

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Testing a Process Model of Causal Reasoning With Inhibitory Causal Links

Permalink

<https://escholarship.org/uc/item/1qg548kt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Rehder, Bob

Davis, Zachary J

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Testing a Process Model of Causal Reasoning With Inhibitory Causal Links

Bob Rehder (bob.rehder@nyu.edu)

Department of Psychology, NYU, 6 Washington Pl, New York, NY, 10003, USA

Zachary J. Davis (zach.davis@stanford.edu)

Department of Psychology, Stanford University, 450 Jane Stanford Way, Stanford, CA, 94305, USA

Abstract

In this paper, we test people’s causal judgments when the graphs have inhibitory causal relations. We find evidence that a particularly important class of errors known as Markov violations extend to these settings. These Markov violations are important because they are incompatible with causal graphical models, a theoretical framework that is often used as a computational level account of causal cognition. In contrast, the systematic pattern of errors are in line with the predictions of a recently proposed rational process model that models people as reasoning about concrete cases (Davis & Rehder, 2020). These findings demonstrate that errors in causal reasoning extend across a range of settings, and do so in line with the predictions of a model that describes the process by which causal judgments are drawn.

Keywords: causal graphical models; causal reasoning; Markov violations; rational process model

Introduction

A large literature exists documenting the profound effect that causal knowledge has on human cognition (Rehder, 2017a; 2017b; Rottman & Hastie, 2014; Sloman & Lagnado, 2015; Waldmann, 2017). Moreover, the theoretical framework known as *causal graphical models* (Pearl, 2009) has been shown to provide a generally good account of human performance in numerous tasks. While causal graphical models have been broadly successful at accounting for human performance, recent work has cast doubt on their success as a model of human causal reasoning. One reason for this doubt is that people systematically violate a core prediction of these models known as the *Markov condition*. In this paper we find new violations of the Markov condition across a range of judgments, suggesting that the failure of CGMs to account for human causal reasoning behavior is more pervasive than previously thought.

To understand the source of these previously undocumented errors in causal reasoning, we compare people’s behavior against the predictions of a recent rational process model of the cognitive mechanisms that underlie many causal judgments (Davis & Rehder, 2020). The first section below describes this model—the “mutation sampler”—and how it explains some classic causal reasoning errors. We then derive new predictions for causal inferences when the reasoner’s causal knowledge includes a mix of both *generative* (a cause makes its effect more likely) and *inhibitory* (a cause makes its effect less likely) causal relations. We then report the results of a new experiment testing these predictions.




Figure 1: (A) A common cause graph with generative causal relations. (B) Possible states of a common cause graph. Filled circles indicate a variable instantiated with a value of 1 (or present), open circles one with a value of 0 (or absent). Edges denote reachable states as defined by the mutation sampler’s proposal distribution $Q(q'|q)$.

The Mutation Sampler

The mutation sampler carries out resource-constrained inference using causal graphical models. It assumes that, when reasoning about a causal system, people think about concrete cases—states of a causal system in which all relevant variables are instantiated with values. Consider the causal graph in Fig. 1A in which variable C is the cause of variables E_A and E_B . Because the variables in this network are assumed to be binary, the *state space* of this graph consists of the eight states shown in Fig. 1B. The mutation sampler assumes that reasoners sequentially *sample* these states using Markov chain Monte Carlo (MCMC) methods, in particular, the Metropolis-Hastings (MH) rule, a computationally efficient rejection sampling method for estimating probability distributions (Hastings, 1970; Griffiths et al., 2015; Van Ravenzwaaij et al., 2018). MCMC methods ensure that the generated samples will, after normalization, approximate the true distribution, with convergence guaranteed as the length of the chain of samples grows large. Whereas MCMC models often deal with a continuous state space, the mutation sampler samples over the discrete states of causal graphs like the one in Fig. 1A.

The mutation sampler makes two additional assumptions. First, MCMC sampling specifies a *proposal distribution*




Figure 2: Joint probability distributions for the common cause graph in Fig. 1A. Graph states are presented on the horizontal axis. The blue line represents the true joint distribution entailed by the causal graphical model under a the parameterization described in Footnote 2. Red lines represent the joint distributions implied by the mutation sampler, with thicker lines meaning fewer samples (thick: 4; medium: 8; thin: 32).

$Q(q'|q)$ that determines which graph state should be proposed as the next state in the chain. The mutation sampler assumes a $Q(q'|q)$ that restricts reachable states q' to those that differ from the current state q by the value of one binary variable. The mutation sampler derives its name from the fact that potential proposals are those formed by “mutating” the value of a single variable. Edges in Fig. 1B denote reachable states from some starting state. Reachable states have an equal probability of being selected as a proposal.

The second assumption is that there is a bias in the starting point for sampling: Sampling always starts from one of the causal graph’s *prototype states*. For example, when the causal relations in Fig. 1A are all generative, the prototype states are those in which nodes are either all present (000) or all absent (111), corresponding to the states at the bottom left and top right corners of Fig. 1B¹. The mutation sampler assumes these starting points because they are guaranteed to be qualitatively consistent with the causal relations. Because the prototypes include no instances in which a cause is present but an effect absent (or vice versa), the reasoners can identify them as consistent with the causal relations without attending to aspects of the causal graph such as the strength, direction,

¹Throughout, a string such as “101” refers to the states of C , E_A , and E_B , respectively, namely that, $C = 1$, $E_A = 0$, and $E_B = 1$.




Figure 3: Average conditional probability judgments from Rehder & Waldmann (2017), in which participants judged the presence of variable E_A on a 0-100 scale. The three judgments on the left are conditioned on C ’s absence whereas those on the right are conditioned on its presence. In each set, E_B is either absent (0), unknown, or present (1). Fits of the mutation sampler (red lines) and the normative model (blue lines) are superimposed on the empirical ratings (gray bars). Details of the fitting procedure are available in Davis & Rehder (2020).

or functional form of the causal relations.

Although the chain of MCMC samples converges to the causal graph’s true joint distribution, the mutation sampler assumes that people’s capacity for sampling is limited and thus take only a few samples (on the order of a dozen rather than thousands or millions). The effects of limited sampling are illustrated in Fig. 2, which presents multiple joint distributions for the common cause graph in Fig. 1A. The blue line represents the graph’s true joint distribution—that is, for each of the eight graph states, the (joint) probability of that state—under a particular graph parameterization.² As expected, the prototype states in which variables are either all absent (000) or all present (111), shown on the far left and right, respectively, are highly probable states. The red lines represent the approximation to the true joint distribution derived from the mutation sampler for three different *chain* (i.e., sampling) lengths: 4, 8, and 32.³ Examination of the mutation sampler’s

²The joint probabilities are generated assuming that $p(C, E_A, E_B) = p(E_A|C)p(E_B|C)p(C)$ and that $p(E_i|C) = 1/(1 + \exp(w_{CE_i}C + w_{E_i}))$ where w_{CE_i} is the strength of the causal relationship between C and E_i and w_{E_i} is the strength of exogenous causes of E_i . In Fig. 2, $p(C) = .5$, $w_{CE_A} = w_{CE_B} = 1$, and $w_{E_A} = w_{E_B} = 0$.

³To make the predictions of the mutation sampler comparable

approximate joint distributions reveals that the probabilities for the prototype states 000 and 111 are greater than those in the true joint distribution, a consequence of sampling beginning at those states. Conversely, the probabilities for the remaining graphs states are less than those in the true joint. That the differences between the true joint distribution and the mutation sampler's approximations becomes small as the chain length increases confirms that MCMC sampling converges to the true joint as the number of samples grows large.

The fact that the mutation sampler reproduces the general shape of the true joint distribution while also introducing distortions explains both why people draw approximately veridical causal inferences and the common errors they make. For example, Rehder and Waldmann (2017, Experiment 2) instructed participants on two generative causal relationships taken from the domain of either economics, sociology, or meteorology that formed a common cause network and then asked them to draw a number of causal inferences. The key inference ratings from that experiment in which subjects predicted the presence of E_A ($E_A = 1$) conditioned on the states of C and E_B are presented in Fig. 3.⁴ On the positive side, subjects correctly judged that E_A was more likely to be present when C was present versus absent (e.g., that $p(E_A = 1|C = 0) < p(E_A = 1|C = 1)$). However, their ratings also violated the defining property of causal graphical models known as the *Markov condition*.

The Markov condition stipulates that the two effects of a common cause graph should be independent conditioned on the cause (e.g., when the state of C is known, the probability of E_A should be unaffected by the state of E_B). That is, knowledge of C *screens off* E_A from E_B . The blue lines in Fig. 3, which represents the best fits of the normative causal graphical model to these data, confirm that the probability of E_B 's presence or absence should not change the probability of E_A when C 's value is known. In contrast, Fig. 3 reveals that participants consistently violated the Markov condition by judging that E_A was more probable when E_B was present (e.g., that $p(E_A = 1|C = 0, E_B = 0) < p(E_A = 1|C = 0) < p(E_A = 1|C = 0, E_B = 1)$), a finding that has been replicated in multiple studies (see Ali et al., 2011; Lagnado & Sloman, Fernbach & Rehder, 2013; Mayrhofer & Waldmann, 2015; Park & Sloman, 2013; Rehder & Burnett, 2005; Rehder, 2014; 2018; Rottman & Hastie, 2016; among others). Importantly, the best fit of the mutation sampler shown in Fig. 3 reveals that these independence violations are predicted by mutation sampling. This prediction is a direct consequence of the distorted joint distribution induced by the biased starting points combined with a relatively small number of samples (e.g., Fig. 2). Stated intu-

to the true joint, the samples it generates have been normalized by dividing the number of visits to each state by the total number of samples.

⁴Because subjects were provided no information that distinguished $C \rightarrow E_A$ and $C \rightarrow E_B$, the causal graph they were given is symmetrical and so inferences regarding the presence E_B are logically equivalent to those regarding E_A (e.g., $p(E_A = 1|C = 1)$ is logically equivalent to $p(E_B = 1|C = 1)$). Thus, the data in Fig. 3 are averaged over both E_A and E_B inferences.

tively, whereas the Markov condition states that E_A and E_B are independent conditioned on C , commencing sampling at one of the two prototype states (000 or 111) introduces a positive correlation between E_A and E_B conditioned on C , one that manifests itself on the conditional probability judgments derived from the sampled joint for short sampling chains.

Predictions for Inhibitory Causal Links

We now derive new predictions for the mutation sampler by varying the *sign* of the causal relations, that is whether they are generative (as they were in Rehder & Waldmann, 2017), or inhibitory. The four experimental conditions shown in Fig. 4 are a result of independently varying the sign of the two causal relationship of a common cause network. It also shows the prototype states that the mutation sampler assumes for each graph. The charts in the bottom row of Fig. 4 shows the qualitative predictions for both the normative model (blue lines) and the mutation sampler (red lines) for the same conditional probability judgments as in Fig. 3.

Note that in all conditions the predictions of the normative model consist of horizontal lines, as required by the Markov condition. In contrast, those of the mutation sampler (red lines) vary with condition. When both causal relations are generative (Fig. 4A), the prototype states are 000 and 111. As already demonstrated in Fig. 3, the consequence of those prototype states is that the mutation sampler yields Markov violations such that E_A is more likely when E_B is present versus absent even conditioning on C .

However, a different pattern emerges in Fig. 4B when $C \rightarrow E_B$ is inhibitory. Recall that the mutation sampler defines prototype states as those that are guaranteed to be qualitatively consistent with the causal relations. Because (a) the inhibitory relation between C and E_B suggest that those two variables should have opposite values (i.e., present and absent or vice versa) and (b) the generative relation between C and E_A suggest they should have the same values, the prototypes for this network are 001 and 110. The lower panel of Fig. 4B reveals that the mutation sampler continues to predict Markov violations but, critically, the direction of those violations has reversed: conditioned on C , E_A is now *less* likely when E_B is present versus absent. This result obtains because the two prototypes have introduced a *negative* correlation between E_A and E_B conditioned on C .

In Fig. 4C it is now $C \rightarrow E_A$ that is inhibitory and so the prototype states are 010 and 101. As in Fig. 4B, the negative correlation between E_A and E_B introduced by these prototypes yields Markov violations with the reverse direction compared to Fig. 4A. Note that because $C \rightarrow E_A$ is now inhibitory, E_A is now less likely to be present when C is (cf. Fig. 4A and Fig. 4B).

Finally, in Fig. 4D both causal relations are inhibitory and so the prototypes are 011 and 100. Because the correlation between E_A and E_B introduced by these prototypes is once again positive, the direction of the Markov violations is as it was in Fig. 4A.




Figure 4: Qualitative predictions of the normative model (blue lines, solid plot points) and the mutation sampler (red lines, open plot points) for four experimental conditions. The top row shows the condition's causal network and corresponding prototype states. (A) Both causal links are generative. (B) $C \rightarrow E_A$ is generative whereas $C \rightarrow E_B$ is inhibitory. (C) $C \rightarrow E_A$ is inhibitory whereas $C \rightarrow E_B$ is generative. (D) Both causal links are inhibitory. The weights on the generative and inhibitory relations were 1 and -1, respectively, and $w_{E_A} = w_{E_B} = 0$ (see Footnote 2).

These novel predictions are now tested in an experiment that manipulates whether each causal relation in a common cause structure was generative or inhibitory.

Method

Materials. Three domains were tested: economics, meteorology, and sociology. Subjects were first told that the domain they were about to study included three binary variables. In the domain of economics the variables were interest rates (either low or normal), trade deficits (small or normal), and retirement savings (high or normal). In the domain of meteorology, the variables were ozone level, air pressure, and humidity. In sociology they were degree of urbanization, interest in religion, and socio-economic mobility.

Subjects were then presented with a verbal description of two causal relations that formed a common cause network. Depending on condition the causal relationships were either generative or inhibitory and included a description of the mechanism responsible for that relationship. For example, one of the economics generative relationships was “Low interest rates cause small trade deficits. The low cost of borrowing money leads businesses to invest in the latest manufacturing technologies, and the resulting low-cost products are exported around the world.” One of the inhibitory relationships was “Low interest rates prevents high retirement savings. The good economic times produced by the low interest rates leads to greater confidence and less worry about the future, so people are less concerned about saving for retirement.”

Procedure. Subjects first studied several screens of information about the domain, including the cover story, a description of the three variables, the two causal links, and a diagram of those links. Participants then took a multiple-choice test of this knowledge. While taking the test, participants could return to the information screens they had studied but doing so obligated them to retake the questions they missed.

Subjects were then presented with the inferences test. Each trial presented the values of one or two variables and asked the subject to predict the state of another. For example, a subject might be told that an economy has low interest rates and a normal trade deficit and be asked the probability of it having a high level of retirement savings. Subjects entered their response by moving a tick on a rating scale whose ends were labeled 0% and 100%. To ensure that subjects did not have to rely on their memory, the causal relationships were repeated on the bottom half of the screen. A total of 24 inference trials were presented, including those shown in Fig. 4 relevant to assessing Markov violations. The order of the 24 test questions was randomized for each participant.

Design and Participants. The experiment consisted of a 3 (domain: economics, meteorology, or sociology) $\times$ 4 (causal network) between-subject design. Subjects were randomly assigned to these $3 \times 4 = 12$ between-subject cells subject to the constraint that an equal number appeared in each cell. 72 New York University undergraduates received course credit for participating. These sample sizes are similar to those in Rehder and Waldmann (2017).




Figure 5: Empirical conditional probability ratings in four experimental conditions (gray bars). The blue and red lines are the fits of the normative model and the mutation sampler, respectively. Error bars are 95% confidence intervals.

Results

Initial analyses revealed no effect of domain and so the key inference ratings in each condition are presented in Fig. 5 collapsed over that factor.⁵

In all conditions ratings varied in the expected manner with the state of variable C : When $C \rightarrow E_A$ was generative (conditions A and B), E_A was judged more likely when C was present versus absent whereas when $C \rightarrow E_A$ was inhibitory (conditions C and D), E_A was judged more likely when C was absent versus present. The key result concerned the direction of the Markov violations in each condition. In condition A, E_A was judged more likely when E_B was present versus unknown versus absent, replicating previous research (e.g., Fig. 3). But when one of the two causal links was inhibitory (conditions B and C), the direction of the Markov violation reversed: E_A was generally judged *less* likely when E_B was present versus unknown versus absent. Finally, when both causal relations were inhibitory (condition D), the sign of the Markov violation was in the same direction as in condition A (albeit the magnitude of the violation appears smaller, especially when C was present). This overall pattern of results—with an apparent positive correlation between effects when the causal relations were the same sign and a negative correlation for one inhibitory and one generative link—was of course the one predicted in Fig. 4.

⁵Because of the symmetry mentioned in Footnote 4, data in panels A and D have been averaged over both E_A and E_B inferences. Analogously, E_A inferences in condition B are logically equivalent to E_B inferences in condition C and so those results have been combined and presented in panel B. E_A inferences in condition C are logically equivalent to E_B inferences in condition B and so those results are presented in panel C.

These effects were assessed statistically with a model in which C and the linear and quadratic effects of E_B , plus their interactions, served as predictors. The results are presented in Table 1. Unsurprisingly, in every analysis there was a large effect of C , indicating that the judged probability of E_A varied depending on the state of C . More importantly, the linear effect of E_B was also significant in every condition, confirming the presence of Markov violations.

Table 1 shows that another effect obtained in all conditions was the interaction between the C and the quadratic effect of E_B . For example, in condition A when C was absent, the rating for $p(E_A = 1|C = 0)$ (28.1) was slightly lower than the average of $p(E_A = 1|C = 0, E_B = 0)$ and $p(E_A = 1|C = 0, E_B = 1)$ (17.5 and 38.8, respectively). But when C was present, the rating for $p(E_A = 1|C = 1)$ (85.0) was slightly *higher* than the average of $p(E_A = 1|C = 1, E_B = 0)$ and $p(E_A = 1|C = 1, E_B = 1)$ (72.5 and 90.0, respectively). That is, the sign of the quadratic effect of E_B varied with the state of C . This pattern obtained in all four conditions. Note that this quadratic interaction is also apparent in the predictions of the the mutation sampler shown in Fig. 4.

Theoretical Modeling

Both the normative model and the mutation sampler were fit to each subject's 24 conditional probability judgments. The results of these fits averaged over subjects are shown in Fig. 5 superimposed on the empirical data. As expected, the fits of the normative model (horizontal blue lines) reflect the absence of Markov violations. In contrast, the fits of the mutation sampler (red lines) reproduce the sign of the Markov violations exhibited by subjects: positive in Conditions A and D

Table 1: Inferential statistics testing the effect of the states of C and E_B on the inference of E_A .

Effect	Condition A			Condition B			Condition C			Condition D		
	F	p	BF	F	p	BF	F	p	BF	F	p	BF
C	200	$< 10^{-10}$	$> 10^8$	303	$< 10^{-18}$	$> 10^{15}$	246	$< 10^{-16}$	$> 10^{14}$	136	$< 10^{-8}$	$> 10^6$
E_B												
Linear	59.2	$< 10^{-6}$	$> 10^4$	46.4	$< 10^{-17}$	$> 10^5$	20.6	$< 10^{-4}$	373	11.2	.003	11.8
Quadratic	0.14	.710	.259	21.8	$< 10^{-4}$	526	0.81	.374	.261	1.65	.216	.494
$C \times E_B$												
Linear	7.17	.015	3.61	5.53	.024	2.00	1.98	.168	.442	3.81	.068	1.14
Quadratic	8.67	.009	5.74	11.7	.002	20.8	6.83	.013	3.37	9.84	.006	8.08

and negative in Conditions B and C. And, the mutation sampler reproduces the interaction between C and the quadratic effect of E_B . These effects resulted in the mutation sampler yielding a better fit ($AICs$ of 103.1, 121.2, 120.9, and 118.6 in conditions A, B, C, and D, respectively) as compared to the normative model (120.3, 127.2, 127.0, and 123.5). In the mutation sampler fits the median value of the chain length parameter across the four conditions was 7.9.

Discussion

This article tested predictions regarding the presence of new causal reasoning errors. Previous demonstrations of Markov violations have all exhibited a characteristic direction: in a common cause network, the presence of an effect is judged more likely when the other effect is also present even when the state of the common cause was known. This article’s empirical contribution is to demonstrate that this result does not generalize to all common cause graphs. When one of the two causal relations is inhibitory, the direction of the Markov violations reverses such that an effect is judged to be *less* likely when the other effect is present.

These findings are inconsistent with some past characterizations of Markov violations. For example, Rottman and Hastie (2016) proposed that causal inferences are guided by a *monotonicity principle* in which the strength of a causal inferences is a function of the number of graph variables present minus the number of variables absent. Relatedly, Rehder and Waldmann (2017) proposed a *rich-get-richer principle* according to which reasoners assume that one variable is more likely to be present to the extent that other variables in the causal model are also present. The current results indicate that neither of these principles generalize to causal graphs with inhibitory links.

The reversal of the direction of the Markov violations was predicted by the mutation sampler. As a rational process model the mutation sampler explains why people draw approximately correct causal inferences but also why they make systematic errors. In particular, starting MCMC sampling at one of a graph’s prototype states combined with short sampling chains yields a distorted joint distribution and, in turn, causal inferences that violate the Markov condition. We demonstrated how the prototype states for causal graphs with

a mix of generative and inhibitory causal relations are sufficient to reverse the usual direction of Markov violations (i.e., rather than richer, the rich got poorer).

We also showed that the prototype states for a graph with nothing but inhibitory links explained the traditional direction of the Markov violation observed in that condition (Fig. 5D). Nevertheless, note that the magnitude of those violations was somewhat smaller than in the other conditions. Interestingly, we also note that the effect of C was smallest in Condition D. One possible explanation for these results is that the verbal descriptions of the inhibitory relations led them to be viewed as the less efficacious as compared to the generative relations. It also possible that people generally reason less confidently with inhibitory relations.

The mutation sampler shares some similarities with mental models theory (MMT) (Johnson-Laird, 1980), which also posits that the fundamental units of reasoning are concrete possibilities and gives special status to certain states (referred to as *initial mental models* in MMT). Recently, Johnson-Laird et al. (2015) defined a sampling process in which mental models are stochastically generated, which then form the basis for a causal inference. Nevertheless the mutation sampler and MMT posit different initial states in some circumstances (Davis & Rehder, 2020). More fundamentally, the mutation sampler but not MMT is an example of a rational process model that approximates the normative standard to the extent that sufficient cognitive resources are available.

Markov violations have been a key data point in assessing the suitability of the causal graphical model framework for a psychological theory of causal cognition. The current results indicate that those violations are both more prevalent than previously thought and vary systematically with the relative mix of generative and inhibitory causal relations. The mutation sampler, a model designed to account for Markov violations with purely generative causal relations, in fact correctly anticipated the novel direction of those violations when inhibitory relations are present. The mutation sampler makes novel predictions for other causal network topologies that will be tested in future work, in addition to other predictions that can be derived from a process level account, such as response times and the within-person variability of causal judgments.

References

- Ali, N., Chater, N., & Oaksford, M. (2011). The mental representation of causal conditional reasoning: Mental models or causal models. *Cognition*, 119(3), 403–418.
- Davis, Z., & Rehder, B. (2020). A process model of causal reasoning. *Cognitive Science*, 44.
- Fernbach, P. M., & Rehder, B. (2013). Cognitive shortcuts in causal inference. *Argument & Computation*, 4(1), 64–88.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in cognitive science*, 7(2), 217–229.
- Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1), 97–109.
- Johnson-Laird, P. N. (1980). Mental models in cognitive science. *Cognitive science*, 4(1), 71–115.
- Johnson-Laird, P. N., Khemlani, S. S., & Goodwin, G. P. (2015). Logic, probability, and human reasoning. *Trends in cognitive sciences*, 19(4), 201–214.
- Mayrhofer, R., & Waldmann, M. R. (2015). Agents and causes: Dispositional intuitions as a guide to causal structure. *Cognitive science*, 39 1, 65–95.
- Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the markov property in causal reasoning. *Cognitive Psychology*, 67(4), 186–216.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive psychology*, 72, 54–107.
- Rehder, B. (2017a). Concepts as causal models: Classification. In M. R. Waldmann (Ed.), *The oxford handbook of causal reasoning (edited volume)*. Oxford University Press.
- Rehder, B. (2017b). Concepts as causal models: Induction. In M. R. Waldmann (Ed.), *The oxford handbook of causal reasoning (edited volume)*. Oxford University Press.
- Rehder, B. (2018). Beyond markov: Accounting for independence violations in causal reasoning. *Cognitive psychology*, 103, 42–84.
- Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, 50(3), 264–314.
- Rehder, B., & Waldmann, M. R. (2017). Failures of explaining away and screening off in described versus experienced causal learning scenarios. *Memory & cognition*, 45(2), 245–260.
- Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological bulletin*, 140(1), 109.
- Rottman, B. M., & Hastie, R. (2016). Do people reason rationally about causally related events? markov violations, weak inferences, and failures of explaining away. *Cognitive Psychology*, 87, 88–134.
- Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual Review of Psychology*, 66(3), 485.
- Van Ravenzwaaij, D., Cassey, P., & Brown, S. D. (2018). A simple introduction to markov chain monte-carlo sampling. *Psychonomic bulletin & review*, 25(1), 143–154.
- Waldmann, M. R. (2017). *The oxford handbook of causal reasoning*. Oxford University Press.