

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Statistical and Mechanistic Information in Evaluating Causal Claims

Permalink

<https://escholarship.org/uc/item/2vp4v6ps>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Johnsol, Samuel G. B.

Keil, Frank C.

Publication Date

2017

Peer reviewed

Statistical and Mechanistic Information in Evaluating Causal Claims

Samuel G. B. Johnson & Frank C. Keil
(sgbjohnson@gmail.com, frank.keil@yale.edu)

Department of Psychology, Yale University, New Haven, CT 06520 USA

Abstract

People use a variety of strategies for evaluating causal claims, including mechanistic strategies (seeking a step-by-step explanation for how a cause would bring about its effect) and statistical strategies (examining patterns of co-occurrence). Two studies examine factors leading one or the other of these strategies to predominate. First, general causal claims (e.g., “Smoking causes cancer”) are evaluated predominantly using statistical evidence, whereas statistics is less preferred for specific claims (e.g., “Smoking caused Jack’s cancer”). Second, social and biological causal claims are evaluated primarily through statistical evidence, whereas statistical evidence is deemed less relevant for evaluating physical causal claims. We argue for a pluralistic view of causal learning on which a multiplicity of causal concepts lead to distinct strategies for learning about causation.

Keywords: Causal reasoning; concepts and categories; information evaluation; statistical reasoning.

Introduction

Causal knowledge is crucial for understanding and controlling the world, and strategies for evaluating causal claims are central to gatekeeping that crucial knowledge. Humans seem especially prone to two strategies—a *mechanism* strategy, on which we consider potential mediating causal links as evidence favoring a causal connection; and a *statistical* strategy, on which we look for correlations between a cause and effect. For example, Jack is assessing the risk that smoking causes cancer. He can assess this claim mechanistically by considering the plausibility of potential mediating mechanisms that explain relationship between smoking and cancer. Or he can assess the claim statistically by observing whether the frequency of cancer is higher in a population that smokes compared to a population that does not.

There is ample evidence that people use both of these strategies, though different theoretical approaches to causal cognition emphasize different types of information.

According to mechanism-based approaches to causal cognition, we learn about causal relations primarily by searching for generative mechanisms through which causes can produce their effects. Several convergent lines of evidence are consistent with causal relations being represented in terms of underlying mechanisms (see Johnson & Ahn, *in press*). Knowledge of underlying mechanisms affects whether discounting or conjunction effects occur in causal attribution (Ahn & Bailenson, 1996), whether the Markov principle is applied to causal networks (Park & Sloman, 2013), and whether causal chains are judged to be transitive (Johnson & Ahn, 2015).

If we learn about causation by searching for plausible

mechanisms, then people would seek out evidence of underlying mechanisms when determining whether one thing causes another. Indeed, people do sometimes assess causal hypotheses by forming a mechanistic narrative that would lead from X to Y and assessing the plausibility of that narrative (e.g., Fernbach, Darlow, & Sloman, 2011; Kahneman & Tversky, 1982; Taleb, 2007). For example, Jack might imagine some physiological mechanism by which smoking and lung cancer could be connected, then evaluate the plausibility of these steps. Moreover, people prefer mechanism evidence overwhelmingly in causal *attribution*—that is, in determining which cause to assign to an effect (Ahn, Kalish, Medin, & Gelman, 1995).

In contrast, statistics-based approaches to causal learning emphasize the role of statistical knowledge in inferring causal relationships. These theories hold that causal relationships are primarily discovered through information about the co-occurrence of the cause and effect, although individual theories differ in the details of how these inferences work (e.g., Cheng, 1997; Gopnik et al., 2004; Griffiths & Tenenbaum, 2005). These theories do not necessarily claim that causal relations are *represented* in terms of statistical patterns, but often hold that causal relations are represented in terms of abstract causal powers underlying the connection between cause and effect, which are then *inferred* through statistical means (Cheng, 1997; Pearl, 2000). Nonetheless, statistical approaches do claim that causal relations are primarily *learned* through co-occurrence information, and there is abundant evidence that people are often able to learn from statistical evidence (e.g., Gopnik et al., 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003).

Moreover, statistical evidence could be an antidote to the shallowness of people’s knowledge of causal mechanisms: Even though people do *use* mechanism knowledge in evaluating causal relationships when it is available, we do not seem to have extensive knowledge of mechanisms. People greatly overestimate their knowledge of how everyday devices such as flush toilets work, revealing misconceptions and pervasive gaps in understanding (Rozenblit & Keil, 2002). People’s beliefs in mechanisms underlying causal relationships are more likely to take the form of generic or highly unspecified ‘placeholders’, akin to our beliefs in abstract category essences (Medin & Ortony, 1989). Such skeletal representations are difficult to square with a strong mechanism view on which people seek a detailed understanding of how causal relationships work and use that understanding to guide inference, but they might seem to be more consistent with statistical approaches to causal thinking on which covariation is used to infer the

existence of abstract underlying mechanisms without being committed to particular mechanistic details.

The mechanism and the statistical approaches, however, need not be in conflict and can be mutually compatible with a third approach known as *causal pluralism* (Cartwright, 2004; Danks, 2005; Hitchcock, 2003; Lombrozo, 2010; Waldmann & Mayrhofer, 2016). According to this approach, people might use a multiplicity of causal concepts and a concordant variety of learning strategies in systematic, context-dependent ways. Some *prima facie* support for the pluralistic position comes from experiments where people used mechanism and statistical evidence in an interactive manner (Fugelsang & Thompson, 2000; Spellman, 1996).

Yet, little is known about contextual factors that lead each type of evidence to predominate. Here, we look at two dimensions along which causal relations can vary—in its level of abstraction and its domain. Because people seem to use different sorts of causal *concepts* for representing these relations, we anticipated that people may also use different strategies to learn about these relations. If you need to decide whether something is a banana, the best question to ask would be about its shape, whereas if you need to decide whether something is a peach, the best question would be about its texture. And just as we must consult our concept of ‘banana’ when deciding whether something is a banana and our concept of ‘peach’ when deciding whether something is a peach, we must consult our concept of ‘cause’ when deciding whether a relationship is causal. When we deploy different causal concepts across contexts, this can lead to different learning strategies.

General and Specific Causation. *General* causal claims refer to generic causal patterns (“Smoking causes a person to get lung cancer”), whereas *specific* claims refer to concrete occasions when a pattern was instantiated (“Smoking caused Jack to get lung cancer”). The inferences supported by general and particular claims differ in several ways. General claims are associated with more essentialist inferences (Cimpian & Erickson, 2012) and, in the domain of human behavior, with more neuroscientific rather than psychosocial explanations (Kim, Ahn, Johnson, & Knobe, 2016). Might these claims also differ in the evidence used for their evaluation?

General claims refer to an entire category of causal relationships (i.e., a set of event pairs), whereas specific claims refer to an instance of that category (one single event pair in that set). Thus, general claims necessarily *quantify* over multiple instances and intrinsically carry statistical content, whereas specific claims do not. We suggest that this conceptual difference could lead statistical evidence to be privileged more for general rather than specific causal claims.

This pattern of evidence preferences can lead to non-normative behavior. If we are not privy to the particulars of Jack’s case, the *only* strategy for evaluating a token-causal claim will be to look for a more general causal

pattern between smoking and cancer—to evaluate the general claim. Thus, evidence relevant for evaluating the general claim would be equally relevant for evaluating the specific claim. Imagine that a tobacco company is being sued under one of two different circumstances: (1) a class action suit (the plaintiffs’ lawyers arguing that “Smoking causes a person to get lung cancer”), or (2) Jack’s single party action (his lawyers arguing that “Smoking caused Jack to get lung cancer”). In both cases, jurors might be confronted with mechanism evidence, such as a biologist’s testimony concerning biochemical mechanisms, or with statistical evidence, such as an epidemiologist’s testimony comparing cancer rates across populations. It seems difficult to justify a difference between these two cases in jurors’ relative weighing of mechanistic and statistical testimony. Yet, if people rely on different processes for evaluating general and specific claims, then the jurors may well behave differently.

Causation across Domains. People use different intuitive theories of causality across domains. Whereas physical causation is typically conceptualized in terms of force propagating down branching causal chains, social and biological causation are thought of as webs of interconnected influences. People tend to identify physical events as having one cause but many effects, whereas social events are seen as having many causes and many effects (Strickland, Silver, & Keil, 2016; see also Johnson, Valenti, & Keil, 2017). Likewise, even young children seem to view biological systems as causally interacting parts in homeostatic balance (Keil, 1989). Thus, the simple, linear causal pathways thought to be at play in the physical world give way to more complex causal structures in the social and biological domains.

Similarly, social (e.g., psychological or economic) causation is often goal-directed (Lombrozo & Carey, 2006) or equipotential (Heider, 1958)—the same ends can often be brought about through many different means. For this reason, people’s causal theories of social (and likely biological) systems often focus on counterfactual dependence (Lombrozo, 2010), whereas their theories of physical systems are characterized more by ideas about physical force and transference of conserved quantities.

Given the relatively linear and force-based conceptions of physical causation, and the relatively web-like and dependence-based conceptions of causality in biological and social causation, people may use a more deterministic concept of physical causation and a more stochastic concept of social and biological causation (Johnson, Valenti, & Keil, 2017). Thus, people may rely more on mechanistic strategies when learning about physical systems and more on statistical strategies when learning about biological and social systems.

Overview of Studies. Two studies test differences in evidence-seeking between general and specific causal claims, with the studies differing in the framing of the claims. After testing this hypothesis about general versus specific causation, we present an analysis of evidence-

seeking preferences across domains, aggregating across studies. In the General Discussion, we assess the prospects for a pluralistic view of causal learning.

Experiment 1

In Experiment 1, we tested what kind of information people thought most relevant for assessing general causal claims (e.g., “Eating polar bear liver causes a person to become dizzy”) and specific causal claims (“Eating polar bear liver caused Bill to become dizzy”). The mechanism view holds that we learn about causal relationships primarily by searching for underlying mechanisms, leading to a preference for mechanism evidence, whereas the statistical view holds that we learn about causal relationships primarily through contingency information, leading to a preference for statistical evidence. In contrast to both positions, we predicted that, whatever people’s baseline preferences for one or the other type of evidence, the preference for statistical evidence would be stronger when evaluating general rather than specific claims.

Method

We recruited 80 participants from Mechanical Turk, and excluded 5 from data analysis because they incorrectly answered more than 33% of the check questions.

Participants saw either the general or specific version of each of 24 causal claims, presented in a box. For each item, participants were asked “Which of the following types of evidence would be most helpful to you in determining whether the statement in the box is true?” as a forced-choice. For the polar bear item, the options read:

Statistical: “Measurements of the frequency of dizziness of many people after they eat or do not eat polar bear liver.”

Mechanism: “An explanation of why eating polar bear liver would cause a person to become dizzy.”

Anecdotal: “Knowing whether there is another occasion on which a person ate polar bear liver and then they felt dizzy.”

We assumed that few people would choose the weak anecdotal evidence, and used this option to assess the degree to which participants used poor causal reasoning. The order of the options was randomized for each item, and the items were presented in a random order.

Results and Discussion

As shown in Table 1, statistical evidence was chosen more frequently when evaluating general compared to specific claims. Due to non-normality, Mann-Whitney U -tests were used to compare the number of items for which participants chose each evidence type in each condition.

These tests showed that statistical evidence was chosen for more items when evaluating general claims than when evaluating specific claims [$U = 496.5$, $p = .028$, $r = .25$]. This corresponded to relatively fewer mechanism responses for the general claims than for the specific claims and fewer anecdotal responses for the general

claims than for the specific claims. Thus, responses shifted relatively more toward statistical evidence for the general than for the specific claims.

This result indicates that people use pluralistic causal learning strategies. Specifically, it appears that the conceptual differences between general and specific claims had downstream consequences for evidence-seeking preferences: Because general claims quantify over instances, statistical evidence is seen as more relevant to evaluating such claims, compared to specific claims, and mechanism evidence is seen as less relevant.

Table 1: Results of Experiments 1 and 2

	Statistical	Mechanism	Anecdotal
Exp. 1			
General	55.3%	36.8%	7.9%
Specific	41.6%	47.5%	11.0%
Exp. 2			
General	62.0%	30.4%	7.6%
Specific	47.6%	41.4%	11.0%
Domain Analysis			
Physical	47.5%	46.6%	5.9%
Biological	53.6%	38.9%	7.5%
Psychological	51.6%	36.0%	12.4%
Economic	53.7%	34.7%	11.6%

Note. Entries indicate the proportion of choices of each evidence type in each experiment. For the domain analysis, the proportion of participants choosing each evidence type was calculated for each item in Experiments 1 and 2, and those proportions were averaged across all items in each domain.

Experiment 2

Experiment 2 sought to generalize the effect of general versus specific causation to contexts where it is known that the events in the specific causal relationship actually occurred. That is, participants in Experiment 1 evaluated claims such as “Smoking cigarettes caused Jack to get lung cancer” without knowing whether or not Jack in fact smoked and whether or not he had cancer. In such contexts, both statistical and mechanism information may seem irrelevant, since a crucial part of evaluating this claim is establishing first that the cause and effect both occurred. In contrast, Experiment 2 examined contexts where it is known that both cause and effect occurred (e.g., by prefacing the causal claim with the statement “Jack smoked cigarettes, and then Jack got lung cancer”), where the primary concern is distinguishing causation from coincidence (see Cartwright, 2017) and where the available evidence would be seen as more relevant.

Method

We recruited 80 participants from Mechanical Turk, and excluded 5 from data analysis because they incorrectly

answered more than 33% of the check questions.

Participants responded to a new set of 24 causal claims. The format of these items differed from Experiment 1 in that contextual information was given for each claim, establishing that the cause and effect occurred. This information was printed above the box containing the claim. For example, one general item read (background information in regular typeface, claim in italics):

Researchers sometimes observe that a person consumes large amounts of meat, and then that the person develops kidney stones.

Consuming large amounts of meat causes a person to develop kidney stones.

The specific version of that item read:

Researchers observed that Tom consumed large amounts of meat, and then that Tom developed kidney stones.

Consuming large amounts of meat caused Tom to develop kidney stones.

The procedure was otherwise identical to Experiment 1.

Results and Discussion

Although participants preferred statistical information overall, this preference was far stronger when evaluating general than when evaluating specific claims [$U = 481.5$, $p = .019$, $r = .27$], consistent with Experiment 1. They correspondingly chose mechanism evidence less frequently for general than for specific claims and anecdotal evidence less frequently for general than for specific claims, as shown in Table 1.

These two experiments together are consistent with the idea that people use different learning strategies depending on what causal concept they are consulting. However, there are other differences between general and specific causation that could plausibly account for some of the variance. First, the reference class from which the statistical evidence is drawn may be more relevant for the general than the specific claim, and second, plurality may have been more salient for the general than for the specific claims. We conducted an additional experiment with artificial stimuli to rule out these two alternative explanations, in which both the general and specific claims were prefaced by a statement about the reference class (e.g., “There is a group of 100 Garbotrons”), with the general claim then made about the entire group and the specific claim about an arbitrary member of that group. This equated the reference class and the salience of plurality, yet produced a similar shift across conditions.

These experiments do not fully tease apart whether the difference is due to a statistics preference for general claims or a mechanism preference for specific claims. We conducted two additional studies to answer this question, one in which participants answered an open-ended question about what evidence they would want to use, and another in which participants rated the two types of evidence on independent scales. Consistent with our claim that these differences arise due to more stochastic

representations of general causation, the condition differences were significant for statistical evidence but not for mechanism evidence in both cases.

Domain Differences

In Experiments 1 and 2, we drew our causal claims from four domains—physical, biological, psychological, and economic—across which causal representations are likely to differ. People typically conceptualize physical causation as flowing in branches, with each event having few causes but many effects, and social (and perhaps biological) causal systems as interconnected webs, in which events have many causes and many effects (Strickland et al., 2016). Similarly, people may use more *transference-based* (or mechanistic) causal concepts in the physical domain, and more *dependence-based* (counterfactual or statistical) causal concepts in the social domain (Lombrozo, 2010). Thus, physical systems may be seen as more deterministic and social systems as more stochastic. According to the pluralistic position, these conceptual differences across domains could translate into different learning strategies: We would expect relatively greater reliance on statistical information for social and biological systems and less for physical systems.

We tested this possibility by comparing preferences for statistical evidence across all 48 items used in Experiments 1 and 2, collapsing across the general and specific versions. For each item, a *statistics preference* score was computed by taking the difference between the proportion of participants choosing statistical evidence for that item and the proportion choosing mechanism evidence for that item. An ANOVA on these scores with domain (physical, biological, psychological, or economic) as a between-items variable uncovered a marginally significant main effect of domain [$F(3,44) = 2.22$, $p = .099$, $\eta_p^2 = .13$], with the preference for statistics evidence smallest for the physical items [$M = 0.01$, $SD = 0.17$], followed by the biological [$M = 0.15$, $SD = 0.19$], psychological [$M = 0.16$, $SD = 0.21$], and economic [$M = 0.19$, $SD = 0.17$] items. Independent-samples *t*-tests revealed that items from the physical domain had a smaller statistics preference than did items from the combined other domains [$t(46) = -2.56$, $p = .014$, $d = 0.85$], while the biological, psychological, and economic domains did not differ from one another [$ts < 1$, $ps > .50$].

This result further supports the pluralistic position, suggesting that differences in causal concepts used across domains translated into different learning strategies.

General Discussion

Cognition requires us to attend to and integrate various sources of information into coherent representations of the world. Our representations of causal systems are particularly critical because they allow us to predict and understand events, and to plan interventions on the world to achieve goals. Humans use two distinct strategies for making inferences about causal claims—evaluating the

plausibility of mediating causal mechanisms, and evaluating statistical evidence for contingencies between cause and effect. What factors lead people to favor one strategy over the other?

First, general causal statements, which refer to a category of causal events, are seen as more compatible with statistical evidence than are specific causal statements, which refer to only an individual causal event. We hypothesized that this would occur because representations of general claims intrinsically include statistical content, and people would seek evidence that conforms to their representation of the causal concept.

Second, statistics were seen as more relevant for biological and social systems than for physical systems, whereas mechanistic evidence was more important for physical systems. We predicted this effect because causal representations vary across domains. Whereas physical systems are seen as more linear and force-based, social and biological systems are seen as more branching and counterfactual-based (Lombrozo, 2010; Strickland et al., 2016). Thus, concepts of biological and social causation would be more stochastic than concepts of physical causation, leading people to favor statistical evidence.

Causal Pluralism. Our causal representations subserve a variety of cognitive functions, and exhibit a concordant variety of properties that sometimes appear contradictory (Johnson & Ahn, *in press*). For instance, causal representations seem to have many of the properties of associations (Shanks, 1987), yet causal inferences exhibit directional biases that are inconsistent with symmetric associative representations (Waldmann & Holyoak, 1992). These shortcomings of associative theories have led to the suggestion of causal models or Bayesian networks as the representation over which causal reasoning operates (e.g., Pearl, 2000; Sloman, 2005). Yet, other evidence suggests that people often fail to make the transitive inferences predicted by Bayesian networks (i.e., that *A* causes *C*, given that *A* causes *B* and *B* causes *C*), and that these failures occur when the connection between *A* and *C* is not seen as a coherent, schematized mechanism (e.g., sex causes pregnancy, which causes nausea, but sex does not cause nausea; Johnson & Ahn, 2015). Thus, causal representations appear to have some association-like properties, some network-like properties, and some schema-like properties. Add to this evidence that causal relations are represented with some properties of forces (Wolff, 2007), icons that support mental simulation (Hegarty, 2004), and metacognitive placeholders (Rozenblit & Keil, 2002), and it becomes clear that people do not represent causation using one unified representation (see Markman & Dietrich, 2000).

Despite the overwhelming evidence for representational pluralism, it does not follow that people use distinct strategies for learning about different varieties of causal concepts. People may not tailor their learning strategies to the representation at hand, but could instead apply a single learning strategy across all types of causal systems,

such as statistical learning algorithms (Pearl, 2000).

However, the current experiments demonstrate learning patterns that are not only pluralistic, but appear to be tailored to the underlying representation. In the cases of specific causation we used, there is no prior knowledge, so the only option is to learn about the general causal relation anew. If the best strategy for learning about the general claim is statistics, then the best strategy for learning about the specific relation is also statistics. Yet, participants shifted dramatically from statistics when learning about specific claims—a signal that they had applied a heuristic, matching statistical representations of general claims to statistical information. Therefore, any view of causal learning and representation that focuses on a single representation or learning mechanism will fail to capture important aspects of our causal thinking.

In addition to clarifying the debate between mechanism and statistical views of causation, causal pluralism may also be a helpful framework for thinking about debates over causal semantics. Theories of causal semantics embrace diverse accounts based on physical forces (Wolff, 2007), on probability (Good, 1961), and on logic (Lewis, 1973). Teasing these accounts apart has been difficult because they often make similar empirical predictions (Barbey & Wolff, 2007; Goldvarg & Johnson-Laird, 2001; Sloman, Barbey, & Hotaling, 2009).

However, in a pluralistic framework, it may not only be difficult but in fact *impossible* to capture all of causal semantics using a single representational format. Our causal representations differ not only in reference (general or specific) and domain (physical, biological, social), but along many other dimensions as well, in potentially interconnected ways—among deterministic, chaotic, and indeterministic systems; among the past, present, and future; between observed, unobserved, and unobservable causes and effects; between categorically or continuously valued causes and effects; and among various potential causal structures. A useful strategy going forward may be to investigate the manner in which such variation in causal meaning propagates to causal learning processes.

References

- Ahn, W., & Bailenson, J. (1996). Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *Cognitive Psychology, 31*, 82–123.
- Ahn, W., Kalish, C.W., Medin, D.L., & Gelman, S.A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition, 54*, 299–352.
- Barbey, A.K., & Wolff, P. (2007). Learning causal structure from reasoning. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science, 71*, 805–20.
- Cartwright, N. (2017). Single case causes: What is

- evidence and why. In *Philosophy of Science in Practice*. Cham, Switzerland: Springer International.
- Cheng, P.W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Cimpian, A., & Erickson, L.C. (2012). The effect of generic statements on children's causal attributions: Questions of mechanism. *Developmental Psychology*, *48*, 159–70.
- Danks, D. (2005). The supposed competition between theories of human causal inference. *Philosophical Psychology*, *18*, 259–72.
- Fernbach, P.M., Darlow, A., & Sloman, S.A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General*, *140*, 168–85.
- Fugelsang, J.A., & Thompson, V.A. (2000). Strategy selection in causal reasoning: When beliefs and covariation collide. *Canadian Journal of Experimental Psychology*, *54*, 15–32.
- Goldvarg, E., & Johnson-Laird, P.N. (2001). Naive causality: A mental model theory of causal meaning and reasoning. *Cognitive Science*, *25*, 565–610.
- Good, I. J. (1961). A causal calculus (I). *The British Journal for the Philosophy of Science*, *11*, 305–318.
- Gopnik, A., Glymour, C., Sobel, D.M., Schulz, L.E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.
- Griffiths, T.L., & Tenenbaum, J.B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–84.
- Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, *8*, 280–285.
- Heider, F. (1958). *The psychology of interpersonal relations*. Hillsdale, NJ: Erlbaum.
- Hitchcock, C. (2003). Of Humean bondage. *The British Journal for the Philosophy of Science*, *54*, 1–25.
- Johnson, S.G.B., & Ahn, W. (2015). Causal networks or causal islands? The representation of mechanisms and the transitivity of causal judgment. *Cognitive Science*, *39*, 1468–503.
- Johnson, S.G.B., & Ahn, W. (in press). Causal mechanisms. In *The Oxford Handbook of Causal Reasoning*. Oxford, UK: Oxford University Press.
- Johnson, S.G.B., Valenti, J.J., & Keil, F.C. (2017). Opponent uses of simplicity and complexity in causal explanation. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.
- Keil, F.C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Kim, N.S., Ahn, W., Johnson, S.G.B., & Knobe, J. (2016). The influence of framing on clinicians' judgments of the biological basis of behaviors. *Journal of Experimental Psychology: Applied*, *22*, 39–47.
- Lewis, D. (1973). Causation. *The Journal of Philosophy*, *70*, 556–67.
- Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, *61*, 303–32.
- Lombrozo, T., & Carey, S. (2006). Function explanation and the function of explanation. *Cognition*, *99*, 167–204.
- Markman, A.B., & Dietrich, E. (2000). In defense of representation. *Cognitive Psychology*, *40*, 138–171.
- Medin, D.L., & Ortony, A. (1989). Psychological essentialism. In *Similarity and analogical reasoning*. Cambridge, UK: Cambridge University Press.
- Park, J., & Sloman, S.A. (2013). Mechanistic beliefs determine adherence to the Markov property in causal reasoning. *Cognitive Psychology*, *67*, 186–216.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
- Rozenblit, L., & Keil, F.C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, *26*, 521–562.
- Shanks, D.R. (1987). Associative accounts of causality judgment. *Psychology of Learning and Motivation*, *21*, 229–261.
- Sloman, S.A. (2005). *Causal models: How people think about the world and its alternatives*. Oxford, UK: Oxford University Press.
- Sloman, S.A., Barbey, A.K., & Hotaling, J.M. (2009). A causal model theory of the meaning of *Cause*, *Enable*, and *Prevent*. *Cognitive Science*, *33*, 21–50.
- Spellman, B.A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, *7*, 337–42.
- Steyvers, M., Tenenbaum, J.B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–89.
- Strickland, B., Silver, I., & Keil, F.C. (2016). The texture of causal construals: Domain-specific biases shape causal inferences from discourse. *Memory & Cognition*.
- Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. New York, NY: Random House.
- Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222–236.
- Waldmann, M. R., & Mayrhofer, R. (2016). Hybrid causal representations. In *Psychology of Learning and Motivation* (Vol. 65). New York, NY: Academic Press.
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, *136*, 82–111.