

Determinants and Consequences of the Need for Explanation

Emily G. Liquin (emilyliquin@berkeley.edu)

Tania Lombrozo (lombrozo@berkeley.edu)

Department of Psychology, University of California, Berkeley
Berkeley, CA 94720 USA

Abstract

Much of human learning throughout the lifespan is achieved through seeking and generating explanations. However, very little is known about what triggers a learner to seek an explanation. In two studies, we investigate what makes a given event or phenomenon stand in need of explanation. In Study 1, we show that a learner's judgment of "need for explanation" for a given question predicts that learner's likelihood of seeking an answer to this question. In Study 2, we explore several potential predictors of need for explanation. We find that the need for explanation is greater for questions expected to have useful answers that require expert understanding, and that "need for explanation" can be differentiated from general curiosity.

Keywords: explanation; curiosity; information search

Why does the moon cause tides? Why do continents move? Why does yawning affect our ears? Much of human learning, from childhood through adulthood, is achieved by asking questions. Children begin to ask information-seeking questions before the age of two, and they ask increasingly more *explanation-seeking* questions between the ages of two and three (Chouinard, Harris, & Maratsos, 2007; Hickling & Wellman, 2001). Explanation remains an important learning mechanism throughout development: children ages two to six use self-generated explanations to guide their exploration (Legare, 2012), and adults use explanatory principles to guide their inferences (Lombrozo, 2016).

Despite the importance of explanation to learning and inference throughout the lifespan, very little research has attempted to identify what *triggers* explanation search in the first place. This could be because explanation-seeking behavior, by its very nature, is difficult to capture in the lab: research has typically relied on a small number of observations "in the wild" or on highly artificial contexts in the lab. The present research circumvents this challenge by drawing upon a novel source of data: large-scale online databases of user-generated questions that allow us to better capture explanation search in the course of everyday life.

Following Grimm (2008) and Wong and Yudell (2015), we propose that explanation-seeking behavior is triggered by *need for explanation* (NFE), the sense that a given event, phenomenon, or claim demands an explanation. In our first study below, we show that people's judgments of NFE indeed predict their explanation-seeking behavior: the questions that generate high ratings of need for explanation are those that

participants are more likely to select when given the option to reveal a subset of answers.

In our second study, we distinguish NFE from general curiosity and test several accounts of the need for explanation. These accounts are drawn from prior work on need for explanation within philosophy, on psychological accounts of attribution and curiosity, and from prior work on explanation and learning. Before turning to our studies, we review these accounts below.

Potential Determinants of NFE

Only two theoretical proposals, both within philosophy, have directly addressed need for explanation. Grimm (2008) suggests that an event stands in need of explanation to the extent that there is a relevant alternative way the world could have been. He poses this theory by considering the contrast between *fact* – what actually occurred – and *foil* – what could have occurred instead. Need for explanation is determined by the salience of this contrast, and an explanation satisfies this need if it identifies what made a difference to the occurrence of the fact as opposed to the foil. Elements of this proposal are echoed by Bruckmüller, Hegarty, Teigen, Boehm, and Luminet (2017), who consider what makes a historical event stand in need of explanation. Among other factors, they suggest that an event is judged by its departure from *background event norms*, and that events that are more different from these norms tend to demand explanation.

Wong and Yudell (2015) offer an alternative to Grimm's "fact-and-foil" account, which they call the "map" account. They suggest that an event stands in need of explanation to the extent that the event does not fit one's theory (or *map*) of the world. The map account is best illustrated by example: imagine that a colleague wins the lottery exactly one day after he received a fortune cookie stating that he would win the lottery the next day. In this example, the event stands in need of explanation because it does not fit our map of prediction in this domain: according to this map, fortune-telling provides generally unreliable predictions because it is not based on real-world causal relations.

While Wong and Yudell (2015) pose their account as normative rather than descriptive, their motivating idea is supported by several related bodies of research in psychology. For example, there is evidence that explanation is linked to *surprise*¹ or *violation of expectation*. Young children choose to generate explanations for events that are

¹ Grimm (2008) and Wong and Yudell (2015) are quick to reject surprise as an adequate determinant of need for explanation. However, they take surprising events to be merely *low probability* events, while we are open to more complex psychological accounts

of surprise, such as that proposed by Foster and Keane (2015). These accounts of surprise are more similar to Wong and Yudell's map account than to basic low probability accounts.

inconsistent with their hypotheses (Legare, Gelman, & Wellman, 2010). Additionally, research on spontaneous attribution finds that people often explain surprising outcomes, as well as outcomes with *negative valence* (Weiner, 1985). Foster and Keane (2015) demonstrate that judgments of surprise are related to the extent to which an event cannot be easily explained, again demonstrating a strong link between surprise and explanation.

Several other possible determinants of need for explanation can be derived from related literatures. A growing body of research shows that explanations are generated and evaluated by the extent to which they adhere to principles like simplicity (Bonawitz & Lombrozo, 2012; Lombrozo, 2007; Pacer & Lombrozo, 2017) and breadth (Johnson, Johnston, Toig, & Keil, 2014; Preston & Epley, 2005). If explanation-seeking processes are well calibrated to these explanatory preferences, need for explanation could be triggered when one believes that a simple or broad explanation is likely to be available. Similarly, we might consider the *function* of explanation, as need for explanation could be triggered when this function is likely to be realized. One plausible function of explanation is to drive learning and increase understanding (Gopnik, 2000; Lombrozo, 2006, 2016). If this is the case, NFE could arise when one detects that there is something to be learned from an event or phenomenon.

Another proposal can be derived by considering the process of explanation search. Often, questions seem to be directed towards those with relevant expertise: we ask doctors about medical treatment, and chefs about food. This pattern of explanation search aligns with the idea of a *division of cognitive labor* (e.g., Keil, 2003; Keil, Stein, Webb, Billings, & Rozenblit, 2008; Lutz & Keil, 2002): knowledge is clustered in other minds, so we must be able to track which people have expertise in which areas. If explanation-seeking processes are sensitive to this division of cognitive labor, need for explanation could be triggered when expertise in some area is needed to answer a given question.

Finally, NFE seems to be strongly linked to *curiosity*, though not all curiosity is directed towards the acquisition of explanations (e.g., Berlyne & Borsa, 1968; Nicki, 1970). If this is the case, we might expect that the determinants of curiosity could also be determinants of NFE. While several theories of curiosity have been proposed, there is currently no consensus as to what constitutes curiosity (Kidd & Hayden, 2015). Loewenstein (1994) proposed the well-known “information gap” account of curiosity: one’s curiosity is piqued when there is a modest amount of information that can be gained. Other classic proposals emphasize how novelty can invoke curiosity – both in non-human animals (e.g., Pavlov, 1927) and in young children (e.g., Smock & Holt, 1962). Finally, a recent proposal (Dubey & Griffiths, 2017) unites these two perspectives, arguing that curiosity functions to increase the value of an agent’s current knowledge. An agent should thus direct their curiosity towards phenomena with maximal future utility. While these theories of curiosity still remain up for debate, each of the factors highlighted above could affect need for explanation.

Present Research

In the present research, we first verify that need for explanation predicts actual explanation-seeking. If this is the case, we would expect people to preferentially search for answers to questions that they also judge to be most in need of explanation.

Second, we evaluate the accounts of need for explanation previously described: fact-and-foil, map mismatch, surprise, negative valence, simplicity/breadth, anticipated learning, expertise, information gap, novelty, and future utility. We also confirm that NFE and curiosity are related but distinct constructs, by demonstrating that these potential determinants differentially predict NFE and curiosity.

Study 1

In Study 1, we investigate whether people’s judgments of need for explanation (NFE) predict their search for explanations. To do this, we present participants with several questions and ask them to rate each question on the extent to which it demands an explanation. Subsequently, participants are asked to select the questions for which they would *most* like to receive an explanation. Verifying the link between need for explanation and explanation search is critical if we hope to derive general conclusions about explanation-based learning.

Method

Participants Participants in Study 1 were 72 adults recruited from Amazon Mechanical Turk (38 male, 34 female, ages 19-69, $M = 37$). Participation was restricted to MTurk workers with IP addresses in the United States, who had completed at least 50 HITs with a minimum approval rating of 95%. Eight additional participants completed the study but were excluded from analyses due to failure to pass an attention check (described below).

Materials Thirty “why” questions were sampled from the Reddit *Explain Like I’m Five* webpage (www.reddit.com/r/explainlikeimfive). On this page, Reddit users submit questions and other users provide easy-to-understand explanations. The 30 questions were edited lightly for grammar and readability. Sample questions include: “Why are most pencils painted yellow? Why not a different color?”; “Why do ice cubes crackle when liquid is poured on them?”; and “Why does the bowl always get hotter than the food in the microwave?” For each question, we selected an answer provided by another user on the basis of its succinctness and accuracy; these were also edited lightly for grammar and readability.

Procedure Each participant saw eight questions, randomly selected from the 30 questions described above. For each question, participants were asked to rate the extent to which the question demands an explanation on a 7-point scale, with the endpoints marked “definitely does not demand an explanation” and “definitely demands an explanation.” In a

separate block, participants were also asked to rate the extent to which they accepted the premise of each question (“Please rate your agreement with the following: that ice cubes crackle when liquid is poured on them”) and the extent to which they believed they knew the answer to the question (“How confident are you that you know the answer to this question?”). These two measures were included as controls, to ensure that any effect of NFE on explanation search could not be accounted for by acceptance of the premise or by perceived knowledge. These two blocks (NFE and control measures) were presented in a random order.

After completing these ratings, participants completed a short distractor task, which involved performing simple arithmetic problems. This task doubled as an attention check, as the task required participants to pay attention and remember what they had seen previously. Participants who made more than two errors were excluded.

After the distractor task, participants were presented with the eight explanations they had rated, and were asked to select three questions that they wanted answered. This was the primary dependent variable of interest.

Finally, participants read the explanations to the three questions they selected and completed a satisfaction rating after each. This ensured that they remained engaged and did not exit the survey prematurely. After reading the three explanations, participants provided their age and gender.

Results

To investigate the relationship between NFE and explanation search, we fit a generalized linear mixed-effects model, using a logit link function. NFE was included as a fixed effect, predicting the probability of choosing to reveal the answer to a question. The control measures (perceived knowledge and acceptance of each question’s premise) were also included as fixed effects. Random slopes were included for participant, and random intercepts were included for item. The model coefficient for NFE was highly significant, $b = 0.18$, 95% CI [0.09, 0.29], $z = 3.59$, $p < .001$. The model coefficient for perceived knowledge of each question’s answer was also significant, $b = -0.10$, 95% CI [-0.19, -0.01], $z = -2.21$, $p = .03$, but the control measures only made a marginally significant contribution to the full model over the baseline model containing only NFE, $\chi^2(2) = 5.11$, $p = .08$. Each 1-unit increase in NFE for a given question predicted a 120% increase in the probability of selecting that question to be explained (see Figure 1). These results suggest that NFE predicts explanation search.

Discussion

Study 1 confirms that judgments of need for explanation are predictive of explanation-seeking behavior: the probability of choosing to receive an explanation for a given question was strongly predicted by the extent to which that question was judged to demand an explanation, even controlling for other potentially relevant factors. Having established this functional consequence of NFE, we next turn to investigate potential determinants of NFE.

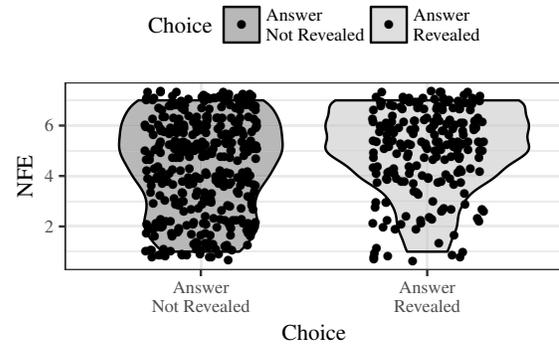


Figure 1: Study 1 NFE judgments, plotted according to whether the question was later selected to be answered.

Study 2

In Study 2, we test potential determinants of NFE judgments. To do so, we present participants with several questions, for which they judge NFE and additionally complete ratings related to the hypotheses outlined in the introduction. We can then evaluate which ratings best predict judgments of NFE.

We also clarify the relationship between NFE and curiosity. One possibility is that NFE and curiosity are effectively equivalent, at least in the context of our task. A second possibility is that NFE is a sub-type of general curiosity: it is the “explanation-seeking” flavor of curiosity. A third possibility is that NFE and curiosity are distinct constructs. In Study 2, we begin to investigate these possibilities.

Method

Participants Participants in Study 2 were 74 adults recruited from Amazon Mechanical Turk (40 male, 33 female, 1 other; ages 18-60, $M = 33$). Participation was restricted to MTurk workers recruited as in Study 1. Five additional participants who failed to pass an attention check (described below) were excluded.

Materials Fifty-three new “why” questions (that were not used in Study 1) were sampled from the Reddit *Explain Like I’m Five* webpage. Again, the questions were edited lightly for grammar and readability.

Procedure Each participant saw ten questions, randomly selected from the 53 questions described above. For each question, participants were asked to complete the same NFE rating as in Study 1. Additionally, participants rated their curiosity in response to each question: “To what extent are you curious about the answer to this question?” Finally, participants rated each question on eleven other dimensions, summarized in Table 1. These 13 judgments were presented in a random order.

Next, participants were shown 10 questions. Four of these questions were questions that they had seen previously, while

Table 1: Potential determinants of NFE and their relationship with NFE in Study 2 and Study 2 replication.

Potential Determinant	Item (answered on a seven-point scale)	Study 2	Replication
Fact-and-Foil	How easily can you imagine a world in which it is not the case that [premise]?	-0.04	-0.007
Map Mismatch	How well does the claim that [premise] fit in with your current beliefs about the world?	0.08*	0.05
Surprise	To what extent do you think it is surprising that [premise]?	0.34**	0.33**
Negative Valence	To what extent do you think the main claim of the question (that [premise]) is negative, positive, or neutral?	-0.08*	-0.02
Simplicity	Do you think the answer to this question is likely to be simple or complex?	-0.43**	-0.38**
Breadth	Do you think the answer to this question would tell you something that applies only to what is being explained, or would it tell you something that applies more broadly to other cases that are similar?	0.14**	0.12**
Anticipated Learning	To what extent do you think the answer to this question would teach you something new?	0.43**	0.47**
Expertise	Do you think that answering this question requires special expertise in some domain?	0.44**	0.43**
Information Gap	How much do you know about the topic of this question?	0.16**	0.17**
Novelty	How novel for you is the claim that [premise]?	0.27**	0.26**
Future Utility	To what extent would knowing the answer to this question be useful to you in the future?	0.39**	0.40**

* $p < .05$; ** $p < .001$

six were new “why” questions sampled from the Answers.com “science” category (www.answers.com).

Participants were asked to select the questions they had seen in the previous part of the study, and were excluded if they made more than two errors.

Results

First, to investigate general relationships between the 12 measures of interest and NFE, independent of each measure’s correlation with other measures, we fit 12 separate regression models to the data, predicting NFE with each measure as a fixed effect in a separate model. Each model also included random intercepts for participant and item. Regression coefficients for each model were standardized, to enable comparison between measures. The regression coefficients for each measure (as well as the same coefficients from a replication study²) are reported in Table 1. These results reveal the potential importance of several measures. Notably, fact-and-foil was *not* significantly related to NFE in both Study 2 and the Study 2 replication. In the replication, map mismatch and negative valence also failed to reach significance.

Next, we turned to an investigation of the differential impact of these predictors on NFE and curiosity. As many of the measures were correlated (with r as high as 0.60) we performed exploratory factor analysis to find a smaller number of orthogonal dimensions, which could then be used to predict NFE and curiosity in a single regression analysis. Initial eigenvalues for the first four factors were over one, and

explained 27%, 17%, 10%, and 9% of the variance, respectively. Inspection of the scree plot suggested a three-factor solution. Three and four factor solutions were inspected, using varimax rotation. The four-factor solution was retained due to simple structure and greater interpretability.

Three potential determinants – fact-and-foil, negative valence, and breadth – were eliminated because they failed to load onto any factor with a factor loading above 0.4. The final four-factor solution with varimax rotation, excluding these three items, explained 59% of the total variance. The four

Table 2: Factor loadings for final four-factor solution. Loadings over 0.4 are printed in bold.

Potential Determinant	Factor 1	Factor 2	Factor 3	Factor 4
Surprise	0.776	0.133	0.161	0.184
Map Mismatch	0.565	-0.051	0.121	-0.169
Information Gap	0.244	0.084	0.558	-0.103
Novelty	0.667	0.094	0.289	0.147
Future Utility	0.024	0.184	0.021	0.674
Anticipated Learning	0.258	0.258	0.736	0.364
Expertise	0.064	0.515	0.189	0.307
Simplicity	-0.049	-0.985	-0.109	-0.103

² This study was also replicated with an additional sample of 98 participants (plus two excluded from analysis) from Amazon Mechanical Turk, using a different set of 50 questions from *Explain*

Like I’m Five. All findings that were not successfully replicated are noted in the Results.

factors corresponded roughly to “prior knowledge” (surprise, map mismatch, and novelty), “expertise” (simplicity – negative factor loading – and expertise), “information content” (anticipated learning and information gap), and “future utility” (future utility). The factor loading matrix for this analysis is presented in Table 2.³

Thompson scores were extracted for use in a simultaneous regression predicting NFE and curiosity. This regression model included prior knowledge, expertise, information content, and future utility as fixed effects, along with a fixed effect for dependent measure (NFE vs. curiosity, effects coded). The interactions between each factor and the dependent measure were also included as fixed effects. Random intercepts were included for participant and item. This analysis revealed a significant interaction between factor score and dependent measure (NFE vs. curiosity) for three factors: expertise, $b = -0.12$, 95% CI [-0.19, -0.05], $t(1351) = -3.26$, $p = .001$, information content, $b = 0.20$, 95% CI [0.11, 0.28], $t(1351) = 4.42$, $p < .001$, and future utility, $b = 0.18$, 95% CI [0.08, 0.27], $t(1351) = 3.64$, $p < .001$. These results (illustrated in Figure 2) indicate that expertise is a stronger predictor of NFE than of curiosity, while information content and future utility are stronger predictors of curiosity than of NFE. While the interaction between prior knowledge and dependent measure was not significant, the regression coefficient for prior knowledge was significant, $b = 0.40$, 95% CI [0.30, 0.49], $t(974.30) = 8.21$, $p < .001$, indicating that prior knowledge was modestly predictive of both curiosity and NFE.

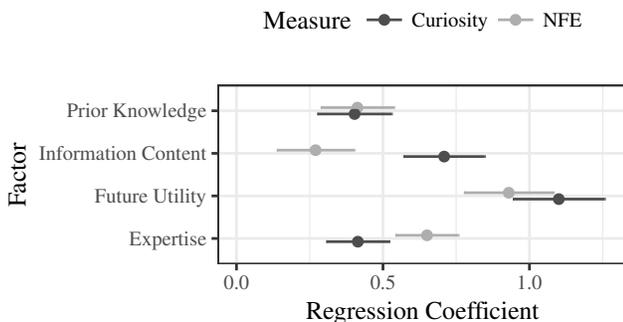


Figure 2: Study 2 model coefficients (\pm 95% CI) for each factor, from regression analysis predicting NFE & curiosity.

Discussion

These results suggest that some of the tested determinants of NFE are more important than others. Notably, the fact-and-foil account (Grimm, 2008) and the negative valence account (Weiner, 1985) fail to capture substantial variance in NFE ratings. The map mismatch account (Wong & Yudell, 2015) also explained little variance in NFE on its own.

However, a regression analysis on factor scores revealed

³ There were modest differences in the results of this factor analysis on the replication dataset. While the data were slightly more consistent with a three-factor solution, a four-factor solution produced more interpretable results. After removing the measures that did not load onto any factor in the initial analysis (fact-and-foil,

several important predictors of NFE: future utility (composed solely of future utility scores) and expertise (composed of expertise scores and reverse-scored simplicity scores – i.e., complexity scores) were highly predictive of NFE. Prior knowledge (composed of surprise, map mismatch, and novelty scores) and information content (composed of anticipated learning and information gap scores) were also modestly predictive of NFE.

These results also suggest that NFE and curiosity are closely related, but that NFE is not equivalent to curiosity. While curiosity judgments were predicted more strongly by information content and future utility, NFE judgments were predicted more strongly by expertise. These interactions are consistent with the idea of a division of cognitive labor: explanation search is critical when desired knowledge is stored in another person’s mind (Keil, 2003). Curiosity, on the other hand, may be satisfied by any number of exploratory behaviors: experimentation, observation, or explanation search, to name a few. If this is correct, it makes sense that NFE might depend more crucially on expertise and expected complexity, while curiosity might track information content and future utility more generally.

General Discussion

In this research, we sought to investigate *need for explanation*: why some events or phenomena seem to demand an explanation, while others do not. Here, we answer three questions about need for explanation. First, is NFE predictive of actual explanation-seeking behavior? Second, what factors determine NFE judgments? And third, how (if at all) is NFE distinct from curiosity?

To the first question, the answer is yes: NFE judgments were highly predictive of explanation-seeking behavior, even controlling for possible related confounds. Our investigation of the second two questions also revealed several important insights: NFE can be predicted by prior knowledge, expertise, information content, and future utility. Further, it is distinct from curiosity in that it more closely tracks judgments of expertise, while curiosity more closely tracks information content and future utility.

There are several limitations of these studies. One important limitation is that the experimental materials consisted of questions that were actually posed. It is natural to assume that these “why” questions already meet some threshold of NFE beyond which explanations are sought. This may limit our conclusions, especially in Study 2: some potential determinants of NFE may have fared poorly because of this restricted range.

Additionally, our investigation was restricted to particular kinds of explanations, namely those offered in response to “why” questions about predominantly factual matters. Just as explanations vary in form and content, so too might NFE.

negative valence, and breadth), the factor structure was nearly equivalent to that reported above, aside from the failure of information gap to exceed 0.4 in its loading on the “information content” factor.

Third, our study designs were correlational. For the time being, we cannot draw strong conclusions about whether the factors identified are in fact *determinants* of NFE as opposed to consequences of NFE, or covariates predicted by some common cause. Future research must employ experimental methods to determine any causal influence of the factors of interest identified in this preliminary research. That said, our findings offer the first empirical demonstration that need for explanation is a construct that predicts explanation-seeking and differs from general curiosity, and we offer an initial glimpse into its determinants.

Acknowledgements

This work was supported by a McDonnell Foundation Scholar Award in Understanding Human Cognition to TL and an NSF Graduate Research Fellowship to EL [DGE 1752814]. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the McDonnell Foundation or the National Science Foundation.

References

- Berlyne, D. E., & Borsari, D. M. (1968). Uncertainty and the orientation reaction. *Attention, Perception, & Psychophysics*, 3(1), 77–79.
- Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology*, 48(4), 1156–1164.
- Bruckmüller, S., Hegarty, P., Teigen, K. H., Böhm, G., & Luminet, O. (2017). When do past events require explanation? Insights from social psychology. *Memory Studies*, 10(3), 261–273.
- Chouinard, M. M., Harris, P. L., & Maratsos, M. P. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development*, 72(1), 1–129.
- Dubey, R., & Griffiths, T. L. (2017). A rational analysis of curiosity. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 307–312). Austin, TX: Cognitive Science Society.
- Foster, M. I., & Keane, M. T. (2015). Why some surprises are more surprising than others: Surprise as a metacognitive sense of explanatory difficulty. *Cognitive Psychology*, 81, 74–116.
- Gopnik, A. (2000). Explanation as orgasm and the drive for causal knowledge: The function, evolution, and phenomenology of the theory formation system. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 299–323). Cambridge, MA, US: The MIT Press.
- Grimm, S. R. (2008). Explanatory inquiry and the need for explanation. *The British Journal for the Philosophy of Science*, 59(3), 481–497.
- Hickling, A. K., & Wellman, H. M. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology*, 37(5), 668.
- Johnson, S. G. B., Johnston, A. M., Toig, A. E., & Keil, F. C. (2014). Explanatory scope informs causal strength inferences. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 707–712). Austin, TX: Cognitive Science Society.
- Keil, F. C. (2003). Folkscience: coarse interpretations of a complex reality. *Trends in Cognitive Sciences*, 7(8), 368–373.
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science*, 32(2), 259–300.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3), 449–460.
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83(1), 173–185.
- Legare, C. H., Gelman, S. A., & Wellman, H. M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning: Causal explanatory reasoning in children. *Child Development*, 81(3), 929–944.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75.
- Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, 10(10), 464–470.
- Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, 55(3), 232–257.
- Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, 20(10), 748–759.
- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, 73(4), 1073–1084.
- Nicki, R. M. (1970). The reinforcing effect of uncertainty reduction on a human operant. *Canadian Journal of Experimental Psychology*, 24(6), 389–400.
- Pacer, M., & Lombrozo, T. (2017). Ockham's razor cuts to the root: Simplicity in causal explanation. *Journal of Experimental Psychology: General*, 146(12), 1761.
- Pavlov, I. P. (1927). *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. Oxford University Press.
- Preston, J., & Epley, N. (2005). Explanations versus applications: The explanatory power of valuable beliefs. *Psychological Science*, 16(10), 826–832.
- Smock, C. D., & Holt, B. G. (1962). Children's reactions to novelty: An experimental study of "curiosity motivation." *Child Development*, 33(3), 631–642.
- Weiner, B. (1985). "Spontaneous" causal thinking. *Psychological Bulletin*, 97(1), 74–84.
- Wong, W., & Yudell, Z. (2015). A normative account of the need for explanation. *Synthese*, 192(9), 2863–2885.