesian framework. In this paper we attempt to formalize category-based causal background knowledge in a similar hierarchical Bayesian model. This model then motivates an empirical demonstration which parallels the opening example, showing fine grained sensitivity of “Markov violation” to context.

## Modeling

In this section we build up a hierarchical Bayesian model which captures a reasoner’s knowledge about both overt causal structure, and the structured causal background. We make the simplifying assumption that the overt causal structure, given by specific world-knowledge or experimental instructions, is that of a simple common-cause network (as in Fig. 1a). This structure predicts an influence of the cause on each of the effects, but predicts that the effects are otherwise independent (a consequence of the Markov condition). However, there may be many unobserved factors which could potentially affect any system—preventive causes, missing enabling conditions, etc.—and these factors can cause additional correlations between the effect variables. We will capture these background correlations by augmenting the overt causal structure with an unobserved common noise source that represents all of the additional influences which potentially affect the different effect variables. Further, unobserved factors are more likely to be shared by objects within a category than those in different categories. In our model this kind of sensitivity to category structure is a consequence of a refined common noise source: we include a separate noise source for each category, shared by effect variables pertaining to category members. Finally, we will augment the model to handle both uncertainty over the causally relevant categories, and the role of perceptual features in suggesting relevant categories.

### Background correlation

We begin with the standard common-cause network shown in Fig. 1a, which we assume to be parametrized as in a noisy-or causal structure: the cause  $C$  has power  $w_C$  to bring about each effect independently (see Cheng, 1997). To capture unobserved influences common to all of the effect variables we introduce an additional variable  $PN$  which also acts as a causal parent to all of the effects, preventing  $C$  from bringing about its effects via a noisy-and-not parametrization<sup>2</sup>. With what probability will the final effect  $E_n$  be on, given the states of the common cause  $C$  and the other effects  $E_1 \dots E_{n-1}$  (but not the unobservable background variable  $PN$ )? This is given

<sup>2</sup>It is possible to also include a common generative influence, we have omitted this for reasons of simplicity.

by<sup>3</sup>:

$$\begin{aligned} &P(E_n=1|C, E_1, \dots, E_{n-1}) \\ &= \sum_{PN} P(E_n=1|C, PN, E_1, \dots, E_{n-1})P(PN|C, E_1, \dots, E_{n-1}) \\ &= \sum_{PN} P(E_n=1|C, PN)P(PN|C, E_1, \dots, E_{n-1}) \end{aligned} \quad (1)$$

In the first step we have introduced the hidden variable  $PN$  and marginalized over its possible values; in the second step we have used the independence of the effect variables given all of their parent variables (this is the Markov condition, which once again holds formally after accounting for the background influences via  $PN$ ). The reasoning captured by Eq. 1 can be considered as a two step process: First the state of  $PN$  is inferred, then, given that state, the state of the target effect  $E_n$  is inferred.

### Structuring the background correlation

Let us say now, that we are given the knowledge that the effect variables partition into distinct categories, and that all the relevant unobserved causal factors are specific to a single category. How can we capture this additional structure? Instead of a single hidden variable  $PN$  we will have one such variable  $PN_{z_j}$  for each category  $z_j$  of the partition  $\mathbf{z}$  (see Fig. 1c). Writing  $E_{z_j}$  (excluding  $E_n$ ) for the effect variables in the  $z_j$ th category, we now have:

$$\begin{aligned} &P(E_n=1|C, E_1, \dots, E_{n-1}, \mathbf{z}) \\ &= \sum_{PN_{z_n}} P(E_n=1|C, PN_{z_j})P(PN_{z_j}|C, E_{z_n}) \end{aligned} \quad (2)$$

Notice that the state of  $E_n$  only depends on the states of other variables within its category.

### Uncertainty over relevant categories

However, we may suspect that there are causally relevant categories, but not know what they are. For instance, if the effects pertain to particular objects which have features suggesting separate categories. We may maintain uncertainty over the correct categorization, while still gaining the ability to leverage any categories which present themselves, by placing a prior distribution on categories (this is very similar to the construction of causal schemata in Kemp, Goodman, and Tenenbaum (2007)). We use a Chinese Restaurant Process prior  $P(\mathbf{z})$  on the partition  $\mathbf{z}$  (see Pitman, 2002), and we additionally assume that there are features  $F$  associated with the objects that also depend on the partition through the like-

<sup>3</sup>The learner in our experimental situations doesn’t know the true value of the causal strengths  $w_C$  and  $w_{PN}$ . We place a Beta prior over these causal strengths to represent this uncertainty, and must then integrate over the strengths, giving:  $P(E_n=1|C, E_1, \dots, E_{n-1}) = \int P(E_n=1|C, E_1, \dots, E_{n-1}, \mathbf{w})d\mathbf{w}$ . These integrals all have standard solutions, and we omit them from our exposition for clarity.




Figure 1: a) An example of a simple causal Bayes net, here a common-cause model. The cause variable  $C$  is connected to its effects via causal arrows. The state of each effect node depends only on the state of the cause node. Due to the Markov condition, the effects are independently influenced by  $C$ . b) Modeling the Markov violation by inserting a common preventive noise node  $PN$ . This hidden node interacts with  $C$ . If  $PN$  is present, the causal influence of  $C$  is lowered for all effects. Thus, if  $E_1, \dots, E_{n-1}$  were observed absent, even if the cause is present,  $PN$  likely, and this lowers the probability of  $E_n$ . c) Adding structure to the simple  $PN$ -model. For each cluster effect variables a common preventive noise node is added. This leads to dependent conditional probabilities within clusters but not between clusters.

likelihood  $P(F|\mathbf{z})^4$ . Hence:

$$\begin{aligned}
 P(E_n=1|C, E_1, \dots, E_{n-1}, F) \\
 = \sum_{\mathbf{z}} P(E_n=1|C, E_1, \dots, E_{n-1}, \mathbf{z}) P(F|\mathbf{z}) P(\mathbf{z}) \quad (3)
 \end{aligned}$$

### Model predictions

What inferences do we expect on the basis of this model? In the case without categories (Fig. 1b) this is quite simple. When the common cause is present, if observed effects are off it will provide evidence that  $PN$  is on (preventing the activity of  $C$ ). Hence, the conditional probability of the target effect should decrease with an increasing number of absent collateral effects. That is, we predict a violation of the Markov condition—or rather the Markov condition as it would hold for a simple model lacking  $PN$ . In the case that the common cause is absent, the conditional probability of the target effect will be independent of the states of the collateral effects (since the preventer has nothing to prevent). In this case the probability of the target effect should only depend on its base rate.

Whenever a causally relevant category is known with certainty (Fig. 1c), the same predictions apply with respect to those effects belonging to the same category, but not to the effects of the other category. That is, we predict only *within category* “Markov violations”.

When there is uncertainty over the relevant categories the influence of the  $PN$  nodes is marginalized out over all possibilities. This leads to graded influences, since there will be some weight on partitions that puts any group of effects together. The more probable a partition given the features, the higher the influence of category members within this category on each other. This leads to the prediction that the more similar two effects are, the more dependent their conditional probabilities should be.

<sup>4</sup>We take the dependence of features on category to be given as in the infinite mixture model (see Pitman, 2002)—a model of categorization related to Anderson’s rational model (Anderson, 1991). Other feature-based categorization models could be substituted.

## Experiment

We wished to test whether the “Markov violation” phenomenon in human causal inference is sensitive to category structure, as predicted by the model. For this purpose we used a version of the “mind sending” task used by Waldmann et al. (2007), which adapted the mind reading alien task of Steyvers et al. (2003) and combined it with the reasoning task introduced by Rehder and Burnett (2005). In this task, subjects were presented with four aliens, Gonz, Murks, Brxxx, and Zoohng, who mostly think of nothing but sometimes think of “POR” (“food” in the alien language), further, Gonz is able to transmit its thoughts of POR to the other aliens. Thus the thoughts of Murks, Brxxx, and Zoohng are statistically dependent on the thoughts of Gonz. In the test phase subjects rated the conditional probability of a target alien, e.g. Murks, thinking of POR given the thoughts of the other aliens. The results showed that participants’ estimate of the target effect’s probability are—in this case—strongly influenced by the thoughts of the other effect aliens.

The model predicts that the presence/absence of events which are similar to the target event have more impact on estimated probability than the presence/absence of events dissimilar to the target event. To investigate the influence of category structure we adapted the task of Waldmann et al. (2007), by increasing the number of effects aliens and introducing a (binary) color feature for the effect aliens. As in the original task, subjects had to reason about the thoughts (POR/nothing) of a target effect alien in a common-cause structure with respect to the thoughts of the common cause alien and the other effect aliens.

### Method

**Participants** 60 students from the University of Goettingen (33 female, 27 male; mean age 25.1 years) participated in exchange for candy.

**Procedure and Materials** In the instruction phase we presented participants with instructions about six aliens: Gonz,

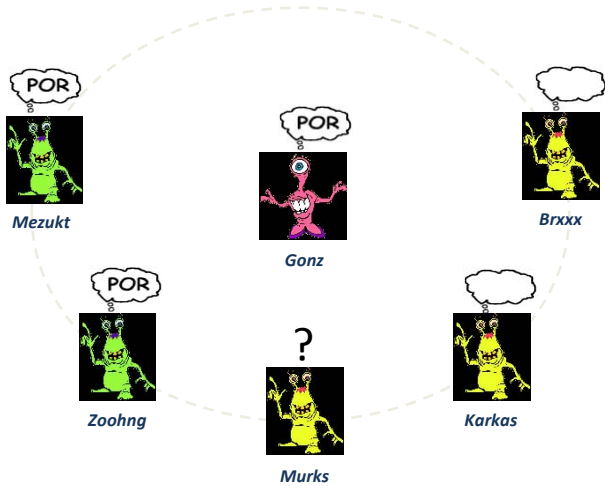


Figure 2: An example of a test panel used in the experiment. Experimental manipulation varied the thoughts of the mind sending alien Gonz (POR or nothing) and the number of aliens of each color thinking POR. As dependent measure subjects were requested to rate the probability that the target alien (indicated by a question mark above its head) thinks of POR in the given situation. (Thus 0, 1, or 2 aliens of the same color as the target alien might successfully think of POR, and 0, 1, or 2 aliens of the other color.)

Murks, Mezukt, Zoohng, Brxxx, and Karkas, who usually think of nothing (illustrated by an empty bubble) and sometimes think of POR (indicated by a bubble containing “POR”). Participants were told that Gonz is able to transmit its thoughts into the heads of the other aliens (i.e., Gonz served as cause in a common-cause structure with five effects). It was emphasized that the effect aliens frequently think of POR if Gonz thinks of POR, and that Gonz isn’t perfect, which means that he sometimes fails to transmit his thoughts due to inattention; the four possible cases with two aliens thinking of POR or nothing were shown as illustrations. In addition, it was pointed out that there were two kinds of effect aliens and that these kinds are indicated by the color of the aliens: green and yellow; but it was not specified what influence this might have (leading to uncertainty about the relevance of this feature). Identical pictures were used for all effect aliens—with two different colors—in order to exclude influences of implicit features. There were always two green collateral effect aliens and two yellow collateral effect aliens. The target effect alien’s color was counterbalanced across subjects. In the test phase subjects were confronted with several test panels with all the non-target aliens thinking of POR or nothing (for an example see Fig. 2). For each panel, subjects were requested to imagine ten situations with the given configuration and then to judge, in how many of these situations the target alien—Murks—would probably think of POR. (It was pointed out that this is merely a proba-

bility judgment.)

**Design** The predictions were tested in a 2x3x3 repeated measurement ANOVA design, with the following factors: the thoughts of the cause alien (nothing vs. POR, i.e.,  $C=0$  and  $C=1$ ); the number of effect aliens of the same category (i.e., same color) as the target alien that are thinking of POR (0, 1, or 2—three levels); the number of effect aliens of the opposite category (i.e., other color) thinking of POR (0, 1, or 2—also three levels). Because this would lead to 18 test cases, subjects were tested on both levels of the  $C$  factor in a balanced partial confounded randomized block design (see Kirk, 1995). So, within the 3x3-combinations on each level of  $C$  only three cases were given to each participant, chosen such that each level of each factor was realized once for each subject (so, three subjects constitute a full data set). The order of the test cases was randomized across subjects.

**Predictions** Let us clarify the predictions of our model applied to this particular task. First, there should be a strong influence of the state of the common cause ( $C=0$  vs.  $C=1$ ), as was found by Waldmann et al. (2007). When the common cause is absent ( $C=0$ ), the conditional probability of the target effect should be nearly independent of the states of other effects (also found by Waldmann et al.). In the case that the common cause is present ( $C=1$ ), and hence the preventive noise sources come into play, the model predicts that a strong “Markov violation” should be found with respect to aliens with the same color; a smaller violation should be found with respect to the differently colored aliens. The latter prediction follows because, although we have only one feature, there is uncertainty about the causal relevance of this feature, and the possible existence of causally relevant categories is not betrayed by obvious features.

## Results

The main prediction is about different effect sizes of the main effect (i.e., different amounts of “Markov violation”), hence the results are displayed as marginalized main effects in Fig. 3. The left hand side of the figure shows the data, the right hand side the model predictions<sup>5</sup>.

In general, the ratings are much higher when the cause alien thinks of POR, which reveals a strong influence of the  $C$  factor, as expected from causal Bayes net theory. If the cause alien thinks of nothing (the lower two lines), the number of effect aliens thinking of POR has no influence on the ratings; the number of same colored effect aliens thinking of POR has no effect ( $F = 1.91$ ), and neither does the number of different colored effect aliens thinking of POR ( $F < 1$ ). This is in line with the model predictions. If the cause alien

<sup>5</sup>To compute the predictions a Monte Carlo simulation with 50,000 repetitions was performed. Only the background rate of the effects and the general strength of the preventive noise nodes were adjusted by hand and set to a Beta(1,3) and a Beta (5,1) distribution. The other parameters of the model were set to values given by the simple model applied to the data of Waldmann et al. (2007) and by Kemp, Goodman, and Tenenbaum (2007).

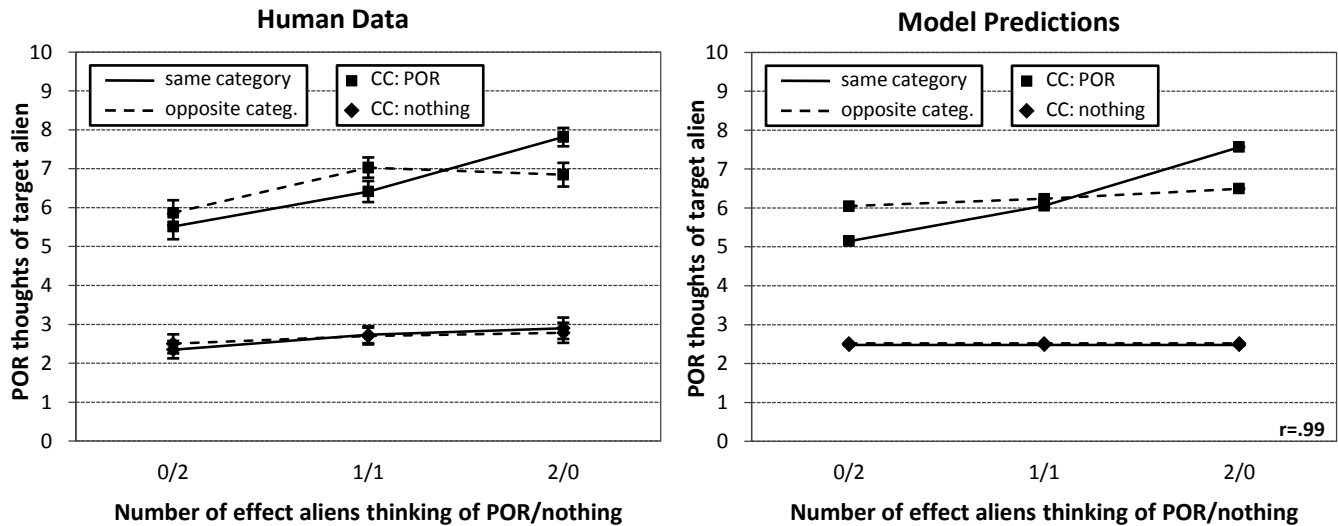


Figure 3: Mean rating and standard error of the number of times the target aliens thinks of POR plotted against the number of times collateral aliens think of POR (in the left panel the experimental data, in the right panel the model predictions). The upper two lines in each panel correspond to the sending alien thinking of POR, the lower two lines to the sending alien thinking of nothing. Results shown as marginalized main effects: The solid lines average over the number of opposite category aliens thinking of POR (showing in column 1/1, for instance, the average ratings over conditions in which the number of opposite-same aliens thinking of POR was: 0-1, 1-1, and 2-1.) The dashed line averages over the number of same category aliens. The columns indicate the number of effect aliens within the relevant group thinking of POR. (Note that the solid and dashed series thus represent the same data, averaged in two different ways in order to examine the effects within each category separately. The crossing of the lines thus cannot be interpreted as an interaction.)

thinks of POR and can therefore transmit its thoughts (the upper two lines), both types of effect aliens have an influence. The number of same category effect aliens thinking of POR reveals a stronger effect ( $F_{2,112} = 22.53, p < 0.0001$ ) than the number of opposite category effect aliens thinking of POR ( $F_{2,112} = 6.60, p < 0.01$ ). (The interaction of these factors isn't significant ( $F_{4,112} = 1.44, p = 0.224$ )<sup>6</sup>, as predicted by the model.)

To test the main hypothesis—whether the number of same color effect aliens thinking of POR has the greater influence on the subjective conditional probability of the target effect—a test must be constructed, because no standard procedure is available. The interesting question is, whether the same category factor explains more variance than the opposite category factor, which simply means that it has more influence on the dependent measure. So, an appropriate statistic for testing this is the ratio of the variances explained by the first and the second factor. The empirical variance ratio is 3.413. To test its significance, we constructed a bootstrap distribution of this variance ratio under the null hypothesis that the variances are equal (i.e., the ratio is 1). For this purpose we drew

<sup>6</sup>Because Figure 3 only shows marginalized main effects, the interaction cannot be obtained by simple examination. The slopes in the figure display the different strengths of the factor's influence on the ratings (the higher the slope the more variance is explained by the corresponding factor).

bootstrap samples from the data under the assumption that the factors are exchangeable (i.e., that color doesn't matter)<sup>7</sup>. The p-value of the empirical variance ratio with respect to the bootstrap distribution is 0.026. Hence, the number of same color effect aliens thinking of POR has a greater influence on the subjective conditional probability rating of the target effect than the number of different color aliens thinking of POR. This is in line with the main prediction of the model.

## Discussion

Our experimental results replicate the finding of “Markov violation”, but also shows that this phenomenon is sensitive to subtle contextual factors—“Markov violations” are stronger within categories of effects than between them. This is consistent with the idea that people have abstract intuitions about influences of the underlying causal background structure. Other current theories of human causal reasoning suggest that Markov violations express irrationalities of the human mind, and offer no compelling account of the sensitivity of these violations to context. In contrast, we have developed a hierarchical Bayesian model that provides a rational account of these findings. The model was designed to capture abstract causal background structure associated with causal types of objects, and does so by extending the causal Bayes nets picture with a form of abstract knowledge.

<sup>7</sup>10,000 bootstrap samples were drawn.

Our experiment was designed to demonstrate that similarity structure affects causal reasoning in “Markov violation” situations. For this purpose we limited the task to a single relevant feature, the aliens’ color. In this setting the ambiguity of causally relevant categories is minimized. By contrast, consider the earlier example, when no computer will connect to the wireless network in your favorite cafe—while you are pushing to finish your CogSci paper. These computers have many features—Macs vs. PCs, laptops with white cases vs. those with black cases, etc. These features do not weigh equally on your intuitions about whether a final computer will work: it’s brand may matter, but probably not it’s size. Our model allows that different features might have different causal relevance for the inference task of interest. Indeed, confronted with evidence the model will learn which features matter with respect to any structured dependencies in the data and which features do not. However, little is known about how and whether people acquire such abstract knowledge about the causal background; further experiments are clearly needed.

We believe that the flavor of causal reasoning studied here—which accounts for causal background structure—is necessary to capture real-world human reasoning abilities. In most natural contexts there are a myriad of influences which cannot each be considered individually, but whose overall influence is significant. Yet it is likely that richer forms of causal background knowledge than we have considered will be needed to capture human intuitions. These might, for instance, include notions of agent-patient relationships or domain restrictions on causal influence. In addition, it is important to consider other modeling approaches to these issues, such as the combination of causality and similarity described in Kemp, Shafto, Berke, and Tenenbaum (2007).

We have suggested that experimental tests showing “Markov violations” do not reflect fundamental human irrationality, but the invocation of an important form of abstract knowledge about the causal background structure. If our interpretation is correct, then a great deal of important psychology is hidden within “Markov violations” and similar phenomena.

### Acknowledgments

The research was supported by a grant of the Deutsche Forschungsgemeinschaft (DFG: WA-621/20-1,2) [RM, MRW] and by the J. S. McDonnell foundation causal learning collaborative initiative [NDG]. In addition, the authors would like to thank Marie-Theres Kater for her assistance in data collection.

### References

Ahn, W. (1998). Why are different features central for the natural kinds and artifacts? the role of causal status in determining feature centrality. *Cognition*, 69, 135-178.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409-429.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51, 354-384.

Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2007). Learning causal schemata. In *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*.

Kemp, C., Shafto, P., Berke, A., & Tenenbaum, J. B. (2007). Combining causal and similarity-based reasoning. *Advances in Neural Information Processing Systems*, 19.

Kirk, R. E. (1995). *Experimental design: Procedures for the behavioral sciences* (3rd ed.). CA: Brooks/Cole.

Lien, Y., & Cheng, P. W. (2000). Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, 40, 87-137.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, MA: Cambridge University Press.

Pitman, J. (2002). *Combinatorial stochastic processes. Lecture notes for Saint Flour Summer School*.

Rehder, B. (2003). Categorization as causal reasoning. *Cognitive Science*, 27, 709-748.

Rehder, B., & Burnett, R. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, 50, 264-314.

Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. Oxford: Oxford University Press.

Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search* (2nd ed.). New York: Springer.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453-489.

Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for bayesian cognitive science*. Oxford: Oxford University Press.

Waldmann, M. R., & Hagmayer, Y. (2006). Categories and causality: The neglected direction. *Cognitive Psychology*, 53, 27-58.

Waldmann, M. R., & Martignon, L. (1998). A Bayesian network model of causal learning. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society*.

Waldmann, M. R., Mayrhofer, R., & Hagmayer, Y. (2007). *Mind reading aliens: Causal capacities and the Markov condition*. Unpublished manuscript.

Walsh, C. R., & Sloman, S. A. (2007). Updating beliefs with causal models: Violations of screening off. In M. A. Gluck, J. R. Anderson, & S. M. Kosslyn (Eds.), *A Festschrift for Gordon H. Bower*. New York: LEA.