

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Explain, Explore, Exploit: Effects of Explanation on Information Search

Permalink

<https://escholarship.org/uc/item/0dq3c3wf>

Journal

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

Authors

Liquin, Emily G.

Lombrozo, Tania

Publication Date

2017

Peer reviewed

Explain, Explore, Exploit: Effects of Explanation on Information Search

Emily G. Liquin (EmilyLiquin@Berkeley.Edu)

Tania Lombrozo (Lombrozo@Berkeley.Edu)

Department of Psychology, 3210 Tolman Hall
Berkeley, CA 94720 USA

Abstract

How does actively seeking explanations for one's observations affect information search over the course of learning? Generating explanations could plausibly lead learners to take advantage of the information they have already obtained, resulting in *less* exploration. Alternatively, explaining could lead learners to explore *more*, especially after encountering evidence that suggests their current beliefs are incorrect. In two experiments using a modified *observe or bet task*, we investigate these possibilities and find support for the latter: participants who are prompted to explain their observations in the course of learning tend to explore more, especially after encountering evidence that challenges a current belief.

Keywords: explanation; exploration; learning; decision making

In the decades leading up to his publication of *On the Origin of Species*, Charles Darwin recorded the titles of 687 books of English non-fiction that he read. According to analyses by Murdock, Allen, and DeDeo (2017), Darwin's reading fell into three epochs, each defined by a certain pattern of *exploration*, or broad search across new topic areas, and *exploitation*, or extended examination of texts within a similar topic area. Darwin's example raises questions about the relationship between explanation and information search. In searching for an explanation (in Darwin's case, a scientific explanation for the diversity of living things), do people pursue evidence broadly (i.e., exploring), or restrict their search to align with their current beliefs (i.e., exploiting)? Do these tendencies shift over time as new evidence is acquired? And if so, how?

Lombrozo and colleagues have proposed that when engaged in explanation, children and adults recruit explanatory considerations as evaluative constraints, rendering them more likely to generate and favor hypotheses that support "good" explanations – namely those that are simple, broad, and exhibit other explanatory virtues (Lombrozo, 2016; Williams & Lombrozo, 2010, 2013). There is also evidence that the hypotheses one generates and considers influence information search (Bonawitz, van Schijndel, Friel, & Schulz, 2012). Combining these proposals thus predicts that patterns of information search could be affected by engaging in the process of explanation.

To date, few studies have investigated the relationship between explanation and information search. In one study, Legare (2012) found that children's explanations for an unexpected piece of evidence predicted their subsequent exploratory behavior. In more recent work, Ruggeri, Lombrozo, and Xu (in prep) found that prompting children

to explain relationships in a target domain prepared them to ask more efficient questions on a subsequent 20-questions task in that domain. Neither study, however, was designed to test the causal influence of generating explanations on decisions to explore in a dynamic learning task, nor were they designed to examine adults' exploratory behavior.

In two experiments, we investigate how explanation generation affects patterns of exploration by prompting adult learners to explain as they search for information over the course of a category learning task. To accomplish this, we draw on research from the reinforcement learning literature on the *explore-exploit dilemma* (Cohen, McClure, & Yu, 2007; Kaelbling, Littman, & Moore, 1996). As defined in this literature, exploration involves seeking new information, while exploitation involves seeking reward (by taking advantage of the information one has already acquired). For example, in the *observe or bet task* (Navarro, Newell, & Schulze, 2016; Tversky & Edwards, 1966), agents must choose between "observing" which of two bulbs lights up on a given trial (without receiving any reward) or "betting" which bulb they think will light up for the chance to earn a reward (without observing that trial's result). Each bulb lights up with some fixed probability that the learner must infer through a period of observation. In this task, observation is equated with exploration (i.e., information seeking with no potential for reward) and betting is equated with exploitation (i.e., reward seeking with no potential for information).

The Present Research

In the present research, we propose a new method, the *contextual observe or bet task* (inspired by the contextual multi-armed bandit task; Langford & Zhang, 2008). In this task, a set of "context variables" (i.e., features of the options that vary across trials) can be used to predict the option that will provide a reward on each trial. Successful performance depends on learning to identify and use these variables. This method integrates a more complex, real-world learning task into an active, dynamic learning environment. We can answer the question "when do explainers choose to explore?" by measuring when learners choose to observe rather than bet.

To develop a contextual observe or bet task well suited to exploring the effects of explanation on information search, we adapt the stimuli from Williams and Lombrozo (2010). In a set of three studies, Williams and Lombrozo presented learners with four exemplars from each of two novel categories. Category members could be classified by a




Figure 1: a. Typical trial (Expts. 1 and 2): Both robots can be classified by the 100% rule (foot shape) and the 75% rule (body/head shape); b. Non-obvious anomaly trial (Expts. 1 and 2): Both robots can be classified by the 100% rule but not the 75% rule; c. Obvious anomaly trial (Expt. 2): Both robots can be classified by the 100% rule, but only one robot can be classified by the 75% rule. Category labels were only displayed if a participant chose to observe and are included here for clarity.

salient rule that accounted for 75% of exemplars or a subtle rule that accounted for 100% of exemplars. Participants who were asked to explain the category membership of each exemplar were more likely to discover the 100% rule than participants who engaged in a control task.

For our contextual observe or bet task, we present learners with *pairs* of category exemplars over a number of trials. On each trial, learners can choose to “observe” the category labels of the exemplars or “bet” which exemplar they believe belongs to a given category. Learners are thus free to determine when to seek information (exploration/observation) and when to seek reward (exploitation/betting) as they learn the features that predict category membership (context variables).

Prior work motivates two hypotheses regarding the effects of explanation generation on explore-exploit decision making. By Hypothesis 1, explaining could lead learners to greater exploitation. Previous research suggests that people use the first explanation they receive as a benchmark by which to judge subsequent explanations (Ihme & Wittwer, 2015) and use their current explanation to decide between competing hypotheses for new data (Johnson & Krems, 2001). Learners may thus prefer the first explanation they generate. This tendency towards accepting the first explanation in a series could lead people to switch to exploitation after arriving at an initial explanation, even if it is based on scant evidence. We suggest that learners may thus be more willing to quickly settle on a hypothesis that aligns with their initial beliefs based on the first pieces of information gathered, leading to increased exploitation.

By Hypothesis 2, explaining could lead learners to greater exploration. This hypothesis is consistent with one interpretation of the findings from Williams and Lombrozo (2010): when prompted to explain, participants continued to “search” the stimuli until they found a good explanation, rather than settling for the salient but imperfect 75% rule. Relatedly, Williams, Lombrozo, and Rehder (2013) found that explainers seemed to persevere in looking for a perfect classification rule, even when none was available. If explainers explore until they find a good explanation, then evidence that a candidate explanation is inadequate could be a critical cue that leads explainers to engage in further exploration. Indeed, Williams and Lombrozo (2010) found that explaining anomalies (i.e., exceptions to the 75% rule) was particularly powerful in encouraging learners to reject

an imperfect rule and discover a better alternative (see also Williams, Walker, & Lombrozo, 2012). However, this finding was not experimentally linked to an increase in exploration or information search, which makes it possible that anomalous evidence influenced discovery via other mechanisms. It is thus an open question whether explanation has a causal impact on exploration, and if so whether this impact is most pronounced when the evidence that is being explained contradicts one’s current beliefs.

For our contextual observe-or-bet task, Hypothesis 1 thus predicts that relative to control participants, those who are prompted to explain will make more “bet” choices. In contrast, Hypothesis 2 predicts that relative to control participants, those prompted to explain will be more likely to observe, especially on trials following the observation of information that is anomalous with respect to initial beliefs, which we expect to correspond to the obvious rule that accounts for 75% of exemplars. In two experiments, we test these hypotheses.

Experiment 1

Method

Participants Participants for both experiments were recruited from Amazon Mechanical Turk and paid \$0.85 for participating in the 8.5-minute study. Participation in both experiments was restricted to users in the United States with a 95% or higher approval rating based on at least 50 previous tasks. Participants in Experiment 1 were 302 adults (143 males and 159 females) ranging from 18 to 74 years of age ($M_{age} = 34$) and were randomly assigned to the *explain* condition ($N = 151$) or the *control* condition ($N = 151$). Ninety-four additional participants (44 in the *explain* condition and 50 in the *control* condition) were excluded for failing to pass two attention checks (see below).

Materials Thirty-two images of “alien robots” (see Figure 1) were designed based on the stimuli used by Williams and Lombrozo (2010). Robots varied along four dimensions: foot shape, body/head shape, left-half color, and right-half color. Twenty-two different foot shapes were used, each of which appeared on no more than two robots. All Glorps had feet that were pointy on the bottom surface, and all Drents had feet that were flat on the bottom surface. Overall, 75% of Glorps had round bodies/heads, 25% of Glorps had

square bodies/heads, 75% of Drents had square bodies/heads, and 25% of Drents had round bodies/heads. The color dimensions were irrelevant to category membership.

Foot shape (pointy/flat) was a “100% rule” that accounted for the category membership of all robots, and body/head shape (round/square) was a “75% rule” that only accounted for the category membership of 75% of the robots.

Procedure Participants were introduced to Glorp robots and Drent robots. On each of 16 trials, participants were shown a Glorp-Drent pair. Robots were paired such that no color appeared more than once in a pair, and all atypical Glorps were paired with atypical Drents. The side on which Glorps and Drents appeared was counterbalanced across trials. Pairs were presented in a random order, aside from the first four trials. For these trials, typical exemplars were presented on trials one, two, and three (“typical trials”), and atypical exemplars were presented on trial four (“anomaly trial”).

On each trial, participants were given the choice to “observe” – offering the opportunity to gain information but no reward – or “bet” – offering the opportunity to gain reward but no information. If a participant chose to observe, the participant was shown which robot from that pair was a Glorp and which was a Drent. Participants in the *explain* condition were asked to explain why the indicated robot was a Glorp robot, while participants in the *control* condition were asked to write down any thoughts they had about that trial. Participants were required to spend at least 10 seconds completing these tasks before advancing to the next trial. No points were awarded when a participant chose to observe.

If a participant chose to bet, the participant was asked to indicate which robot they thought was a Glorp. If their choice was correct, the participant would gain one point, and if their choice was incorrect, they would lose one point. However, no feedback was given on bet trials; participants were not shown their scores until the task was complete.¹

Participants were instructed to attempt to maximize their score, but also to learn how to differentiate Glorps and Drents. All participants were explicitly told that they would be asked to report any patterns that could help differentiate Glorps and Drents at the end of the task. Participants were not incentivized on the basis of their score, a point to which we return in the General Discussion.

After the 16-trial observe or bet task, participants reported any patterns they had found that differentiated Glorps and Drents and indicated what percentage of robots they thought could be accurately characterized using that pattern. Participants could report up to eight patterns. Participants also completed an attention check in which they had to distinguish between a robot they had seen during the previous task and three robots that they had not seen before. All novel robots were obviously different in appearance from Glorps and Drents. A second attention check required

¹ Participants were also prompted to report their confidence after both observe and bet trials; in the interest of space we do not report analyses of confidence here.

participants to read the instructions from the first attention check, which directed them to ignore the question that followed and instead type a specific word into the answer textbox.

Results

Rule Discovery In the *explain* condition, 17% of participants reported the 100% rule after completing the observe or bet task, while only 6% of participants in the *control* condition reported this rule. A chi-square analysis revealed that this difference was significant, $\chi^2(1) = 8.27, p = .004$. While these discovery rates seem quite low, they are not inconsistent with previous research (Williams & Lombrozo, 2010). Additionally, these results replicate Williams and Lombrozo’s (2010) finding that generating explanations promotes the discovery of broad rules.




Figure 2: Experiment 1 choices by condition and trial. Vertical line indicates first anomaly trial. Error bars: 1 SE.

Explanation and Exploration We used a generalized linear mixed effects model to predict participants’ observe/bet choices over time, with a random intercepts term to capture individual differences. The choice to observe rather than bet was significantly predicted by linear and quadratic effects of trial number, analysis of deviance: $\chi^2(1) = 212.70, p < .001$ and $\chi^2(1) = 78.04, p < .001$, respectively. Increasing trial number led to more betting at a decreasing rate over time. Condition was not a significant predictor of observe/bet choices (see Figure 2). This indicates no overall differences between the two conditions in explore/exploit decisions.

Next, we investigated the effect of explaining anomalies on subsequent exploration. For participants who observed on the first anomaly trial and thus received information contradicting the 75% rule, *explain* condition participants were more likely than *control* condition participants to continue to observe on the following trial, at a level approaching significance, $\chi^2(1) = 3.49, p = .06$. There was no condition difference in observation on the trial following the first anomaly for participants who did not observe on the anomaly trial, and thus did not receive information contradicting the 75% rule, $\chi^2(1) = 0.55, p = .46$. However, a logistic regression predicting observation following the anomaly trial by condition and observation on the anomaly

trial did not reveal a significant interaction ($b = 0.71$, $OR = 2.03$, $z = 1.24$, $p = 0.21$), likely due to the small proportion of our sample that observed the anomaly (27%).

To analyze whether the increased exploration following an observed anomaly led to increased discovery of the 100% rule, we performed a logistic regression predicting 100% rule discovery by condition and observation on the trial following the first anomaly. Both condition ($b = 1.12$, $OR = 3.05$, $z = 2.73$, $p = .006$) and post-anomaly observation ($b = 0.80$, $OR = 2.23$, $z = 2.17$, $p = .03$) were significant predictors. Thus, while explanation had an effect on 100% rule discovery above and beyond the effects of post-anomaly observation, this increase in exploration following the observation of an anomaly also predicted rule discovery.

Discussion

In Experiment 1, we found that after observing information that contradicted the 75% rule, participants who were asked to explain tended to explore more often than control participants. This exploration increased the probability of discovering a broad rule that accounted for all observations. These findings support Hypothesis 2: explanation led learners towards greater exploration after receiving evidence that current beliefs were wrong or incomplete.

These results leave open two possibilities for how explanation interacts with anomalous information. We previously suggested that explaining anomalies encourages learners to reject their current (imperfect) hypothesis, prompting subsequent exploration in the service of finding a more satisfactory alternative. On this view, explanation affects the downstream processing that *follows* the detection of an anomaly. However, there is also evidence that generating explanations can help learners articulate and recognize their current beliefs, rendering a conflict between those beliefs and subsequent information more apparent (Chi, de Leeuw, Chiu, & LaVancher, 1994). Extending these ideas, it could be that generating explanations on trials that preceded an observed anomaly helped learners recognize the anomalies as such, and that increased sensitivity to anomalies is what drove effects of explanation.

In Experiment 2, we evaluate these alternatives by introducing violations of the 75% rule that are detectable whether or not the participant chooses to observe on that trial (“obvious anomalies”). When an atypical exemplar from one category (e.g., a round Drent) is paired with a typical exemplar from the other category (e.g., a round Glorp), both robots have the same shape. Since each trial contains one Glorp and one Drent, the trial is a clear violation of the 75% rule. If explanation enhances anomaly detection by solidifying learners’ initial beliefs, we would expect participants who are prompted to explain to observe at a higher rate on the first obvious anomaly trial relative to those who are not prompted to explain. On the other hand, if *explaining* an anomaly is instead what prompts learners to reject prior beliefs and seek out better alternatives, we would expect effects of explanation to emerge only after an anomaly has been observed, and to manifest as an increase

in observation on trials *following* obvious anomalies. In Experiment 2, we test these potential accounts.

We additionally varied the point at which the first obvious anomaly was introduced – on trial 4 versus trial 8 – to test whether the timing of anomalous information affects rule discovery and/or interacts with explanation. If the power of explaining anomalous information emerges from the conflict between the novel information and prior beliefs, then introducing anomalous evidence later in the task (after beliefs have been solidified) should lead to a larger effect of explanation on exploration. If, however, explanation biases learners towards their prior beliefs (Walker, Lombrozo, Williams, Rafferty, & Gopnik, 2017; Williams & Lombrozo, 2013), increasing the strength of learners’ beliefs by increasing the amount of consistent evidence prior to introducing an anomaly could decrease the effect of anomalous evidence on subsequent exploration.

Experiment 2

Method

Participants Participants were 400 adults (192 males, 204 females, and 4 who did not specify gender) ranging from 18 to 73 years of age ($M_{age} = 34$) who were randomly assigned to the *explain* condition ($N = 203$) or the *control* condition ($N = 197$), as well as *early* anomaly timing ($N = 197$) or *late* anomaly timing ($N = 203$). One hundred fifty-three additional participants (73 in the *explain* condition and 80 in the *control* condition) were excluded for failing to pass two attention checks mirroring those used in Experiment 1.

Materials The 32 alien robot images used were identical to those used in Experiment 1.

Procedure The procedure was largely identical to Experiment 1. Three atypical Glorps were paired with atypical Drents (“non-obvious anomalies”). One atypical Glorp was paired with a typical Drent, and one atypical Drent was paired with a typical Glorp (“obvious anomalies”).

For those assigned to *early* anomaly timing, the first obvious anomaly was presented on trial 4. For *late* anomaly timing, the first obvious anomaly was presented on trial 8. The second obvious anomaly was always on trial 15. Non-obvious anomalies were randomly distributed throughout the remaining trials, excluding the first three trials.

Results

Rule Discovery Within *early* anomaly timing participants, 33% of *explain* participants and 27% of *control* participants reported the 100% rule. Within *late* anomaly timing participants, 30% of *explain* participants and 14% of control participants reported the 100% rule. A logistic regression analysis revealed that participants in the *explain* condition were significantly more likely than participants in the *control* condition to discover the 100% rule ($b = 0.60$, $OR = 1.81$, $z = 2.55$, $p = .01$). Participants with *late* anomaly timing were somewhat less likely than participants with

early anomaly timing to discover the 100% rule ($b = -0.43$, $OR = 0.65$, $z = -1.84$, $p = .07$).

Explanation and Exploration We used a generalized linear mixed effects model to predict participants' observe/bet choices over time, with a random intercepts term to capture individual differences. The choice to observe rather than bet was significantly predicted by condition (*explain* vs. *control*) and linear and quadratic effects of trial number, analysis of deviance: $\chi^2(1) = 4.47$, $p = .03$; $\chi^2(1) = 273.93$, $p < .001$; and $\chi^2(1) = 150.40$, $p < .001$, respectively. Anomaly timing (*early* vs. *late*) was not a significant predictor of observation. *Explain* participants were more likely to observe than *control* participants, and increasing trial number increased the likelihood of betting at a decreasing rate over time (see Figure 3).

Next, we analyzed exploration on the first obvious anomaly trial by performing a logistic regression with task (*explain* vs. *control*) and anomaly timing (*early* vs. *late*) as predictors. Participants with *late* anomaly timing were 56% less likely to observe on the first anomaly trial relative to participants with *early* anomaly timing ($b = -0.82$, $OR = 0.44$, $z = -2.81$, $p = .005$), indicating that fewer participants observed the first obvious anomaly when it was presented later in the task. Condition was not a significant predictor of anomaly observation, nor was the interaction between condition and anomaly timing. These findings suggest that explanation did not exert effects on discovery by increasing the rate at which obvious anomalies were detected.

To analyze exploration following an anomalous observation, we performed a logistic regression predicting observation on the trial following the first obvious anomaly, with condition (*explain* vs. *control*) and anomaly timing (*early* vs. *late*) as predictors. This revealed a marginally significant interaction between task and anomaly condition ($b = 0.88$, $OR = 2.40$, $z = 1.78$, $p = .07$). For *late* anomaly timing, *explain* participants were more likely than *control* participants to observe on the trial following the first obvious anomaly, $\chi^2(1) = 6.85$, $p = .009$. This difference was not significant for *early* anomaly timing, $\chi^2(1) = 0.01$, $p = .92$. These findings support the idea that explanation affects learning by increasing exploration in the face of anomalous evidence; they also challenge the idea that effects of explanation are restricted to anomaly *detection*. Explainers were no more likely to choose to observe an obvious anomaly, but were more likely (in the *late* anomaly condition) to follow up with additional observation.

To analyze whether this increased exploration following an observed anomaly led to increased discovery of the 100% rule, as well as whether condition had an effect on rule discovery above and beyond the effects of such exploration, we performed a logistic regression predicting 100% rule discovery by condition and observation on the trial following the obvious anomaly. Both condition ($b = 0.49$, $OR = 1.64$, $z = 2.10$, $p = .04$) and observation following the first obvious anomaly ($b = 0.79$, $OR = 2.20$, $z = 3.12$, $p = .002$) were significant predictors of 100% rule discovery.

Discussion

These results suggest that explanation generation leads to greater exploration after observing evidence that challenges a current hypothesis. This difference in exploratory behavior does not depend on simply noticing the presence of contradictory information, but instead depends specifically on one's attempts to explain this anomalous information.

We also found a suggestive difference between *early* and *late* anomaly timing. Further research is clearly warranted, but the effect of explaining an obvious anomaly may be more powerful as the degree of conflict between the anomaly and one's current beliefs is increased.




Figure 3: Experiment 2 choices by condition, anomaly timing, and trial. Vertical lines indicate obvious anomaly trials. Error bars: 1 SE.

General Discussion

In two experiments, we investigated how explanation generation affects exploration over the course of a category learning task. Lombrozo and colleagues (Lombrozo, 2016; Williams & Lombrozo, 2010, 2013) have proposed that generating explanations recruits a set of inductive constraints on hypothesis generation and selection, which can lead to the discovery of broad, simple, and generalizable rules and patterns. In the present research, we extend this account, suggesting that this learning outcome is partially dependent upon generating explanations for anomalous observations, which increases a learner's propensity to seek additional evidence.

Our results are consistent with Legare's (2012) finding that children's explanations for surprising events predicted their exploratory behavior. In the present research, however, we establish a causal link between explanation and exploration, and demonstrate that this link holds not only for children's causal learning (as proposed by Legare), but also for adults' category learning.

That said, many open questions remain. For example,

might explanation affect explore/exploit decisions by shifting participants' confidence on each trial (e.g., Auer, 2002)? Does explaining affect motivation, which could also be achieved by incentivizing reward? Equally important is identifying boundary conditions on our effects: are there situations in which explaining anomalies could lead learners to *explain them away* (Chinn & Brewer, 1998), and thus engage in greater exploitation?

While some of the effects reported here failed to reach statistical significance, we did find similar results across two experiments. Unfortunately, the effect sizes are limited by the small proportion of participants who are able to discover the 100% rule. Future work might benefit from more sensitive paradigms. Additionally, the paradigm used here allowed participants to gain some information on each trial without engaging in exploration. Since exemplars from one category were always presented with exemplars from the other category, participants could identify the features that differed between the two categories without choosing to observe. Future work will limit potential learning to observation trials in order to better isolate the effects of explanation on information search.

In sum, our findings suggest that attempting to explain observations that are anomalous with respect to one's current beliefs encourages further exploration. This may be one mechanism by which generating explanations affects learning, and provides compelling evidence that top-down constraints on hypothesis generation and selection affect not only the conclusions that learners draw, but also the ways in which they seek information – whether that information comes from 19th century texts or a robot classification task.

Acknowledgments

This work was supported by a McDonnell Foundation award to TL. EL is also the recipient of a National Science Foundation Graduate Research Fellowship. 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

- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3, 397–422.
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64(4), 215–234.
- Chi, M. T., de Leeuw, N., Chiu, M., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18, 439–477.
- Chinn, C. A., & Brewer, W. F. (1998). An empirical test of a taxonomy of responses to anomalous data in science. *Journal of Research in Science Teaching*, 35(6), 623–654.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942.
- Ohme, N., & Wittwer, J. (2015). The role of consistency, order, and structure in evaluating and comprehending competing scientific explanations. *Instructional Science*, 43(4), 507–526.
- Johnson, T. R., & Krems, J. F. (2001). Use of current explanations in multicausal abductive reasoning. *Cognitive Science*, 25(6), 903–939.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
- Langford, J., & Zhang, T. (2008). The epoch-greedy algorithm for multi-armed bandits with side information. *Advances in neural information processing systems 20* (pp. 817–824).
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83(1), 173–185.
- Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, 20(10), 748–759.
- Murdock, J., Allen, C., & DeDeo, S. (2017). Exploration and exploitation of Victorian science in Darwin's reading notebooks. *Cognition*, 159, 117–126.
- Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore-exploit dilemma in static and dynamic environments. *Cognitive Psychology*, 85, 43–77.
- Ruggeri, A., Lombrozo, T., & Xu, F. (in prep). Explaining prepares preschoolers to ask efficient questions.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71(5), 680–683.
- Walker, C. M., Lombrozo, T., Williams, J. J., Rafferty, A. N., & Gopnik, A. (2017). Explaining constrains causal learning in childhood. *Child Development*, 88(1), 229–246.
- Williams, J. J., & Lombrozo, T. (2010). The role of explanation in discovery and generalization: Evidence from category learning. *Cognitive Science*, 34(5), 776–806.
- Williams, J. J., & Lombrozo, T. (2013). Explanation and prior knowledge interact to guide learning. *Cognitive Psychology*, 66(1), 55–84.
- Williams, J. J., Lombrozo, T., & Rehder, B. (2013). The hazards of explanation: Overgeneralization in the face of exceptions. *Journal of Experimental Psychology: General*, 142(4), 1006–1014.
- Williams, J. J., Walker, C. M., & Lombrozo, T. (2012). Explaining increases belief revision in the face of (many) anomalies. *Proceedings of the 34th annual conference of the cognitive science society* (pp. 1149–1154). Austin, TX: Cognitive Science Society.