Probabilistic Versus Heuristic Accounts of Explanation in Children: Evidence from a Latent Scope Bias

Angie M. Johnston[*1], Samuel G. B. Johnson[*1], Marissa L. Koven[2], & Frank C. Keil[1]
(angie.johnston@yale.edu, samuel.johnson@yale.edu, mlkoven@gmail.com, frank.keil@yale.edu)
[1]Department of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520 USA
[2]Department of Psychology, Emory University, 36 Eagle Row, Atlanta, GA 30322 USA
***The first two authors contributed equally to this work; authorship order was determined by Settlers of Catan**

Abstract

Like scientists, children must find ways to explain causal systems in the world. The Bayesian approach to cognitive development holds that children evaluate explanations by applying a normative set of statistical learning and hypothesis-testing mechanisms to the evidence they observe. Here, we argue for certain supplements to this approach. In particular, we demonstrate in two studies that children, like adults, have a robust *latent scope bias* that conflicts with the laws of probability. When faced with two explanations equally consistent with observed data, where one explanation made an unverified prediction, children consistently preferred the explanation that did *not* make this prediction (Experiment 1). The bias can be overridden by strong prior odds, indicating that children can integrate cues from multiple sources of evidence (Experiment 2). We argue that children, like adults, rely on heuristics for making explanatory judgments which often lead to normative responses, but can lead to systematic error.

Keywords: Cognitive development; causal reasoning; explanation; evidence; probability; philosophy of science.

Beauty is truth, truth beauty,—that is all
Ye know on earth, and all ye need to know.
-John Keats, "Ode on a Grecian Urn" (1819)

Introduction

Children are often characterized as budding scientists. In the first years of life, young children perform inductive feats befitting of a Newton or a Darwin, managing to learn the vocabulary and grammar of one or more natural languages, to carve the world up into useful categories, and to increase their understanding of the causal structure of the physical and social worlds. These accomplishments are all the more remarkable because, unlike mature scientists, children must induce this knowledge without the benefit of formal education or scientific training.

If children truly approach the world like little scientists, gathering evidence and inferring regularities, then perhaps their inferential practices are also similar to those of actual scientists. In order for scientists to make sense out of the world, they must perform *abduction*—inferring the best explanation for a given set of observations (Lipton, 2004). However, within philosophy of science, there is considerable disagreement about what criteria scientists use for evaluating hypotheses or explanations. According to Bayesian confirmation theory (e.g., Jeffrey, 1965), scientists are concerned with inferring the *likeliest* explanation—the hypothesis that has maximum posterior probability after observing the evidence. On the reasonable assumption that seeking the truth requires us to seek the most probable explanation, scientists certainly seem to aspire to this goal.

However, scientists may not always *directly* consider which explanations are most likely, but may instead search for the *loveliest* explanation, in the hope that their instinctual sense of explanatory virtue can be a guide to truth. There is much anecdotal support for the importance of explanatory elegance to the work of scientists. For example, Hermann Bondi describes his experience meeting Albert Einstein (quoted in Zee, 1999):

> What I remember most clearly was that when I put down a suggestion that seemed to me cogent and reasonable, Einstein did not in the least contest this, but he only said, "Oh, how ugly." As soon as an equation seemed to him to be ugly, he really rather lost interest in it and could not understand why somebody else was willing to spend much time on it. He was quite convinced that beauty was a guiding principle in the search for important results in theoretical physics.

Numerous other great scientists and mathematicians have echoed Keats' refrain, that "beauty is truth, truth beauty."

Psychologically, we can think of the "likeliness" strategy as using normative probability theory to evaluate explanations, selecting the explanation that is made most probable by the evidence (Pearl, 1988). In contrast, we can think of the "loveliness" strategy as a heuristic strategy, selecting the explanation that scores highest on a set of "explanatory virtues" such as simplicity, scope, and generality (Lipton, 2004; McGrew, 2003). This more heuristic view, while potentially leading to error at times, has the advantage of being computationally much simpler—it does not require explicit calculations of prior probabilities, likelihoods, or posteriors.

Unfortunately, these views are difficult to disambiguate, because lovely explanations often *are* likelier. Take simplicity, for instance. Adults (Lombrozo, 2007), children (Bonawitz & Lombrozo, 2012), and scientists (Zee, 1999) all prefer explanations invoking fewer causes to explanations invoking more causes. Yet, it is unclear whether this simplicity preference is driven by an attempt to maximize the *likelihood* that an explanation is true, or by a more course-grained heuristic (e.g., "simpler is better") that often approximates normative inferences. All else being equal, simpler explanations have higher prior probabilities than more complex explanations. If Disease A can explain all of a patient's symptoms but Diseases B and C would need to act in conjunction to account for her symptoms, then Disease A is the best explanation not merely because it is simpler or lovelier, but also because a person is far likelier to acquire one disease than two diseases.

If scientists and other humans prefer lovelier explanations, it is fortunate that they are usually more likely to be true—but a preference for simpler explanations would also be consistent with people performing direct probability calculations rather than using a heuristic. There is, however, some evidence favoring the heuristic view, at least for adults. People's simplicity bias goes beyond what is normatively justified, as people require approximately four times more evidence than they normatively should before abandoning a simple explanation in favor of a more complex one (Lombrozo, 2007). Thus, people seem to use simplicity as a heuristic for estimating prior probability. Similarly, people use complexity in the opposite way, as a heuristic for estimating likelihood (i.e., the probability of the evidence given each hypothesis; Johnson, Jin, & Keil, 2014).

But how do these explanatory heuristics arise? Perhaps adults *learn* to use these heuristics because they approximate normative calculations. In other words, adults might develop a preference for explanations that are lovelier simply because they have learned that they are, on average, likelier. Yet, if explanatory heuristics are such critical sense-making tools, then perhaps they are foundational to our cognitive machinery and guide our explanation evaluations from very early on. If this is the case, then we should expect even young children to use the same explanatory heuristics as adults.

Although even infants can carry out some reasoning in a manner consistent with probability theory (Gweon, Tenenbaum, & Schulz, 2010), it is less clear whether young children also use some of the same explanatory heuristics as adults. That said, some preliminary evidence was provided by Bonawitz and Lombrozo (2012), who found that young children, like adults, require disproportionate evidence before abandoning a simple explanation in favor of a more complex one. In their study, 4- to 6-year-old children encountered a toy that had a light and a fan. Children were taught that putting red coins in the machine caused the light to turn on, putting green coins in the machine caused the fan to turn on, and putting blue coins in the machine caused both the fan and the light to turn on. Then, the experimenter 'accidentally' tipped a bag of coins over, so that either one or two coins fell in, causing both the light and fan to activate. Even if there were many red and green coins but only one blue coin in the bag, so that it was actually *more* probable that both a red and a green coin fell into the machine, children nonetheless favored the simple explanation. Thus, like adults, children appear to use a simplicity heuristic for estimating the prior probability of an explanation.

One limitation of the research on simplicity is that *some* degree of simplicity preference is normatively justified, making it more difficult to distinguish the probabilistic and heuristic views. In the current studies, we capitalized on a non-normative explanatory bias shown by adults—the *latent scope* bias (Johnson, Rajeev-Kumar, & Keil, 2014; Khemlani, Sussman, & Oppenheimer, 2011). For example, imagine that your car smelled like antifreeze, and this could be due to one of two problems—a problem with the cooling system or a problem with the exhaust. Suppose that a cooling problem would activate the "check engine" light, but an exhaust problem would not. Clearly, the thing to do is to check the light. But alas, the light is useless, because the bulb has burned out! In this situation, the light is in the *latent scope* of the cooling system explanation—that is, the light would count as evidence in favor of a cooling problem if it were observed, but the prediction is unverified. Normatively, both explanations are equally likely. Yet, in situations like this, adults prefer explanations with *narrower* latent scope—that is, explanations that make fewer unverified predictions (Khemlani et al., 2011). That is, adults would say that the exhaust explanation—which does not predict any additional effects—is more satisfying and more probable.

This non-normative inference appears to result from a combination of two heuristics (Johnson, Rajeev-Kumar, & Keil, 2014). First, when confronted with an explanation that makes an unverified prediction, people apply an *inferred evidence* heuristic to resolve this ignorance, effectively guessing whether the evidence would be observed if they were able to look. In doing so, people rely on the base rates of the unverified effect, even if the prior probabilities are explicit in the problem (this is what makes the inference non-normative). Second, they apply an *explanatory scope* heuristic, preferring explanations to the extent that they account for as many actual and as few non-actual observations as possible (Johnson, Johnston, Toig, & Keil, 2014; Read & Marcus-Newhall, 1993).

Putting these two heuristics together yields a latent scope bias. Most effects in the world (e.g., check engine lights switching on) have low base rates. Therefore, people typically infer that an unverified effect likely would not have occurred, and count this inferred evidence against the explanation that would predict it. This leads to a preference for narrow latent scope. (Indeed, for cases where the unverified effect has a high base rate, people infer that it probably *would* be observed, and have a *wide* latent scope bias; Johnson, Rajeev-Kumar, & Keil, 2014). Although in this experimental situation, this heuristic leads to error, it is a generally adaptive strategy to try to make inferences about unobserved evidence to maximize one's evidential basis for reasoning.

Given that children are also generally reluctant to accept epistemic ignorance (i.e., ignorance residing in their own mind, rather than in the world; Robinson et al., 2006), it is plausible that they would tend to use an inferred evidence heuristic when evaluating explanations and, thus, show the same, non-normative, latent scope bias as adults. However, if children have not yet acquired the inferred evidence heuristic, then they might respond in accordance with Bayesian norms, and, ironically, outperform adults. Normatively, there is no evidence in favor of either the wide or narrow latent scope explanation, regardless of how high or low the base rates are for the unobserved effect. This is because knowledge of the base rates of potential *explanations* "screens off" information about the effect base rates. A Bayesian child would ignore the base rates of the *effects* and instead (1) calculate the prior probabilities of both explanations and their ratio (i.e., the prior odds), (2) calculate the likelihoods of both explanations and their ratio (i.e., how probable the data would be under each hypothesis; when all that varies across explanations is latent scope, this ratio is 1, because the known evidence is predicted by both hypotheses), and then (3) multiply these two ratios (Pearl, 1988). Although this process is in general much more complex than the heuristic process, it ironically leads to a more straightforward answer in one special case: when the likelihood ratio and prior odds both equal one. In this case, the Bayesian computation, indicating that both explanations are equally probable, is at least as simple as the heuristic computation that leads to a latent scope bias.

In Experiment 1, we tested for a latent scope bias in children, capitalizing on this special case by (1) holding the base rates of competing explanations constant (making the prior probability ratio equal one) and (2) manipulating only the unobservable evidence across explanation (i.e., making the likelihood ratio equal one). Thus, to the extent that the task would be too demanding for children, this would make them look like Bayesians rather than like heuristic reasoners. In Experiment 2, we varied the base rates of the explanations to make the wide latent scope explanation more probable, testing whether the latent scope bias, like simplicity (Bonawitz & Lombrozo, 2012), can be overridden by strong prior odds.

Experiment 1

In Experiment 1, children encountered a toy that, like Bonawitz and Lombrozo's (2012), had a fan and a light. Children learned that one color coin turned on the fan (the *one-effect* coin) and that the other color coin turned on both the fan and light (the *two-effect* coin). After several familiarization trials with these coins (in which various parts of the toy were occluded), children were presented with one test trial in which the light was occluded so that they could not tell whether it was on or not. Then, one coin was randomly and covertly put into the machine and children were asked to infer which coin was placed inside. The coin was drawn from a bag containing 5 coins of each color, to ensure that the prior probabilities were equal. If children respond normatively, they should guess at chance, because the fan is not diagnostic (it is consistent with either explanation), and the key piece of information (the light) is unavailable. In contrast, if children show a latent scope bias like adults (Khemlani et al., 2011), they should indicate that the one-effect coin is more likely, since it does not make the additional, unverified prediction that the light would be on.

Method

Participants Thirty-one 4 and 5-year-old children (M = 4 years, 11 months; range = 4 years, 0 months – 6 years, 0 months) participated in Experiment 1. An additional 14 children (11 4-year-olds and 3 5-year-olds) participated but were replaced because they failed the familiarization check questions (see below).

Materials The materials included a machine toy (see Figure 1), constructed from white cardboard. On the top of the machine, facing the child, were a fan that could rotate and a light that could turn on. A slot at the front of the machine was used to drop coins in, which purportedly caused the fan or light to operate. In fact, the fan and light were covertly operated by the experimenter using switches wired to the back of the box, out of view of the child. No child voiced suspicion over the operation of the machine; in fact, a senior museum staff member at one of our testing sites was surprised to learn that the coins did not control the machine.

Procedure The procedure involved three phases: The *introduction*, *familiarization*, and *test* phases.

In the *introduction* phase, the experimenter explained the function of the blue and red coins. One coin (the *one-effect* coin) made just the fan turn, while the other coin (the *two-effect* coin) made both the fan and light turn on. The color of the coins was counterbalanced, such that the one-effect coin was blue for some children and red for others. For each coin, the experimenter put the coin in the slot so the child could witness what the coin caused the box to do. The experimenter then said, "See! The blue

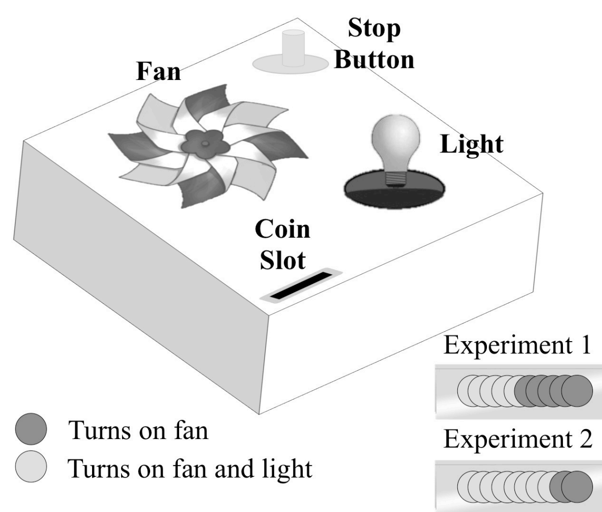

Figure 1: Machine toy used in Experiments 1 and 2, including the coins that operated the machine and their base rates across experiments. The light was occluded on test trials so that children could not observe whether it was on. The toy was oriented so that the child faced the coin slot.

[*red*] coin makes the fan [*both the fan and the light*] go." After introducing each coin, the experimenter gave a card to the child depicting the coin's color and its effects to reduce the task's memory load. The order in which the experimenter introduced the coins was randomized.

Next, in the *familiarization* phase, the child made six predictions—two in which both parts of the toy were visible and four in which one part was occluded—about what would happen if coins were put into the toy. If a child required more than one correction on the same familiarization trial (either visible or occluded), that child did not proceed to the test phase and was excluded from data analysis. On the first set of familiarization trials (i.e., the two *visible* trials), the child was asked to predict what would happen when the red and blue coins were put into the slot. These trials were intended to make sure that the children remembered or could rely on their diagrams to understand how the machine worked. If the child answered incorrectly, the experimenter put the coin in to show the child the correct answer, and the trial was repeated. The order of the two visible trials (for the red and blue coins) was randomized.

On the second set of familiarization trials (i.e., the four *occluded* trials), either the fan or the light was covered up using an opaque cardboard cover, and the child was asked to predict what would occur when each color coin was placed in the slot. These trials were framed as a guessing game, wherein parts of the machine were sometimes covered. This was done in order to break any pedagogical or pragmatic inferences children might be making about what the experimenter was communicating by covering the fan and light, and to ensure that children understood that unobserved effects could still occur. If the child answered incorrectly, the experimenter lifted the cover, and the trial was repeated. The order of the four invisible trials (for the red and blue coins, and with either the fan or light covered) was randomized.

Finally, in the test phase, the light was occluded. The test trial was continuous with the familiarization trials, so that from the child's perspective, covering the light on this trial was no different than covering parts of the machine on the previous familiarization trials. The experimenter showed the child a transparent plastic bag containing five red coins and five blue coins and said:

> We're going to use this bag of coins! See, there are 5 red coins and 5 blue coins in this bag. I'm going to close my eyes and pull one out. Then, I'll put it in the box, and I want you to guess which color went in.

Then, the experimenter and child both closed their eyes, and the experimenter selected a coin at random from the bag, so that the child could not see what coin was selected. The experimenter then placed the coin in the slot and the appropriate effects occurred (i.e., the fan always turned on, and the occluded light did or did not turn on, depending on the coin color). Then, the experimenter asked, "Which color do you think went in?" If children are averse to latent scope, they should choose the one-effect option (the light should be off), but if they prefer latent scope explanations, they should choose the two-effect option (the light should be on). Alternatively, if children are indifferent to latent scope and respond normatively, they should choose the coins equally often.

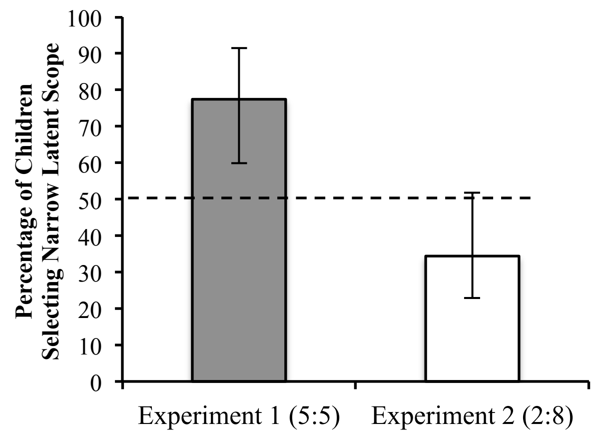

Figure 2: Results of Experiments 1 and 2. Dashed line indicates chance responding, and bars represent 95% CIs.

Results and Discussion

As shown in Figure 2, children preferred the explanation with narrow latent scope—the coin that caused only the fan to turn on. Specifically, on the test trials, 24 out of 31 children (77%) chose the narrow latent scope coin (p = .003, sign test). This preference was equally strong among 4- and 5-year-olds (p = 1.00, Fisher's exact test). These results demonstrate that children as young as age 4 have a robust latent scope bias, suggesting that even very young children are swayed by some of the same non-normative explanatory preferences as adults.

Experiment 2

Children have surprisingly sophisticated probabilistic reasoning skills, starting from infancy (Gweon et al., 2010). In particular, children use the base rates of explanations to calibrate their preference for simpler over complex explanations (Bonawitz & Lombrozo, 2012). Specifically, when the base rates of the simple and complex explanations are made equal by varying the number of colored coins, children (like adults; Lombrozo, 2007) prefer the simple explanation. But when the complex explanation is much more probable than the simple explanation (a 1:6 ratio), children are able to override their simplicity preference and choose the more probable explanation. Would children similarly be able to override their latent scope bias when the base rates favor the wide latent scope explanation?

To test how children integrate explanatory scope and base rates in their explanatory inferences, we manipulated the prior odds using the method of Bonawitz and Lombrozo (2012). Instead of drawing a coin at random out of a bag with 5 two-effect and 5 one-effect coins as in Experiment 1, the bag contained 8 two-effect and 2 one-effect coins. That is, the wide latent scope explanation had a prior probability that was 4 times as high as the narrow latent scope explanation. If children can override their latent scope bias by using probabilistic information, they should choose the more probable two-effect coins. But if overwhelming prior odds are still insufficient to override the latent scope bias, then they should continue to choose the one-effect coins with narrow latent scope.

Method

Participants Thirty-two 4- and 5-year-old children (*M* = 4 years, 11 months; range = 3 years, 11 months – 5 years, 10 months) participated in Experiment 2. An additional 6 children (all 4-year-olds) participated but were replaced because they failed the same familiarization trial at least two times (the same criterion used in Experiment 1).

Materials and Procedure The materials and procedure were identical to those for Experiment 1, except for the test trial. On that trial, the experimenter used a bag of coins with 8 two-effect coins (i.e., wide latent scope) and 2 one-effect coins (i.e., narrow latent scope), in contrast to Experiment 1 where 5 of each type of coin were used.

Results and Discussion

As shown in Figure 2, the results of Experiment 2 differed dramatically from those of Experiment 1. Whereas 24 out of 31 children (77%) in Experiment 1 chose the narrow latent scope coin when the coins were equally probable, only 11 out of 32 children (34%) chose the narrow latent scope coin in Experiment 2 where the narrow latent scope coin was more probable. Thus, children in Experiment 2 chose the narrow latent scope explanation less often than children in Experiment 1 ($p < .001$, Fisher's exact test), and, if anything, showed a preference for the *wide* latent scope explanation ($p = .11$, sign test), a preference that was equally strong among 4- and 5-year-olds ($p = .46$, Fisher's exact test).

These results show that young children are able to combine information about an explanation's scope and its prior probability. Like the simplicity bias, the latent scope bias can be overridden by strong prior odds. Explanatory heuristics therefore are not used blindly, but in concert with other sources of evidence in a flexible manner.

General Discussion

Children may be scientists, but what kind of scientists are they? Do they search for the *likeliest* explanations, like good Bayesians, or do they search for the *loveliest* explanations, as some philosophers of science recommend (Lipton, 2004) and many scientists actually do in practice (Zee, 1999)? In two experiments, we demonstrated that children, like adults, have a non-normative preference for *narrow latent scope* explanations—explanations that make few unverified predictions. The early emergence of this bias constitutes further evidence that explanatory heuristics are not merely quirks of adult cognition, but a fundamental component of explanatory reasoning that may undergird later, more sophisticated behaviors.

In Experiment 1, children preferred narrow latent scope explanations over wide latent scope explanations, even when their probabilities were matched. Experiment 2 was an exact replication of Experiment 1, except that the prior probabilities favored the *wide* latent scope explanation. This change eliminated the latent scope bias (actually reversing it), showing that the bias can be overridden by strong prior odds. This speaks to the flexible manner in which explanatory heuristics can be integrated with other sources of evidence.

However, one possible concern is that children's latent scope bias is not due to adult-like heuristic processing, but instead to a different, lower-level process. Perhaps children chose the one-effect coin merely because that coin corresponded to the one effect they could observe (a perceptual matching bias). However, this interpretation is unlikely to be correct for two reasons. First, this bias was overridden by probabilistic evidence in Experiment 2, meaning that children could integrate multiple sources of evidence rather than blindly perceptually matching. Second, if the results were due to perceptual matching, one would expect stronger effects at younger ages. However, there was no age difference in either Experiment 1 or 2. Further, we conducted an additional test of children's latent scope bias using a different method with 5- to 8-year-olds (Johnson, Johnston, Koven, & Keil, 2015). Not only was the latent scope bias replicated using a different method, but there were once again no age differences even across this wider age range.

The non-normativity of the latent scope bias—as well as the underlying heuristic mechanisms (Johnson, Rajeev-Kumar, & Keil, 2014)—can help to distinguish between probabilistic ("likeliest") and heuristic ("loveliest") accounts of explanatory reasoning. According to Bayesian confirmation theory (e.g., Jeffrey, 1965), the best explanation of a phenomenon is the explanation that is likeliest to have caused it. This idea has been refined by advances in statistics and machine learning (Pearl, 1988), which use Bayesian networks to model the conditional independence assumptions that vastly simplify the computational problem of causal learning and reasoning. Psychological versions of these theories have seen great success in modeling human causal reasoning in children (Gopnik et al., 2004), adults (Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), and even rats (Blaisdell, Sawa, Leising, & Waldmann, 2006). Thus, explanatory inferences often approach normative ideals—a consistent finding across studies that accrues support for the view that children and adults infer the likeliest explanation.

However, in many cases, it may be sufficient to rely on heuristics that typically approximate normative inferences, rather than going through the more cognitively demanding task of explicit probability calculations. Heuristics such as simplicity (i.e., preferring explanations invoking fewer causes, all else equal) and scope (i.e., preferring explanation that explain more of the evidence, all else equal) are normatively grounded, in that following them will lead to rational inferences, yet they are more computationally straightforward than explicit probability calculations. In fact, several studies suggest that children and adults rely on simplicity to estimate prior probabilities heuristically (Bonawitz & Lombrozo, 2012; Lombrozo, 2007) and that adults rely on complexity to estimate likelihood (i.e., the probability of the evidence given each hypothesis; Johnson, Jin, & Keil, 2014).

The current results provide even more powerful evidence for the heuristic approach, in documenting a non-normative behavior by Bayesian standards. In our Experiment 1, 77% of children preferred an explanation that did not posit an unobservable piece of evidence, even though the children could clearly see that the two explanations had equal base rates (i.e., the same number of red and blue coins in the bag from which the coin was randomly selected). Further research could explore whether this bias might extend to even younger ages to further rule out the possibility that it is a learned heuristic.

Though the probabilistic and heuristic views may appear to be competitors, they need not be. Although people do not appear to be Bayesians at an algorithmic level, it is equally clear that people often make sophisticated inferences that are more-or-less normative at the computational level. Since most Bayesian theories are posed at the computational level, the heuristic account need not be in tension with such probabilistic approaches. Rather, heuristics can allow us to implement reasoning that can approximate Bayesian norms, in a way that is tractable given our cognitive limits.

Thus, children's latent scope bias may be best viewed not as an inferential failure, but as one part of a grander method—an arsenal that may contain many explanatory heuristics, working in concert—that we can use to understand our environment, to explain what happens, and to make sense of the world. Contra Keats, beauty may not be the very essence of truth—but the explanatory virtues may suffice to get by, most of the time.

Acknowledgments

We thank the staff of the Yale Peabody Museum and Stepping Stones Museum for Children, the members of the Cognition and Development Lab, and the children and parents who participated in this research.

References

Blaisdell, A.P., Sawa, K., Leising, K.J., & Waldmann, M.R. (2006). Causal reasoning in rats. *Science*, *311*, 1020–1022.

Bonawitz, E.B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology*, *48*, 1156–1164.

Gopnik, A., Glymour, C., Sobel, D.M., Schulz, L.E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.

Gweon, H., Tenenbaum, J.B., & Schulz, L.E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*, 9066–9071.

Jeffrey, R. C. (1965). *The logic of decision*. New York, NY: McGraw-Hill.

Johnson, S.G.B., Jin, A., & Keil, F.C. (2014). Simplicity and goodness-of-fit in explanation: The case of intuitive curve-fitting. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 701–706). Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Johnston, A.M., Koven, M.L., & Keil, F.C. (2015). *Little Bayesians or little Einsteins? Disentangling probabilistic and heuristic accounts of explanation in children.* Manuscript in preparation.

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. 2453–2458). Austin, TX: Cognitive Science Society.

Johnson, S.G.B., Rajeev-Kumar, G., & Keil, F.C. (2014). Inferred evidence in latent scope explanations. 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.

Khemlani, S.S., Sussman, A.B., & Oppenheimer, D.M. (2011). *Harry Potter* and the sorcerer's scope: Latent scope biases in explanatory reasoning. *Memory & Cognition*, *39*, 527–535.

Lipton, P. (2004). *Inference to the best explanation* (2nd Edition.). London, UK: Routledge.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*, 232–257.

McGrew, T. (2003). Confirmation, heuristics, and explanatory reasoning. *The British Journal for the Philosophy of Science*, *54*, 553–567.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Francisco, CA: Morgan Kaufmann.

Read, S. J., & Marcus-Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology*, *65*, 429–447.

Robinson, E. J., Rowley, M. G., Beck, S. R., Carroll, D. J., & Apperly, I. A. (2006). Children's sensitivity to their own relative ignorance: Handling of possibilities under epistemic and physical uncertainty. *Child Development*, *77*, 1642–1655.

Steyvers, M., Tenenbaum, J.B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.

Zee, A. (1999). *Fearful symmetry: The search for beauty in modern physics.* Princeton, NJ: Princeton University Press.