Explaining increases belief revision in the face of (many) anomalies

Joseph Jay Williams (joseph_williams@berkeley.edu)
Caren M. Walker (caren.walker@berkeley.edu)
Tania Lombrozo (lombrozo@berkeley.edu)
Department of Psychology, 3410 Tolman Hall
Berkeley, CA, 94709

Abstract

How does explaining novel observations influence the extent to which learners revise beliefs in the face of anomalies – observations inconsistent with their beliefs? On one hand, explaining could recruit prior beliefs and *reduce* belief revision if learners "explain away" or discount anomalies. On the other hand, explaining could *promote* belief revision by encouraging learners to modify beliefs to better accommodate anomalies. We explore these possibilities in a statistical judgment task in which participants learned to rank students' performance across courses by observing sample rankings. We manipulated whether participants were prompted to explain the rankings or to share their thoughts about them during study, and also the proportion of observations that were anomalous with respect to intuitive statistical misconceptions. Explaining promoted greater belief revision when anomalies were common, but had no effect when rare. In contrast, increasing the number of anomalies had no effect on belief revision without prompts to explain.

Keywords: explanation, self-explanation, learning, generalization, statistics, misconceptions, anomalies.

Introduction

Human learning relies on the ability to use novel observations about the world to revise current beliefs. This raises basic questions about how observations impact beliefs, and in particular how different cognitive processes influence the process of belief revision. The current paper examines how *explaining* observations that are anomalous with respect to a learners' current beliefs influences the nature of belief revision.

Previous research reveals that seeking and generating explanations can play an important role in learning and reasoning across a variety of task and domains, including learning novel categories (Williams & Lombrozo, 2010), inferring causal relationships (Koslowski, 1996), and generalizing the properties of people and objects from known to unknown cases (Rehder, 2009; Sloman, 1994). In addition, generating explanations can drive conceptual development in children (Wellman & Liu, 2006) and has been shown to have important pedagogical benefits in a variety of educational contexts (e.g. Fonseca & Chi, 2011).

Despite widespread appreciation for the impact of explaining on learning and belief revision, current accounts of explanation's effects pose a challenging puzzle. On the one hand, explaining is frequently hypothesized to encourage learners to accommodate novel observations in the context of their prior beliefs (Ahn, Brewer & Mooney, 1992; Chi et al, 1994; Lombrozo, 2006; Walker, Williams, Lombrozo & Gopnik, under review; Williams & Lombrozo, 2010). This suggests that explaining could lead learners to draw on a range of belief-preserving strategies in the face of anomalous observations (Chinn & Brewer, 1993; Kuhn, 1962; Koslowski, 1996; Lord, Lepper, & Ross, 1979). Learners who seek explanations could engage in less belief revision (relative to those who do not explain) if they "explain away" anomalies by discounting them as implausible, or reinterpret observations as consistent with or irrelevant to preferred theories.

On the other hand, explaining anomalies could prompt learners to reject their currently-held beliefs and construct new theories that better accommodate the anomalous observations. For example, explaining could ensure that anomalies are not simply ignored, increasing learners' allocation of attention and processing time to these unanticipated observations (Legare, 2010; Siegler, 2002). In addition, explaining could encourage learners to seek and discover general patterns that go beyond prior beliefs to capture the anomalies being explained (Walker, Williams, Lombrozo & Gopnik, under review; Williams & Lombrozo, 2010).

A third possibility is that explaining has the potential to produce both effects. Whether explaining preserves or revises beliefs could depend on the nature of the evidence provided by the anomalies. For example, relative to other learning strategies, explaining could encourage learners to discount anomalous observations if they are infrequent or there is no alternative theory available to explain them. However, seeking explanations for more extensive inconsistencies could emphasize the limitations of current beliefs and guide learners to alternative theories, thus promoting more radical belief revision.

We investigate these possibilities in the context of a statistical judgment task in which participants learn a system for ranking students' performance across different courses. Participants learn to rank by observing sample rankings, and they are either prompted to provide explanations for the rankings or to engage in free study as a control condition. Previous research suggests that participants' prior beliefs will favor rankings on the basis of statistically problematic principles rather than one that is normatively defensible (e.g., Schwartz & Martin, 2003; Belenky & Nokes-Malach, in press). We can therefore manipulate how many of the sample rankings happen to be consistent with the non-

normative principles, and how many are anomalous and only accounted for by the correct principle. This allows us to examine how the process of belief revision is influenced by explaining the novel observations (the sample rankings), by the number of anomalous observations, and by any interaction between these factors.

Experiment

Participants were instructed to learn how a university ranks their students across different courses by studying five examples of ranked pairs of students. Each example reported which of two students in two different courses was ranked higher, listing each student's (percentage) grade along with each course's mean grade, average deviation (the average absolute deviation), and minimum and maximum grade. We used average deviation as a measure of variability instead of standard deviation to follow past research in avoiding the need for formulas and using a concept more transparent to participants (Schwartz & Martin, 2003; Belenky & Nokes-Malach, in press).

All five examples were ranked according to a *relative-to-deviation* principle: the better student was the one that scored a greater number of average deviations above the mean (see Schwartz & Martin, 2003; Belenky & Nokes-Malach, in press). However, some of the examples were also consistent with three non-normative principles (e.g. the student with higher absolute scores is always ranked higher). These are described further in the *Materials & Procedure* section, below.

To probe how anomalous observations influenced belief revision, we manipulated how many of the five ranked examples were consistent (or anomalous) with respect to the non-normative principles. The relative-to-deviation principle always accounted for all five examples. In the *single anomaly* condition, four of the five example rankings also conformed to the non-normative principles, and there was just one anomaly that was inconsistent with them. In the *multiple anomalies* condition, there were four anomalies and only one of the five rankings was consistent with the non-normative principles.

To examine how explaining interacted with anomalies to impact learning, we also manipulated the extent to which participants engaged in explanation. In the *explain* condition we prompted participants to explain each ranked example. In the *free study* control condition, participants were free to use any study strategy, but were prompted to articulate their thoughts while studying each ranked example. Like explaining, this control condition involved paying attention to the details of the cases and articulating one's thinking in language.

As discussed in the introduction, engaging in explanation and varying the number of anomalies could impact belief revision in several ways. Explaining the anomalies could reduce the revision of prior beliefs and inhibit learning about the relative-to-deviation principle. This effect could be especially potent when there is only a single anomaly to the non-normative principles. On the other hand, explaining

could magnify the effects of anomalies in rejecting belief in the non-normative principles and instead encourage learners to induce and adopt the relative-to-deviation principle.

Finally, the effects of explaining could depend on whether the explained observations include a single anomaly or many. Explaining could have a large impact on belief revision in the context of multiple anomalies, but have no effect or even inhibit belief revision when only a single (and easily discounted) anomaly is present. Alternatively, explaining could boost the impact of a single anomaly that might otherwise be ignored, but have no effect (relative to control) when there are multiple anomalies that make the need for belief revision completely apparent. The design of our experiment allows us to differentiate these possibilities.

Methods

Participants

Participants were 275 adults recruited online through the Amazon Mechanical Turk marketplace and reimbursed for their time.[1]

Materials & Procedure

The materials consisted of five examples of student pairs ranked by the university, ten unranked pre-test pairs, and ten unranked post-test pairs. The experiment involved introduction, pre-test, study, and post-test phases.

Introduction Participants were informed that they would observe pairs of students from different classes whose academic performance had been ranked by the university, and that their goal was to learn the ranking system employed. They were given the definition of "average" – the sum of all scores divided by the number of students in a class – and "average deviation" – the sum of all the (absolute) differences between student scores and the average, divided by the number of students.

Ranked examples for study During study, five examples of ranked student pairs were presented. A ranked example (see Figure 1a and 1b) stated which student was ranked higher by the university, and reported each student's: (1) name (e.g., Sarah); (2) class (e.g., Sociology); (3) class's mean score (e.g., 79%); (4) class's average deviation (e.g., 8%); (5) class's minimum score (e.g., 67%); and (6) class's maximum score (e.g., 90%).

Principles for ranking students Participants could interpret or predict the rank of each student pair using at least four principles. The three *non-normative* principles were incorrect but designed to correspond to intuitive statistical misconceptions.

---

[1] We included a question that assessed whether participants were actually reading instructions (Oppenheimer et al, 2009). The pattern of results was the same if participants who did not pass this test were excluded.

(a) Sarah got 85% in a Sociology class, where the average score was 79%, the average deviation was 8%, the minimum score was 67%, and the maximum score was 90%.

Tom got 69% in a Art History class, where the average score was 65%, the average deviation was 3%, the minimum score was 42%, and the maximum score was 87%.

Sarah was ranked more highly by the university than Tom.

(b) Sarah got 85% in a Sociology class, where the average score was 79%, the average deviation was _**3%**_, the minimum score was 67%, and the maximum score was 90%.

Tom got 69% in a Art History class, where the average score was 65%, the average deviation was _**8%**_, the minimum score was 42%, and the maximum score was 87%.

**Tom** was ranked more highly by the university than _**Sarah**_.

Figure 1: **(a)** A *consistent* ranked example for which all four principles predicted the same ranking. **(b)** An *anomalous* ranked example constructed by switching the class average deviations of the consistent example from Figure 1a. The switch means that the correct relative-to-deviation ranking is now the opposite of what is predicted by the raw-score, relative-to-average, and relative-to-range principles. Emphasis is added for illustration and was not provided to participants.

We term the principles (1) *raw-score*: the higher ranking went to the student with the higher score, irrespective of mean, average deviation, and minimum or maximum score; (2) *relative-to-average*: the higher ranking went to the student whose score was the farthest above (or least below) the class's mean score; (3) *relative-to-range*: the higher ranking went to the student whose score was farther from the average *relative to the range* in class scores, where this was calculated as the difference from the mean divided by the range.

The *relative-to-range* principle privileges the score that is farther from the mean as measured in "range-units," capturing some notion of variability (when range is correlated with variability), and could be approximated by looking at a score's distance from the maximum score.

The fourth and more accurate *relative-to-deviation* principle favored whichever score was a greater number of average deviations above the mean. This was calculated as the difference from the mean divided by the average deviation, and is closely related to normative measures such as the standard deviation and z-score, indicating the person's score relative to the distribution of scores in the class.

Consistent vs. anomalous examples All five ranked examples conformed to the relative-to-deviation principle. However, a ranked study example could be *consistent* with or *anomalous* with respect to the ranking given by the raw-score, relative-to-average, and relative-to-range principles, all of which always generated identical rankings on study examples (see Figure 1a and 1b). Five consistent examples were constructed so that each could be converted to an anomalous example by switching the average deviation of the two students' classes. This permitted a close match between consistent and anomalous examples on all other dimensions (compare Figure 1b to Figure 1a).

In the *single anomaly* condition there were four consistent examples and one anomalous example. The *multiple anomalies* condition had the opposite ratio: one consistent example and four anomalous examples.

Pre-test To provide a baseline measure of belief before study, participants were presented with ten unranked student pairs. They judged which student the university would rank higher, and rated confidence in their judgment on a scale from 1 ("not at all") to 7 ("extremely").

The ten student pairs were designed to identify the principle(s) that participants used to rank students, and thus pitted candidate principles against each other. Specifically, there were two instances of each of the following types of pairs, pitting (1) the relative-to-deviation principle against the three non-normative principles (like anomalous study examples); (2) the raw-score principle against the other three principles; (3) the relative-to-average principle against the other three principles; (4) the relative-to-range principle against the other three principles; and (5) the two principles that were most sensitive to variability, relative-to-range and relative-to-deviation, against the raw-score and relative-to-average principles.

Study Each of the five ranked examples was presented onscreen for exactly 90 seconds in a format similar to Figure 1a and 1b. Participants in the *explain* condition were prompted to explain why the higher-ranked student was ranked more highly, typing their explanation into a text box onscreen. Participants in the *free study* control condition were told to type their thoughts during study into an equivalent text box.

Post-test To assess belief after study, participants' ranking judgments and confidence ratings were solicited for ten unranked student pairs. All names and grades were changed from the pairs used in pre-test, but five points were added to each grade to generate novel numbers while preserving the way in which the items pitted the principles against each other.[2]

---

[2] Additional questions were asked at the end of the experiment (e.g. demographics, sufficient time for task, strategy) but are not further discussed here in the interest of space.

Results

Overall pre- and post-test accuracy Learning was assessed by comparing accuracy on the pre-test and post-test items. Correct responses were considered to be those that were consistent with the relative-to-deviation principle. Figure 2 reports an overall measure of accuracy across all pre-test and post-test items as a function of learning task and number of anomalies. Accuracy improved from pre- to post-test: A 2 (task: explain vs. free study) x 2 (number of anomalies: single vs. multiple) x 2 (test: pre-test vs. post-test) repeated measures ANOVA found a main effect of test (pre- vs. post-) on overall accuracy, $F(1, 269) = 4.33$, $p < 0.05$.

The ANOVA additionally revealed interactions between test and learning task, $F(1,269) = 4.95$, $p < 0.05$, and between test and number of anomalies, $F(1,269) = 3.88$, $p < 0.05$. These effects were driven by a greater boost in pre- to post-test accuracy for participants in the *explain – multiple anomalies* conditions, which in turn was driven primarily by rankings on anomalous items, as discussed further below.

Anomalous items: Change in pre- to post-test accuracy Figure 3 reports accuracy on *anomalous* pre-test and post-test items. These items were analogous to the anomalies at study in pitting the relative-to-deviation principle against all three non-normative principles. Accuracy on these items was critical to testing our hypotheses about the effects of explanation and anomalies on the revision of beliefs to favor the relative-to-deviation principle over the non-normative alternatives. Accuracy on anomalous items improved from pre- to post-test: A 2 (task: explain vs. free study) x 2 (number of anomalies: single vs. multiple) x 2 (test: pre-test vs. post-test) ANOVA on accuracy on anomalous items revealed a main effect of pre- vs. post- test, $F(1, 269) = 63.85$, $p < 0.05$.[3]

Subsequent analyses directly examined the pre-test to post-test *change* in accuracy on the anomalous items. Figure 4 reports performance on this measure, calculated as post-test accuracy minus pre-test accuracy.

A 2 (task: explain vs. free study) x 2 (anomalies: single vs. multiple) ANOVA on the pre- to post-test change in accuracy on anomalous items revealed a significant effect of number of anomalies, $F(1, 266) = 12.8$, $p < 0.05$. This main effect was qualified by an interaction between learning task and number of anomalies, $F(1, 266) = 7.77$, $p < 0.05$. No other effects were significant.
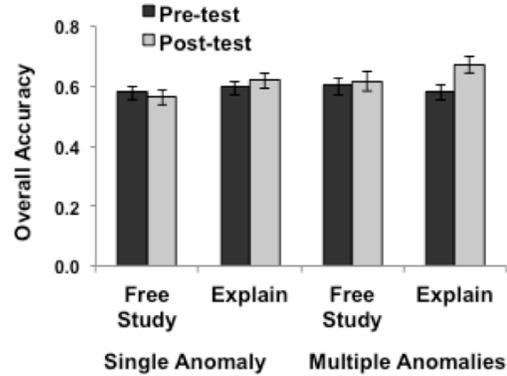

Figure 2: Accuracy on *all* pre-test and post-test items, by learning task and number of anomalies.
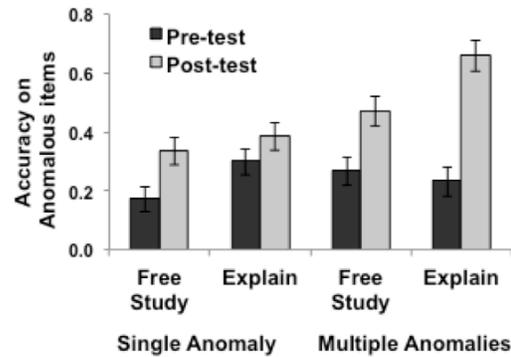

Figure 3: Accuracy on *anomalous* pre-test and post-test items, by learning task and number of anomalies.
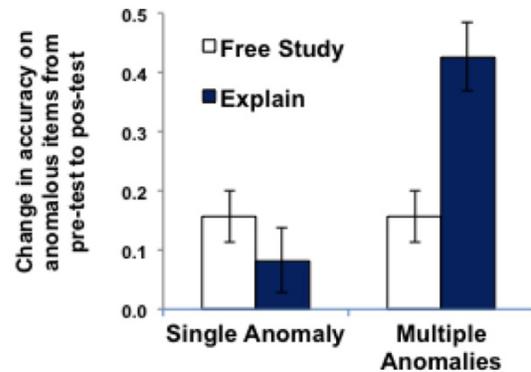

Figure 4: Change from pre-test to post-test in accuracy on *anomalous* items, by learning task and number of anomalies.

When there were multiple anomalies, participants prompted to explain showed greater learning (relative to free study) about the relative-to-deviation principle, $t(126) = 2.44$, $p < 0.05$. But when only a single anomaly was present, there was no significant effect of explanation, $t(143) = 1.13$, $p = 0.26$.

---

[3] The ANOVA on accuracy for anomalous items revealed a number of additional effects, the relevance of which is more readily communicated in our subsequent analyses on the pre- to post- test *change* in accuracy. These include a two-way interaction between test and number of anomalies, $F(1,266) = 12.80$, $p < 0.001$, a three-way interaction between test, learning task and number of anomalies, $F(1,266) = 7.77$, $p < 0.01$, and main effects of task, $F(1,266) = 4.80$, $p < 0.05$, and number of anomalies, $F(1,266) = 8.45$, $p < 0.005$.

In sum, although explaining helped participants learn the challenging relative-to-deviation principle over more intuitive alternatives, this benefit only exceeded that observed in the control condition when many anomalies were explained. It could be that the effects of explaining one anomaly were too small to yield a statistically significant effect. But a more intriguing possibility – suggested by the (non-significant) trend for control participants to show greater learning gains than explain participants in the single anomaly condition (see Figure 4) – is that explaining actually hindered belief revision by encouraging participants to discount the single anomalous observation.

It is worth emphasizing that the learning benefits observed in the *explain – many anomalies* condition cannot be attributed simply to the effects of receiving more anomalies. Receiving multiple anomalies (relative to observing a single anomaly) promoted greater learning when participants were prompted to explain, $t(135) = 2.24$, $p < 0.05$. However, in the free study control condition, observing a greater number of anomalies did not produce a significant learning benefit, $t(134) = 0.55$, $p = 0.58$. Without explaining, the potential learning benefits of anomalous information may not be realized.

Discussion

In an experiment involving statistical judgments, we found that participants who were prompted to explain observations engaged in greater belief revision than participants who engaged in a control task matched for time, attention, and use of language. Specifically, participants who explained showed a greater increase from pre- to post-test in their use of a normative principle for ranking students' performance across courses (the relative-to-deviation principle). However, this benefit was only observed when the observations that participants explained involved multiple anomalies – observations inconsistent with non-normative principles that were arguably more intuitive and more consistent with prior beliefs. When a single anomaly was presented, participants who explained showed comparable or (nonsignificantly) less learning than those in the control condition.

In the introduction we presented several plausible hypotheses about the effects of explaining anomalies. One possibility was that explaining anomalous observations would lead learners to consult current beliefs in making sense of the unanticipated observation, and therefore be more likely to preserve their current beliefs by somehow explaining away or discounting the anomaly (Chinn & Brewer, 1993). In the present experiment, for example, participants could have invoked a clerical error to explain an unanticipated observation, or generated reasons for a ranking that went beyond the information provided (e.g., perhaps a given course was especially difficult and therefore taken by students who were already high achievers).

If explaining encouraged participants to engage in such belief-preserving strategies independently of the number of anomalous observations, then pre- to post-test performance should have increased more for participants in the control condition than for those in the explain condition. While there was a trend in this direction when a single anomaly was presented, the findings do not provide clear support for this hypothesis.

Explanation-induced failures to revise belief in light of anomalies could be more likely in contexts where participants hold stronger prior beliefs. In fact, Williams, Lombrozo, and Rehder (2010) report an "explanation impairment effect" along these lines. It should be noted, however, that under some conditions maintaining current beliefs in the face of anomalous observations could be the correct or rational strategy. For example, when observations are erroneous or generated by probabilistic processes it could be preferable for learners to discount anomalous observations on the basis of (more accurate) prior beliefs.

A second hypothesis that we considered at the outset was that explaining anomalous observations would increase belief revision, perhaps by drawing attention to anomalies or forcing participants to identify patterns that would account for the anomalous observations and past observations in a unified way. While we found support for this hypothesis when many anomalies were presented, explaining did not have a measurable advantage when participants only observed a single anomaly.

Our findings are therefore consistent with a third possibility: that the effects of explanation interact with the number of anomalous observations. The *number* of anomalies per se may not be crucial, but rather serve (in the current experiments) as an indication of the strength of the evidence that current beliefs require revision. It could be that explaining *few* anomalous observations has no effect (relative to control), encourages belief-preserving strategies, or has variable effects across participants, while explaining multiple (or more problematic) anomalies more uniformly increases the extent to which participants revise beliefs to achieve consistency with observations.

It's also noteworthy that in the absence of explanation, encountering additional anomalies was insufficient to increase belief revision: In the free study condition, observing multiple anomalies (80% of observations) did not yield any significant learning benefit beyond observing just one. Chinn and Brewer (1993) point out that anomalies do not always lead to changes in belief given the number of belief-preserving strategies available to learners, and our findings are consistent with this observation. Explaining could therefore be especially valuable as a strategy for ensuring that learners benefit from anomalous observations, especially in pedagogical contexts in which anomalies are likely to highlight misconceptions and point to normative alternatives.

In previous work we have proposed a *subsumptive constraints* account of the effects of explanation on learning (Williams & Lombrozo, 2010), and we interpret the current results as broadly consistent with this proposal. According to this account, explaining does not provide a general boost to processing, but rather exerts a selective constraint to

interpret what is being explained as an instance of a broader pattern or generalization. One substantiated prediction of this account is that explaining guides people towards patterns that apply to more observations – that is, those that render fewer observations anomalous (Williams & Lombrozo, 2010). A second is that explaining increases learners' consultation of prior beliefs to privilege patterns that prior beliefs suggest will generalize to other contexts (Walker et al., under review; Williams & Lombrozo, 2010b; Williams & Lombrozo, under review).

In the current experiment, the correct relative-to-deviation principle involved fewer (zero) anomalies, but the alternative principles were more consistent with most people's prior beliefs concerning ranking. Explaining could therefore have favored the relative-to-deviation principle in the multiple anomaly condition because the evidence indicated that current beliefs were problematic. In contrast, participants in the single anomaly condition – depending on the strength of their prior beliefs – could have been inclined to favor current beliefs or to consider revision in the face of weak evidence for an alternative. These observations raise a number of important questions for future research concerning the precise conditions under which explaining anomalies will revise or entrench current beliefs.

References

Ahn, W., Brewer, W. F., & Mooney, R. J. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory*, & Cognition. *18*(2), 391-412.

Belenky, D. M., & Nokes-Malach, T. J. (in press). Motivation and transfer: The role of mastery-approach goals in preparation for future learning. *The Journal of the Learning Sciences*.

Chinn, C. A., & Brewer, W. F. (1993). The role of anomalous data in knowledge acquisition: A theoretical framework and implications for science education. *Review of Educational Research*, 63, 1-49.

Chi, M.T.H., de Leeuw, N., Chiu, M.H., LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science, 18*, 439-477.

Fonseca, B. & Chi, M.T.H. (2011). The self-explanation effect: A constructive learning activity. In Mayer, R. & Alexander, P. (Eds.), *The Handbook of Research on Learning and Instruction* (pp. 270-321). New York, USA: Routledge Press.

Legare, C.H., Gelman, S.A., & Wellman, H.M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development*, 81, 92-944.

Koslowski, B. (1996). *Theory and evidence: The development of scientific reasoning*. The MIT Press.

Kuhn, T. (1962). *The structure of scientific revolutions* (3rd ed). Chicago, IL: University of Chicago Press.

Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences*, 10, 464-470.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37(11)*, 2098-2109.

Oppenheimer, D. M., Meyvisb, T., & Davidenkoc, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45(4)*, 867-872.

Renkl, A. (1997). Learning from Worked-Out Examples: A Study on Individual Differences. *Cognitive Science*, 21(1), 1–29.

Schwartz, D.L., & Martin, T. (2004). Inventing to Prepare for Future Learning: The Hidden Efficiency of Encouraging Original Student Production in Statistics Instruction. *Cognition and Instruction*, 22(2), 129-184.

Siegler, R. S. (2002). Microgenetic studies of self-explanations. In N. Granott & J. Parziale (Eds.), *Microdevelopment: Transition processes in development and learning* (pp. 31-58). New York: Cambridge University.

Sloman, S. A. (1994). When explanations compete: The role of explanatory coherence on judgments of likelihood. *Cognition*, *52*(1), 1–21.

Wellman, H. M., & Liu, D. (2007). Causal reasoning as informed by the early development of explanations. *Causal learning: psychology, philosophy, and computation*, 261–279.

Walker, C.M., Williams, J.J., Lombrozo, T., & Gopnik, A. The role of explanation in children's causal learning. Manuscript under review.

Williams, J.J. & Lombrozo, T. (2010). The role of explanation in discovery and generalization: evidence from category learning. *Cognitive Science*, 34, 776-806.

Williams, J.J., & Lombrozo, T. (2010). Explanation constrains learning, and prior knowledge constrains explanation. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2912-2917). Austin, TX: Cognitive Science Society.

Williams, J.J., Lombrozo, T., & Rehder, B. (2010). Why does explaining help learning? Insight from an explanation impairment effect. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2906-2911). Austin, TX: Cognitive Science Society.