

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Transferring Novel Causal Knowledge

Permalink

<https://escholarship.org/uc/item/3rx6r915>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Czarnowski, Daniel W

Marsh, Jessecae K.

Publication Date

2022

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

Transferring Novel Causal Knowledge

Daniel W. Czarnowski (dwc319@lehigh.edu)

Department of Psychology, 17 Memorial Drive East
Bethlehem, PA 18015

Jessecae K. Marsh (jessecae.marsh@lehigh.edu)

Department of Psychology, 17 Memorial Drive East
Bethlehem, PA 18015

Abstract

Knowledge of cause and effect allows people to navigate and understand the complex systems of the world. Despite the importance of causal knowledge to everyday reasoning, little is known about how people transfer causal knowledge learned in one situation to novel contexts. In two experiments, we examine when people choose to generalize two types of causal knowledge, causal mechanisms (Experiment 1) and causal strength (Experiment 2), across various domains. We find that people willingly transfer causal knowledge to novel contexts when the entities in those contexts share high categorical relatedness with the source of the causal knowledge. The extent to which people are willing to transfer causal knowledge decreases as category similarity decreases. We discuss future research that could delineate the boundaries of causal transfer.

Keywords: Causal reasoning; category-based induction; analogy

Introduction

Causal knowledge can be used to make predictions, perform interventions, and explain how or why something occurred (Shanks, 2004; Hagmayer & Sloman, 2009; Fernbach et al., 2011). However, people must often navigate novel cause-and-effect relationships despite lacking explicit causal knowledge for that particular situation. For example, a gardener might know that a specific fertilizer causes their rose plants to bloom, but will the same fertilizer cause their tulips to bloom? In such cases, one must draw upon existing, relevant knowledge to reason their way through the novel context. Because the fertilizer has a blooming effect on roses, it may be easily inferred that the fertilizer will have the same effect on tulips, given that roses and tulips are both flowering plants. The simplicity of this inference, however, belies the complexity of how causal knowledge is transferred. Does the fertilizer act on the same mechanism that causes roses and tulips to bloom? Will the fertilizer act on tulips as strongly as it did on roses? In this paper, we investigate people's intuitions of the answers to those questions by investigating when people choose to transfer causal knowledge.

People regularly generalize knowledge and skills (Shepard, 1987) to novel contexts, but the manner in which they do so often depends on the situation they find themselves in. Analogical reasoning, for example, affords knowledge transfer between specific instances (Holyoak & Lee, 2017). Analogies such as the water model for electricity allow one

to understand something novel (i.e., voltage) by conceptualizing it in terms of something already known (i.e., water pressure). This is achieved by identifying components of the two analogues that share common relational structures (i.e., the flow of water through a pipe corresponds with electrical current through a conductor), and using those relations to guide inferences about the novel system (e.g., Structure-Mapping Theory; Gentner, 1983).

People can also transfer knowledge more broadly, such as to or from entire categories. Category-based induction (Osherson et al., 1990) enables one to draw upon their understanding of known categories/exemplars to infer the properties of other categories/exemplars (Hayes & Heit, 2018). For example, knowing that birds possess a particular anatomical or behavioral property allows one to infer that a novel bird will likewise possess that property (e.g., Sloman, 1993). While such inferences can be guided by principles such as similarity (Osherson et al., 1990), diversity (Heit, Hahn, & Feeney, 2004), and typicality (Malt & Smith, 1982), causal relations can also inform category-based inductions by providing causal mechanisms that relate the properties of a category (Murphy & Medin, 1985).

The previous literature on analogy and category-based induction provides a basis for the idea that people could transfer causal knowledge from one novel context to another. Research on causal mechanisms suggests mechanistic knowledge could be a good basis for such transfer. Previous research has demonstrated the importance of mechanism information when evaluating causal relations (e.g., Ahn et al., 1995). Causal mechanisms explain how causes produce effects. People assume the presence of a mechanism that connects the cause to the effect, even if it is not known (Ahn & Kalish, 2000). With respect to causal knowledge transfer, people may use causal mechanisms to understand a source relationship when evaluating whether the information therein can be transferred to the target.

When invoking these causal mechanisms, people may be attempting to fit them into a causal model that represents the relationship in question. Hierarchical Bayesian models have been adopted by some researchers as theoretical models for how people use these causal models to facilitate reasoning, learning, and categorization (e.g., Rehder & Hastie, 2001). Bayesian models of causal induction suggest that people evaluate and update hypotheses about causal relations by

using prior causal knowledge and observed causal relations (Hagmayer & Mayrhofer, 2013; Kemp et al., 2010). Beneath this framework, transferring causal knowledge would reflect one's hypotheses for the plausibility of a novel effect given their prior knowledge of the effect in a related/unrelated source, as well as their understanding of the higher order relations (i.e., causal mechanisms) that permit the effect to occur in the first place. Importantly, Bayesian models of causal induction generally serve to explain the reasoning process computationally, rather than describe it cognitively (Hagmayer & Mayrhofer, 2013; cf. Mayrhofer & Waldmann, 2011). That said, the Bayesian approach suggests both prior category knowledge (i.e., category-based induction) and situational evidence (i.e., analogy) are important components of the inferences that guide causal knowledge transfer.

In addition to the reasoning process informing transfer of causal knowledge, another open question is how broadly people will transfer causal information. One possibility is that people are very miserly in what they are willing to transfer. That is, when learning a causal relationship between a set of events, people may think of this as a very specific, contextually bound relationship (e.g., Heit, 1998; Feeney & Heit, 2011). Taking our gardening example, knowing that a fertilizer works on roses may be a piece of knowledge that is specifically linked to roses and not generalized to other items. In this way, people may not be willing to do the generalization seen in analogy use or in category-based induction.

Restricted willingness to transfer may also vary for reasons unrelated to causal information. The category-based induction literature suggests people should infer properties to be true of similar category members. But, if people do not spontaneously think about the categorical hierarchies an exemplar is in when thinking about causal transfer, they might not consider how exemplars at different levels of categorization are related to one another (Murphy & Brownell, 1985; Markman & Wisniewski, 1997). Consider pink roses and orchids. It is possible that people overlook the fact that these two flowers are similar at the basic level (i.e., flower) and instead simply perceive them as two distinct plants that share few, if any, commonalities. This could discourage transferring knowledge between them. Similarly, if people choose to fixate on differences rather than similarities, transfer might be impaired because people will not identify commonalities on which to base knowledge transfer in the first place (Markman & Gentner, 1996; Miao & Gentner, 2001).

On the other hand, people may be willing to broadly transfer causal properties to novel entities. For example, learning that a fertilizer works on roses may suggest that it could also work on all forms of plants if one possesses the prior belief that a rose is a good representation of plants as a whole (Garcia-Retamero et al., 2009). Given that people are willing to draw analogies between disparate items and that people will infer properties across diverse categories, it is possible that causal transfer is likewise wide-ranging. If people are more promiscuous in how they transfer causal

relationships, then it becomes a question of how far away from a source this transfer will happen. That is, will people transfer to any other target or are there limits to this causal transfer?

In two experiments, we explore how people transfer causal knowledge into novel reasoning scenarios. We tested transfer of a causal mechanism (Experiment 1) and the strength with which that mechanism acts (Experiment 2). We vary the type of targets for transfer by manipulating how related the targets are to the source in a category hierarchy. Through these studies we can gain a better understanding of how experiencing novel causal relationships influences reasoning in other causal situations.

Experiment 1

In Experiment 1, we present a first test of whether people will transfer causal mechanisms to novel targets. Participants were provided with scenarios presenting a cause producing an effect through a described mechanism for a single category exemplar. The extent to which people were willing to transfer the causal mechanism was measured through prompts asking participants whether the causal mechanism in the scenario could produce the same effect in a series of targets that varied in their categorical relation to the exemplar.

Method

Participants Undergraduate students ($N = 55$) enrolled in an introductory psychology course were recruited to participate in exchange for course credit. All participants completed the study remotely on a computer of their choice. Participants were screened to ensure that they had normal or corrected-to-normal vision and were fluent in English. Approval for both studies was obtained from Lehigh University's IRB.

Materials Brief scenarios were constructed that covered the domains of natural phenomena (e.g., gardening, pesticides), personal health (e.g., pain relief, food allergies), and mechanical devices (e.g., internet connectivity, fuel consumption). In total, 12 scenarios were constructed, four for each of the three domains. Each scenario described a cause that brought about an effect on a category exemplar (the transfer source) through a particular causal mechanism. Nine transfer targets were generated for each transfer source. An example scenario and transfer prompt from the natural domain follows, with the italicized item representing one of the possible transfer targets:

Scenario: Chuck was weeding his garden one day when he realized that many of his pink roses were wilting. After some research, he learned that adding used coffee grounds to the soil might help pink roses firm up again by increasing the nitrogen content of the soil, which is associated with healthy pink roses. So, he administered the used coffee grounds to the soil in his garden. The next morning when he woke up, he noticed that the treatment was effective in making the pink roses that were wilting stand straight and healthy.

Transfer Prompt: Chuck was weeding his garden and noticed that many of his *tulips* were wilting. How likely do you feel used coffee grounds will produce healthy *tulips* by increasing the nitrogen content of the soil?

We chose transfer targets that varied in how categorically similar they were to the transfer source. Two of the targets were at the same subordinate level as the transfer source, three transfer targets were in the same basic level category as the source, three were in the same superordinate, and one transfer target represented a category that was in a different superordinate category (see Figure 1 for a general example). For the transfer source example of pink roses, the transfer targets would be as follows: red roses and white roses; roses, tulips, and orchids; flowers, trees, and cactus; and earthworms. For each transfer target, participants were prompted to judge whether the mechanism from the scenario would transfer to the target in question using a 1 (unlikely to transfer) to 7 (extremely likely to transfer) scale.




Figure 1: Example transfer targets by category level for the source pink roses.

Participants also provided ratings for how similar they believed each of the transfer targets were to the transfer source in the scenario. Similarity judgments were made using a 1 (not at all similar) to 7 (extremely similar) scale.

Design and Procedure The study was programmed and distributed using the Qualtrics survey software. Participants first read instructions describing that they would be presented with a series of scenarios containing cause and effect relationships that they would use to make judgments about other situations. A brief example scenario was provided in the instructions. Participants then began the transfer task. In the transfer task, participants were first presented with one of the scenarios, chosen at random. After reading through the scenario, the transfer targets associated with that scenario were presented, one at a time, in a random order. For each transfer target, participants judged the likelihood that the causal mechanism in the scenario would also produce the effect in the transfer target. Each of the 12 scenarios were presented in a random order, with each scenario's

corresponding transfer targets also being presented in a random order. After completing the transfer task, participants provided similarity ratings and then completed a brief demographics questionnaire.

Results

Transfer Judgments We examined the effects of category level and domain on transfer judgments using a 3 (domain: natural, health, mechanical) x 4 (category level: subordinate, basic, superordinate, different-superordinate) repeated measures ANOVA. The ANOVA revealed significant main effects of category level ($F(3, 50) = 66.5, p < .001, \eta_p^2 = .78$) and domain ($F(2, 50) = 19.22, p < .001, \eta_p^2 = .32$), as well as a significant category level by domain interaction, $F(6, 50) = 20.79, p < .001, \eta_p^2 = .27$. We explored the interaction with Sidak-corrected *t*-tests. Comparing judgments between the four levels found significant differences between all four levels for the natural ($ps < .001$) and health ($ps < .001$) domains (Figure 2). However, participants' ratings did not differ between the subordinate and basic level targets ($p = .95$) in the mechanical domain. All other comparisons of transfer judgments between category level were significant in the mechanical domain ($ps < .03$).




Figure 2: Mean transfer judgments for Experiment 1

Effect of Category Hierarchy A byproduct of the category hierarchies we used to produce transfer targets is that some transfer targets belonged to the same hierarchy as the transfer source (e.g., pink roses and flowers), and some transfer targets did not belong to that hierarchy (e.g., pink roses and trees). This occurred at both the basic and the superordinate level. To explore these differences, we conducted a 2 (category hierarchy: inside vs. outside) x 3 (domain) x 2 (category level: basic vs. superordinate) repeated measures ANOVA on transfer judgments. The ANOVA revealed significant main effects of hierarchy membership ($F(1, 50) = 128.9, p < .001, \eta_p^2 = .72$), category level ($F(1, 50) = 98.7, p < .001, \eta_p^2 = .66$), and domain ($F(2, 100) = 20.43, p < .001, \eta_p^2 = .29$). The main effect of category hierarchy suggests that transfer judgments are significantly greater for inside ($M =$




Figure 3: Mean transfer judgments for a) inside- and b) outside-hierarchy targets in Experiment 1.

5.4; $SE = .100$) than outside ($M = 4.47$; $SE = .085$) the category hierarchy. Significant two-way interactions between hierarchy membership and domain ($F(2, 100) = 36.58$, $p < .001$, $\eta_p^2 = .422$), as well as category level and domain ($F(2, 100) = 9.21$, $p < .001$, $\eta_p^2 = .155$) were also found. A significant three-way interaction was also found between hierarchy membership, category level, and domain, $F(2, 100) = 3.57$, $p = .032$, $\eta_p^2 = .067$. Given the three-way interaction, we looked at our follow-up comparisons separately for inside- and outside-hierarchy judgments. Critically, we find significant drops in transfer judgments in all three domains when using Sidak-corrected t -tests to compare the basic to superordinate-level targets ($ps < .001$) inside the category hierarchy of the source item (Figure 3). When comparing transfer judgments between basic and superordinate level transfer targets outside the source item's category hierarchy, we find identical patterns in the natural and health domains ($ps < .001$), but not in the mechanical domain ($p = .085$). This difference in the mechanical domain for outside category hierarchy but not inside category members helps explain the significant three-way interaction.

Discussion

Overall, we found that participants were willing to transfer the action of a novel causal mechanism from one source to other targets. Importantly, this transfer did not happen equally among all targets. Participants provided the highest transfer judgments for the most categorically similar transfer targets (i.e., subordinate level). These judgments then decreased as the targets were more distant from the source in a category hierarchy.

This first study suggests that people are sensitive to the category relatedness of a target when transferring causal mechanism information. We next tested whether other types of causal information would be similarly transferred. Specifically, we test whether the strength of a causal relationship will be thought to transfer along with its mechanism.

Experiment 2

When provided with information that implies the presence of a strong cause, will people similarly believe that the cause will have the same strength in a novel context? In this experiment we present people with transfer prompts that specify that the same causal mechanism applies to a source as a target. We then ask people how likely the same strength of a relationship will also be present. One possibility is that people will transfer any type of causal information they can. Knowing that a mechanism is the same may license thinking that it will produce the effect to the same extent to categorically similar targets (e.g., mutability; Sloman et al., 1998; unbroken mechanism hypothesis; Hagmayer et al., 2011). Alternatively, people may believe that the strength of causal relationships are unique in that a mechanism operates with a particular strength only in a specific context (e.g., Heit, 1998). Moving beyond that context (i.e., to transfer targets) could attenuate the strength of the that mechanism. We test these possibilities in the following experiment.

Method

Participants Undergraduate students ($N=54$) participated in exchange for course credit. All participants completed the study remotely on a computer of their choice. Participants were screened to ensure that they had normal or corrected-to-normal vision and were fluent in English. Approval for both studies was obtained from the authors' university IRB.

Materials, Design, and Procedure Experiment 2 was identical to Experiment 1, save for the following changes. Instead of judging whether the effect in the scenario would be brought about by a shared causal mechanism, participants made ratings of causal strength transfer; that is, whether the effect would be brought about to the same extent in each

scenario's corresponding targets. Additionally, mechanism information was only included in the scenario, not the prompt. We made this change so that judgments were made on the basis of the strength of the relationship, rather than a hybrid of strength and mechanism plausibility (e.g., a different mechanism could produce the effect with the same strength). An example transfer scenario and prompt from the natural domain follows, with the italicized item representing one of the possible transfer targets.

Scenario: One day, when Chuck was weeding his garden, he noticed that many of his pink roses were wilting. After some research, he learned that he could treat the soil in his garden with used coffee grounds. This treatment strongly causes pink roses to become healthy by increasing the nitrogen content of the soil, which is associated with healthy pink roses. So, he administered the used coffee grounds to the soil in his garden. The next morning when he woke up, he noticed that the treatment was effective in making almost all of the pink roses that were wilting stand straight and healthy.

Transfer Prompt: Imagine Chuck was weeding his garden and noticed that many of the *tulips* in his garden were wilting. If Chuck applied used coffee grounds to the soil, to what extent do you believe that the used coffee ground treatment would cause *tulips* to become healthy again? Please respond using the following scale.

In each of the target scenarios, we described the causal strength of the given mechanism as strong for the source item. Participants made transfer judgments for causal strength using an onscreen slider on a 0 (no relationship) to 100 (strongly causes) scale.

Results and Discussion

Transfer Judgments The effects of category level and domain on transfer judgments were examined using a 3 (domain) x 4 (category level) repeated measures ANOVA.




Figure 4: Mean transfer judgments for Experiment 2

The ANOVA revealed significant main effects of category level ($F(3, 141) = 128, p < .001, \eta_p^2 = .73$) and domain ($F(2, 94) = 20.8, p < .001, \eta_p^2 = .31$), as well as a significant category level by domain interaction, $F(6, 50) = 20.79, p < .001, \eta_p^2 = .24$. We explored the interaction with Sidak-corrected *t*-tests. Patterns of transfer judgments in Experiment 2 were quite similar to Experiment 1 (Figure 4). Comparing judgments between the four levels found significant differences between all four levels for the natural ($ps < .001$) and health ($ps < .001$) domains. However, participants' ratings did not differ between the subordinate and basic level targets ($p = .571$) for the mechanical domain. All other comparisons of transfer judgments between category level were significant in the mechanical domain ($ps < .002$).

Effect of Category Hierarchy Again, we explored potential differences between targets belonging to different category hierarchies (see Figure 5). We conducted a 2 (category hierarchy: inside vs. outside) x 3 (domain) x 2 (category level: basic vs. superordinate) repeated measures ANOVA on transfer judgments. The ANOVA revealed significant main




Figure 5: Mean transfer judgments for a) inside- and b) outside-hierarchy targets in Experiment 2.

effects of hierarchy membership ($F(1, 53) = 187.9, p < .001, \eta_p^2 = .78$), category level ($F(1, 53) = 98.8, p < .001, \eta_p^2 = .65$), and domain ($F(2, 106) = 25.01, p < .001, \eta_p^2 = .32$). As in Experiment 1, the main effect of category hierarchy suggests that transfer judgments are significantly greater for inside ($M = 67.8; SE = 1.49$) than outside ($M = 48.6; SE = 1.67$) the category hierarchy. Significant two-way interactions between hierarchy membership and domain ($F(2, 106) = 37.5, p < .001, \eta_p^2 = .35$), hierarchy membership and category level ($F(1, 53) = 10.7, p < .01, \eta_p^2 = .17$), and category level and domain ($F(2, 106) = 12.7, p < .001, \eta_p^2 = .19$), were also found. Unlike Experiment 1, no significant three-way interaction was found ($p = .214$). Looking at the category level by domain interaction with Sidak corrected t -tests, we find that there is a significant drop from the basic to the superordinate in transfer judgments for all domains, $ps < .001$. The lack of a three-way interaction suggested that the pattern for transfer was the same for inside and outside category hierarchy members.

General Discussion

Across two experiments, we demonstrated that people are willing to transfer causal knowledge to novel situations. When presented with a novel causal relationship, people will transfer both the mechanism by which that causal relationship happens and the strength of that relationship to other situations. The degree of this transfer is moderated by how categorically similar the targets are to the transfer source. In this way, people more promiscuously generalize novel causal information to close category members, but still generalize to weaker degrees as category relatedness decreases.

One interesting finding in these data are the differences by domain in transfer patterns. Similar patterns of transfer were observed in scenarios belonging to the natural and health domains. The mechanical domain differed in that participants transferred equally to subordinate and basic level categories. Differences between these two patterns of transfer could be attributed to psychological essentialism (Ahn et al., 2013), in that natural kind categories (represented here as the natural and health domains) are distinguished as possessing rigid category boundaries (Prasada et al., 2012). Artifact categories (represented here as the mechanical domain), however, possess more arbitrary, flexible category boundaries (Brandone & Gelman, 2013). It is possible that the perceived arbitrariness of artifact categories allowed for more transfer across farther reaching category levels.

An open question from our results is how far are people willing to transfer causal knowledge. In all scenarios, non-trivial transfer judgments were produced for even the most categorically dissimilar (i.e., belonging to a different superordinate category) targets. That is, when learning coffee grounds provide nitrogen that is healthy for roses, people believe this nitrogen will be healthy for earthworms to some extent as well. While such transfer could be surprising, people are adept at spontaneously inferring causal links (Hassin et al., 2002). Just given time to think about how

nitrogen could intervene on earthworms may allow people to generate their own causal explanation of how the mechanism works in this category. Future research could elicit participants' confidence in or explanations for their judgments to explore if people are creating their own causal stories that could link these distant targets and how far people are willing to go in making these connections.

A lingering question from our results is whether judgments were informed by category hierarchy knowledge or similarity. We did not present analysis of similarity judgments due to space constraints. Overall, subordinate-level targets were rated to be highly similar, and targets belonging to different superordinate categories were the most dissimilar. Similarity judgments were highly correlated with transfer judgments, producing collinearity issues in the data that made it difficult to specifically partial out the effect of similarity. Overall, similarity ratings mimicked transfer ratings except that we saw a step-wise decrease in similarity for the mechanical domain where transfer judgments did not show the same step-wise decrease in Experiments 1 and 2. Future work is needed to explore the role of similarity separately from category hierarchy. For example, perceptual similarity could be manipulated using visual images. In nature, many unrelated organisms exhibit Batesian mimicry, where an unthreatening or defenseless organism evolves to mimic the appearance of a more threatening organism (e.g., various species of hoverfly, a harmless winged insect, evolved to become nigh-indistinguishable from dangerous bees and wasps). Using stimuli depicting categorically distinct, but perceptually similar, exemplars, would help further explore the boundaries for causal knowledge transfer.

We explored two specific types of causal knowledge that could be transferred, mechanism and strength. Future research should explore what other causal elements can be transferred. For example, if a cause is associated with particular side effects or byproducts, will people similarly believe that those side effects or byproducts will be present in a novel situation? Likewise, if multiple causes are capable of producing an effect, would this multiply determined causality transfer to other items? These questions revolve around whether the causal structure believed to underlie the production of an effect would transfer to new situations. These are questions we are currently exploring.

Across two experiments, we showed that people will transfer causal knowledge to novel contexts. The extent to which this causal knowledge was transferred varied as a function of the transfer target's categorization. People regularly make inferences about how interventions will work on targets that they have no explicit knowledge basis for, such as a fertilizer often used for roses being applied to tulips. As such, the field of causal reasoning benefits from research of this nature that highlights the role of prior knowledge in causal inference. These studies serve as an initial investigation into the thought processes that guide causal inferences informed by existing causal knowledge, and future research will illuminate the underlying process of generalizing causal knowledge.

References

- Ahn, W., Taylor, E. G., Kato, D., Marsh, J. K., & Bloom, P. (2013). Causal Essentialism in Kinds. *Quarterly Journal of Experimental Psychology*, 66(6), 1113–1130. <https://doi.org/10.1080/17470218.2012.730533>
- Brandone, A. C., & Gelman, S. A. (2013). Generic language use reveals domain differences in young children's expectations about animal and artifact categories. *Cognitive Development*, 28(1), 63–75.
- Feeney, A., & Heit, E. (2011). Properties of the diversity effect in category-based inductive reasoning. *Thinking & Reasoning*, 17(2), 156–181. <https://doi.org/10.1080/13546783.2011.566703>
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General*, 140(2), 168–185. <https://doi.org/10.1037/a0022100>
- Garcia-Retamero, R., Müller, S. M., Catena, A., & Maldonado, A. (2009). The power of causal beliefs and conflicting evidence on causal judgments and decision making. *Learning and Motivation*, 40(3), 284–297.
- Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7(2), 155–170. https://doi.org/10.1207/s15516709cog0702_3
- Hagmayer, Y., & Mayrhofer, R. (2013). Hierarchical Bayesian models as formal models of causal reasoning. *Argument & Computation*, 4(1), 36–45.
- Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General*, 138(1), 22–38. <https://doi.org/10.1037/a0014585>
- Hassin, R. R., Bargh, J. A., & Uleman, J. S. (2002). Spontaneous causal inferences. *Journal of Experimental Social Psychology*, 38(5), 515–522. [https://doi.org/10.1016/S0022-1031\(02\)00016-1](https://doi.org/10.1016/S0022-1031(02)00016-1)
- Hayes, B. K., & Heit, E. (2018). Inductive reasoning 2.0. *WIREs Cognitive Science*, 9(3), e1459. <https://doi.org/10.1002/wcs.1459>
- Heit, E. (1998). Influences of prior knowledge on selective weighting of category members. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(3), 712.
- Heit, E., Hahn, U., & Feeney A. (2004) Defending diversity. In Ahn W, Goldstone R. L., Love, B. C., Markman, A. B., Wolff P., editors. *Categorization inside and outside the lab: Festschrift in honor of Douglas L Medin*. Washington, DC: American Psychological Association.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(2), 411.
- Holyoak, K. J., & Lee, H. S. (2017). Inferring causal relations by analogy. Oxford library of psychology. *The Oxford handbook of causal reasoning*, 459–473.
- Kemp, C., Goodman, N. and Tenenbaum, J. B. 2010. Learning to ILearn causal models. *Cognitive Science*, 34: 1185–1243. ([doi:10.1111/j.1551-6709.2010.01128.x](https://doi.org/10.1111/j.1551-6709.2010.01128.x))
- Malt, B. C., & Smith, E. E. (1982). The role of familiarity in determining typicality. *Memory & Cognition*, 10(1), 69–75.
- Mayrhofer, R., & Waldmann, M. (2011). Heuristics in covariation-based induction of causal models: Sufficiency and necessity priors. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97(2), 185.
- Prasada, S., Hennefield, L., & Otap, D. (2012). Conceptual and Linguistic Representations of Kinds and Classes. *Cognitive Science*, 36(7), 1224–1250. <https://doi.org/10.1111/j.1551-6709.2012.01254.x>
- Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: the effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, 130(3), 323.
- Shanks, D. R. (2004). Judging covariation and causation. *Blackwell handbook of judgment and decision making*, 220–239.
- Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology*, 25(2), 231–280.
- Shepard, R. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317–1323.