

The origins of inquiry: inductive inference and exploration in early childhood

Laura Schulz

Massachusetts Institute of Technology, MIT Department of Brain and Cognitive Sciences, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

Analogies between scientific theories and children's folk theories have been central to the study of cognitive development for decades. In support of the comparison, numerous studies have shown that children have abstract, ontologically committed causal beliefs across a range of content domains. However, recent research suggests that the comparison with science is informative not only about how children represent knowledge but also how they acquire it: many of the epistemic practices essential to and characteristic of scientific inquiry emerge in infancy and early childhood.

Core epistemic practices

Science is a historically recent, culturally specific endeavor practiced by a tiny minority of the human species. As such, it might seem a peculiar place to look for universals of human cognition. However, science has a peculiar property: it (often enough) gets the world right. With thoughts that were unthinkable a few hundred years ago – germs, genes, bits and bosons – we predict the future, explain the past, and intervene effectively in the present. How is this possible? What kind of epistemic practices could enable this kind of learning?

The answer to this question, of course, is that we do and do not know. Nevertheless, we can both formally and informally characterize many of the epistemic practices that are fundamental to scientific inquiry across content domains (i.e., that are practiced by experimental physicists, chemists, biologists, paleontologists and psychologists alike). For instance, scientists in every field:

- Rationally infer causal relationships from patterns of statistical evidence.
- Have theories that structure and inform the interpretation of statistical evidence.
- Infer the existence of unobserved variables to explain anomalous data.
- Selectively explore when evidence is confounded or surprising.
- Isolate candidate causes in order to distinguish between competing hypotheses.
- Constrain their generalizations depending on how evidence is sampled.
- Decide when to rely on others' knowledge and when to initiate new investigations.

You could do everything on this list and still not be able to do science. Science requires bringing these inferential processes to bear on detailed knowledge about the world. Moreover, scientists are skilled users of specialized physical and cognitive tools; they can investigate properties of the world that are otherwise intractable (too big, too small, occurring in too fast or too slow a temporal window, or too complex for ordinary inference). However, if scientists did not also engage in the epistemic practices listed above, they would be unlikely to learn anything at all. These abilities are arguably essential to human inquiry. Although many animals explore to learn what is in the world and what can be done with it (e.g., [1–3]), as far as we know, only human beings explore to learn why things happen and only these kinds of inferential processes reliably generate causal discoveries. Although these abilities are both characteristic of and critical to scientific discovery, I will suggest that they are not fundamentally scientific abilities; they are universal abilities at the foundations of human cognition.

Because these are the kinds of inferential abilities that enable rapid, accurate, abstract learning, there are theoretical grounds for supposing that they emerge early in infancy and are continuous throughout development. Empirically however, these abilities have been investigated in children ranging in age from infancy to middle childhood, with the age of study constrained primarily by the complexities of ancillary task demands. This provides an existence proof of core epistemic abilities in childhood, but leaves open the question of whether these abilities are continuous throughout development.

Statistical evidence, folk theories, and rational inference

Starting a decade ago, researchers demonstrated that very young children have an ability basic to scientific discovery: the ability to distinguish variables spuriously associated with outcomes from genuine causes. If, for instance, a red block activates a toy both when it is placed on the toy by itself and when it is placed on the toy together with a blue block, but the blue block only activates the toy when the red block is also present, preschoolers infer that the red block, not the blue one, makes the toy go. Control conditions established that children's inferences depend on the probability of the outcome given the intervention, not the frequency of the outcome or the number of failed interventions [4]. Several studies have since replicated this finding, showing that across a range of tasks, ages, and content

Corresponding author: Schulz, L. (lschulz@mit.edu).

domains, children use the conditional probability of events to make causal judgments ([5–9]; see [10] for review).

Learners can draw accurate inferences from the evidence of just a handful of trials only if their inferences are constrained by more abstract theories [11–15]. Researchers have proposed that children have such theories at many different levels of abstraction, ranging from broad, framework theories of naïve physics, naïve psychology, and naïve biology [14,16–24] to quite local causal beliefs (e.g., about the relationship between balance and mass [25], density and floating [26], and ball properties and tennis serves [27]). At every level of abstraction, these prior beliefs affect learners' judgments about statistical evidence [28–31]. Given, for instance, identical co-variation evidence, preschoolers are more likely to accept candidate causes that are common over those that are rare, and candidate causes that are theory consistent over those that are theory violating [32,33]. However, given sufficient evidence in support of an unlikely cause, preschoolers change their minds (e.g., about whether flipping a switch or talking to a toy will activate the toy [7], whether a toy will activate on contact or at a distance [6], or whether eating food or being scared causes a tummy ache [34,35]). In simple contexts, work has shown that even pre-verbal infants can use sparse data about the co-variation of interventions and outcomes to make rational causal attributions [36] (Figure 1).

Learners' ability to update their beliefs given new evidence can be formally characterized by hierarchical Bayesian inference models (Box 1). These models have motivated many of the empirical studies reviewed here

and successfully predicted both their qualitative and quantitative results.

Inferring unobserved variables

The studies discussed above focused on inferences from observed patterns of data, however, many of the most important scientific discoveries involve inferences about unobserved, and even unobservable, variables. Several recent studies have shown that even very young children introduce unobserved variables to maintain causal beliefs at different levels of abstraction and entrenchment.

Research suggests, for instance, that preschoolers assume that perfect knowledge of the causes of an event should enable perfect prediction of its effects (a strong form of 'causal determinism'). Determinism may not be an accurate claim about the state of the world; nonetheless, the assumption of determinism has played a critical role in scientific discovery. Because scientists typically expect causal mechanisms to behave predictably, they infer unobserved, latent variables when evidence is anomalous with respect to their prior beliefs.

Research suggests that when preschoolers see an apparently probabilistically effective cause, they infer either that a necessary generative cause is sometimes missing or that a sufficient inhibitory cause is sometimes present. Moreover, they rationally trade off these inferences: if they know that a necessary generative cause might be absent, they are less likely to infer that an inhibitory cause might be present [37]. Even toddlers seem to assume determinism; although they faithfully imitate deterministically effective actions, they explore rather than imitate probabilistically effective ones

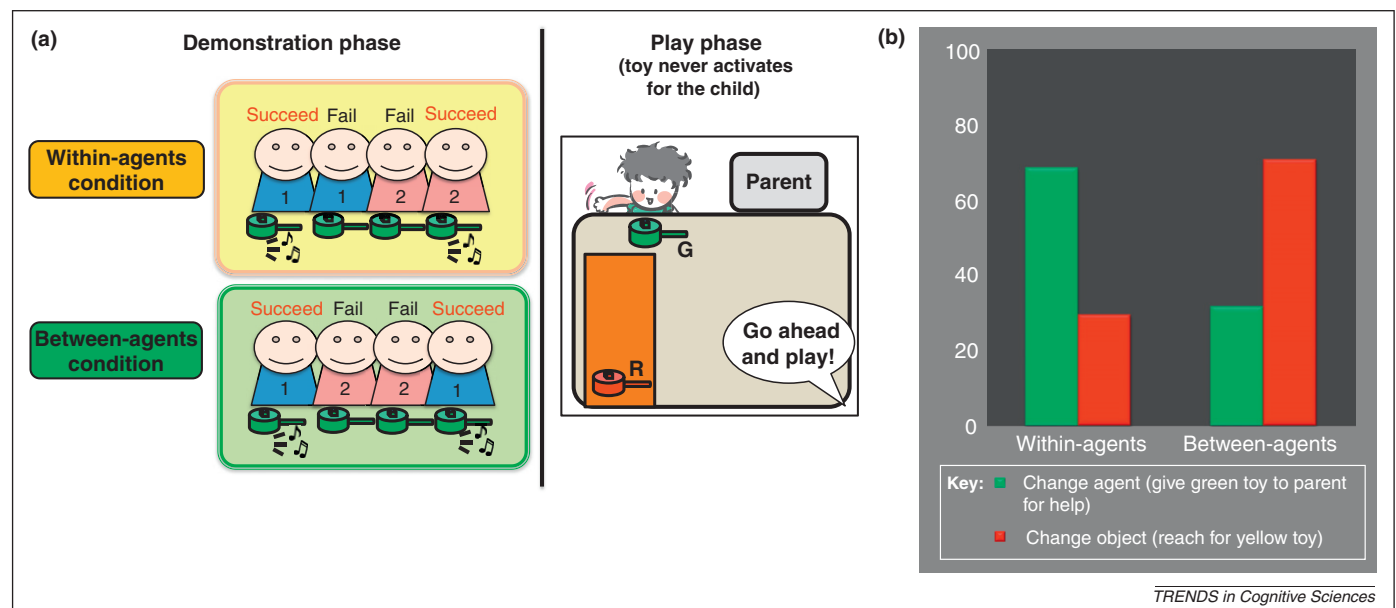


Figure 1. In the context of goal-directed actions, an individual's influence on events and the influence on the outside world are often confounded. This experiment asked whether infants (mean: 16 months; range 13-20 months) could rationally distinguish these causal attributions using minimal statistical data. **(a)** Infants saw two experimenters (labeled 1 and 2 in the figure) push a button on a toy that played music when activated. In the Within-agents condition, both experimenters successfully activated the green toy (G) once and failed once. In the Between-agents condition, one experimenter successfully activated the green toy twice; the other failed twice. The infants were then handed the green toy (G); a red toy (R), identical to the green one except for color, was placed on a cloth within the infants' reach. All infants were seated next to their parents. All infants pushed the button on the toy and the toy always failed to activate for the infants. The outcomes in the Within-agents condition (considering also the infants' failure) vary independently of the agent, suggesting that the failure is due to the object; the outcomes in the Between-agents condition co-vary with the agent independent of the object, suggesting that the failure is due to the agent. **(b)** As predicted, infants were more likely to first change the agent (by handing the toy to their parents) than the object (by reaching for the new toy) in the Between-agents than Within-agents conditions ($p < .05$ by Fisher's Exact test; change agent vs. change object: Within-agents, 29.4% vs. 71.6%; Between-agents, 68.4% vs. 31.6%). These results suggest that infants track the statistical dependence between agents, actions and outcomes and can integrate prior knowledge and statistical data to make rational causal attributions. Reproduced, with permission, from [36].

Box 1. Hierarchical Bayesian inference models

Hierarchical Bayesian inference models provide a formal account of how abstract knowledge can be learned from sparse data. Bayes' law states that the learner's belief in a hypothesis after observing evidence, the posterior probability of the hypothesis, $P(h|e)$, is proportional to (i) its likelihood, $P(e|h)$, that is, the probability that the hypothesis, if true, would have generated the observed evidence, and (ii) its prior probability, $P(h)$, that is, the probability that the hypothesis is generated by the learner's background theories. Formally:

$$P(h|e) \propto P(e|h)P(h) \quad [I]$$

If for instance, a learner observes evidence (e) of a child coughing and considers the possibility (h_1) that the child has a cold, (h_2) that the child has influenza, and (h_3) that the child has lung cancer, the prior probability favors h_1 and h_2 , because colds and flu are more common than lung cancer. However, the likelihood favors h_1 and h_3 , because colds and lung cancer are more likely than the flu to generate coughing. Thus the posterior probability favors the hypothesis that the child has a cold (example from, and detailed account available in [62]).

Bayes' law also provides a formal account unifying what might otherwise seem like very different routes to uncertainty and exploration. Bayes' law suggests that a learner will be uncertain whenever the posterior probability of two or more hypotheses is equivalent:

$$P(h_1|e) \approx P(h_2|e) \quad [II]$$

This can occur if the prior probability favors one hypothesis and the likelihood another:

$$P(e|h_1) < P(e|h_2) \text{ and } P(h_1) > P(h_2) \quad [III]$$

Or if the prior probability and the likelihood of multiple hypotheses are equivalent:

$$P(h_1) \approx P(h_2) \text{ and } P(e|h_1) \approx P(e|h_2) \quad [IV]$$

Line III is a formalization of what it means for evidence to be surprising; line IV, of what it means for evidence to be confounded. Thus Bayes' law provides an intuitive account of why exploration in the face of surprise and confounding derive from a common principle.

These principles follow directly from the simplest form of Bayes' law. However, Bayesian inference can be extended hierarchically, allowing the learner to do probabilistic inference at multiple levels of abstraction. A learner, for instance, may not only have the specific prior belief that colds are a more common cause of coughing than lung cancer, but also have the more abstract knowledge that diseases generate symptoms. This allows the learner to rule out many specific hypotheses about possible causal relations (e.g., that coughing causes lung cancer) that are otherwise consistent with an observed correlation [15].

Moreover, the learner can learn abstract beliefs at the same time, or even prior to, learning the more specific beliefs they constrain. This ability to engage in joint inference can be illustrated intuitively by a thought experiment proposed by the philosopher Nelson Goodman [68]. Imagine walking into a room with thousands of brown paper bags, each containing a hundred marbles. Suppose you reach into one bag and pull out a single marble at random. It is red. You reach into the same bag and pull out a second marble at random. Again it is red. You reach in a third time and pull out a third red marble. Now you move onto a new bag and pull out a single marble. It is blue. You reach in a second time. Again the marble is blue. Then you move onto a third bag. You pull out a green marble, and then a second green marble. At this point you have seen a handful of marbles in a room containing hundreds of thousands of marbles. Nonetheless, from this tiny bit of data you might have drawn a powerful inductive inference: the contents of the bags in the room are homogeneous. This very abstract inference can support even more rapid learning from subsequent data: if you pull a single purple marble from the next bag, there are 99 marbles in the bag that you have not seen; nonetheless, you might guess that they are all purple.

Hierarchical Bayesian inference models provide a formal account of our ability to jointly infer such abstract 'over-hypotheses', together with subordinate hypotheses that constrain the interpretation of subsequent data, showing for instance that learners can simultaneously infer a specific tree-structure and the fact that the data are organized as a tree (e.g., rather than a ring or a chain) [60–64]. Recently researchers have shown that even infants can infer over hypotheses from sparse data [65].

(despite remembering the actions equally well in both contexts). Such flexible behavior could allow children to discover hidden causes that might account for the otherwise unpredictable failures of these actions [38]. Preschoolers will even introduce unobserved variables to preserve abstract beliefs that are themselves derived from minimal data [39] (Figure 2).

Selective exploration and isolating variables

The link between learning and exploration lies at the heart of scientific discovery. Although researchers have long believed that children 'learn by doing', until recently there was relatively little empirical support for this claim and little understanding of how the seemingly unsystematic behavior that characterizes exploratory play might relate to the inferential processes thought to support learning. However, recent work suggests that children's exploratory behavior is driven not just by the perceptual novelty or salience of stimuli, but by formal properties of evidence. In particular, consistent with the idea that learners should be curious about causal hypotheses when the posterior probability of a small number of hypotheses is equivalent (Box 1), my collaborators and I have shown both that children selectively explore both when the prior probability favors one hypothesis but observed evidence is more likely under another (i.e., when evidence is

surprising) and when the evidence is equally consistent with multiple plausible hypotheses (i.e., when evidence is confounded).

For instance, six and seven-year-olds with different beliefs about how objects balance engage in different patterns of both exploration and explanation given identical evidence. Children who have the (correct) mass theory of balance selectively explore a novel toy when given a choice between playing with a novel toy and a familiar asymmetric block balanced over its center of mass. However, children who (incorrectly) believe that all objects balance over their geometric center explore the block instead. Children show the opposite pattern of responses when given the choice between the novel toy and an asymmetric block balancing over its geometric center [40].

Children also explore violations of more abstract beliefs. One virtue of inductive generalization is that it obviates the need for trial and error learning: if you learn that one *lep* is magnetic, you can infer that other *leps* are magnetic without testing them (see, for example, [41,42]). By the same principle however, when such inductive generalizations are violated (e.g., only some *leps* are magnetic), exploration is rational. Consistent with this, preschoolers are more likely to explore perceptually identical objects when a causal property varies within kinds than across kinds ([43]; see also [44]). Moreover, children's exploration




Figure 2. Can preschoolers (mean: 57 months; range: 49-64 months; $n = 16$ per condition) infer abstract laws from minimal data? This study, predicted that, if so, infants would infer the presence of an unobserved entity to account for evidence that would otherwise be anomalous if construed with respect to the abstract law. **(a)** In Phase 1, children saw a few trials of evidence in which red blocks activated blue blocks, blue blocks activated yellow blocks, and other interactions were inert (Conditions A and B) or red blocks activated both blue and yellow blocks and other interactions were inert (Conditions C and D). In Phase 2, children either saw a novel white block activate a blue block (Conditions A and C) or a red block activate the novel white block (Conditions B and D). In Phase 3, all children saw the white block go behind a curtain. A yellow block adjacent to the curtain then activated. If preschoolers are 'concrete' learners they should assume that the white block is the only block behind the curtain. Schulz *et al.* predicted that children would instead make a more abstract inference, inferring from the fact that the white block apparently activated the yellow block that the white block was the same kind of block as the red block. If so, the evidence that the white block apparently activated the yellow block would seem anomalous. **(b)** As predicted, although all the children had seen only the white block go behind the curtain, only when the evidence was anomalous with respect to the causal laws (i.e., in Conditions A and D) did children infer the presence of an additional, unobserved block. All children reached behind the curtain once and were given the white block; would children reach a second time? Children were more likely to reach twice in Condition A than in Condition B (69% vs. 6%; $\chi^2(1, n = 32) = 13.33, p < .001$) and in Condition D than in Condition C (50% vs. 6%; $\chi^2(1, n = 32) = 7.57, p < .01$). There were no differences between Conditions A and D or B and C. Reproduced, with permission, from [39].

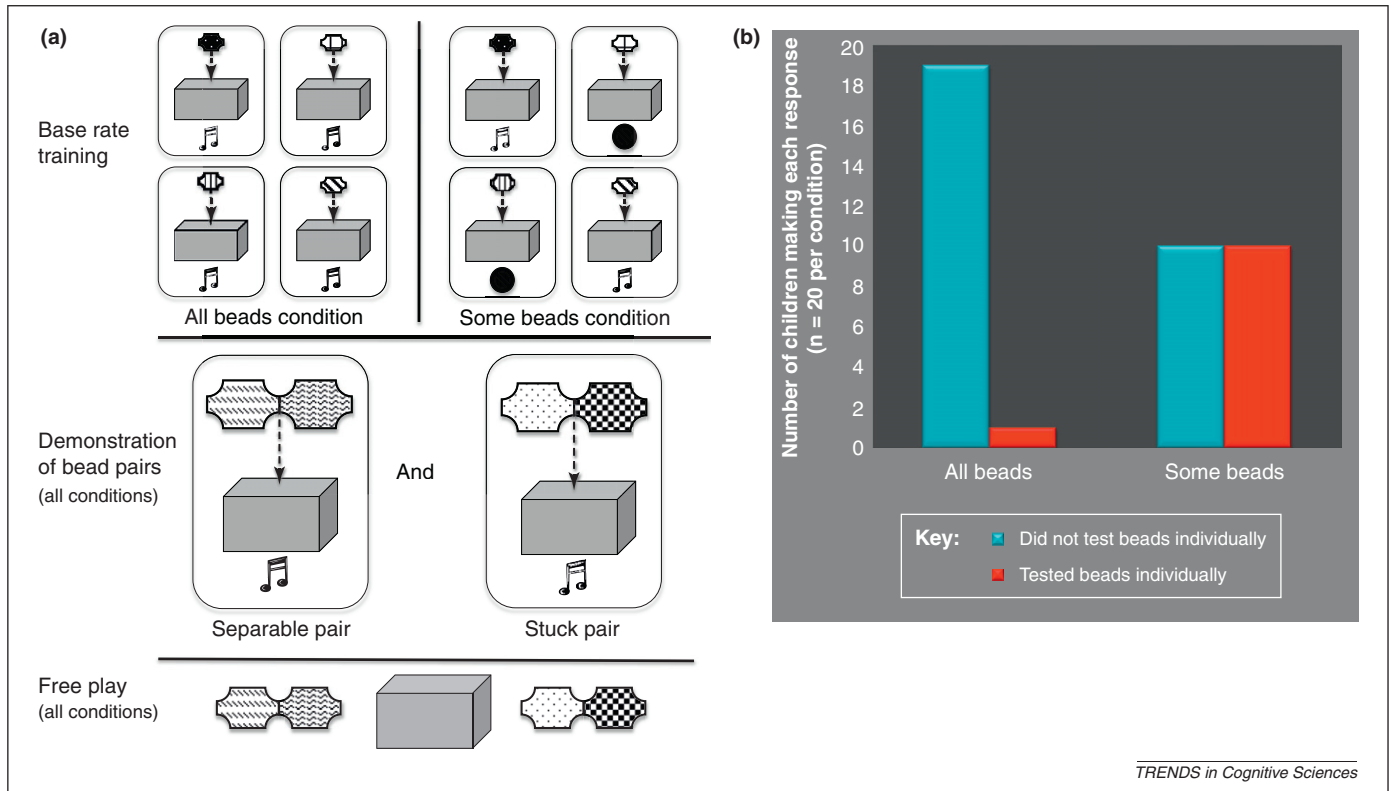
is systematically related to their tendency to generate causal explanations for anomalous evidence. Children who expect *tomas* to activate toys and *blickets* not to activate toys can explain a *toma's* unexpected failure to activate the toy either by referring to a category mistake (e.g., 'It's really a blicket') or a causal mistake (e.g., 'You put it on the wrong side'). Children who generate specifically causal (as opposed to category-switch) explanations engage in more exploratory behavior [45] (see also [46]).

Such studies have looked at children's exploration when evidence is surprising; other studies have looked at preschoolers' tendency to explore confounded evidence. Relative to a novel toy for instance, preschoolers are more likely

to explore a familiar toy when it generates confounded than un-confounded evidence [47]. Moreover, preschoolers are sensitive not only to the relative ambiguity of evidence, but also the potential information gain associated with different interventions. Given ambiguous evidence, children selectively perform interventions that isolate competing candidate causes and maximize information gain [48]; see also [49]) (Figure 3).

Sampling processes

Scientists routinely generalize from small samples of data and draw different generalizations depending on the degree to which a sample is thought to be representative of a



TRENDS in Cognitive Sciences

Figure 3. This experiment investigated whether preschoolers ($n = 20$ per condition; mean: 54 months, range: 46-64 months) would isolate causal variables consistent with a folk theory of contact causality, specifically when there was the possibility of information gain. (a) Children were shown either that 4 of 4 beads (the *All Beads* condition) or 2 of 4 beads (the *Some Beads* condition) activated a toy when the beads were placed, one at a time, on top of the toy. Both groups of children were then shown two bead pairs. One pair could be pulled apart into two individual beads; the other bead pair was glued together. The children saw that both bead pairs, as pairs, activated the toy. In the *All Beads* condition the base rate information supports the hypothesis that both beads in both pairs activate the toy, thus there is relatively little potential for information gain. By contrast, in the *Some Beads* condition, the evidence is genuinely ambiguous: either or both beads might be effective. All children were left alone to play with the beads and the toy. The prediction was that children would be more likely to snap apart the separable bead pair and test each bead individually on the toy in the *Some Beads* conditions than the *All Beads* condition. (b) As predicted, children in the *Some Beads* conditions were more likely than children in the *All Beads* condition to perform the informative action that could disambiguate the evidence (50% vs. 5%; Fisher's Exact, $p < .05$). In a second experiment, children ($n = 20$ per condition; mean: 54 months, range: 43-63 months) received the same base rate training but then were given only the stuck pair of beads to explore. In the *Some Beads* condition, 45% of the children (compared to only one child, 5%, in the *All Beads* condition; Fisher's Exact, $p < .05$) spontaneously designed a novel intervention to isolate the variables: rotating the beads vertically so that only a single bead contacted the toy at a time. These results suggest that preschoolers distinguish potentially informative and uninformative interventions and selectively perform interventions that isolate variables and maximize information gain. Reproduced, with permission, from [48].

target population. By 15 months, infants also constrain their generalizations of object properties depending on whether they think evidence has been sampled randomly or selectively [50] (Figure 4). Abundant work now suggests that a sensitivity to the relationship between samples and populations emerges early in infancy; recent studies suggest as early as six-months [51]. Infants for instance, expect a box containing more red balls than white balls to generate a sample of more red than white balls; similarly, they expect that a sample of mostly red balls came from a population of mostly red balls. Moreover, infants suspend such inferences if the experimenter who pulls the sample first expresses a preference for the minority ball [52] (see also [53]). Slightly older infants make the same kind of inference in reverse: if an agent selects only frogs from a box containing mostly ducks, children infer that the agent has a preference for frogs, and the more improbable the sample, the more likely children are to assume the agent has a preference [54,55]. Thus before they are two, infants seem to understand that evidence can be sampled in different ways, that different sampling processes will generate different evidence, and thus that different generalizations are warranted.

Learning from others

New scientific investigations are informed (and sometimes rendered unnecessary) by studies already part of the literature; what scientists decide to explore depends on what they believe is already known. Such rational inferences about the advantages of learning from others and learning from self-guided exploration also govern the behavior of preschoolers [44,56–58]. If for instance, a teacher freely demonstrates one function of a toy (and functions are rare), the learner can rationally assume that there is only one function; if additional functions were present, and the teacher is knowledgeable and helpful, she would have demonstrated these as well. This suggests a trade-off between instruction and exploration: children who are deliberately instructed in the functions of a toy should explore the toy less than children who are shown identical evidence by an interrupted or naïve agent. In the pedagogical condition, the absence of evidence for additional functions provides uniquely strong evidence for their absence (see [59] for a formal account of such pedagogical sampling assumptions). My collaborators and I tested this prediction by looking at four-year-olds' exploratory behavior in a pedagogical condition ('Look at what my toy does.') or

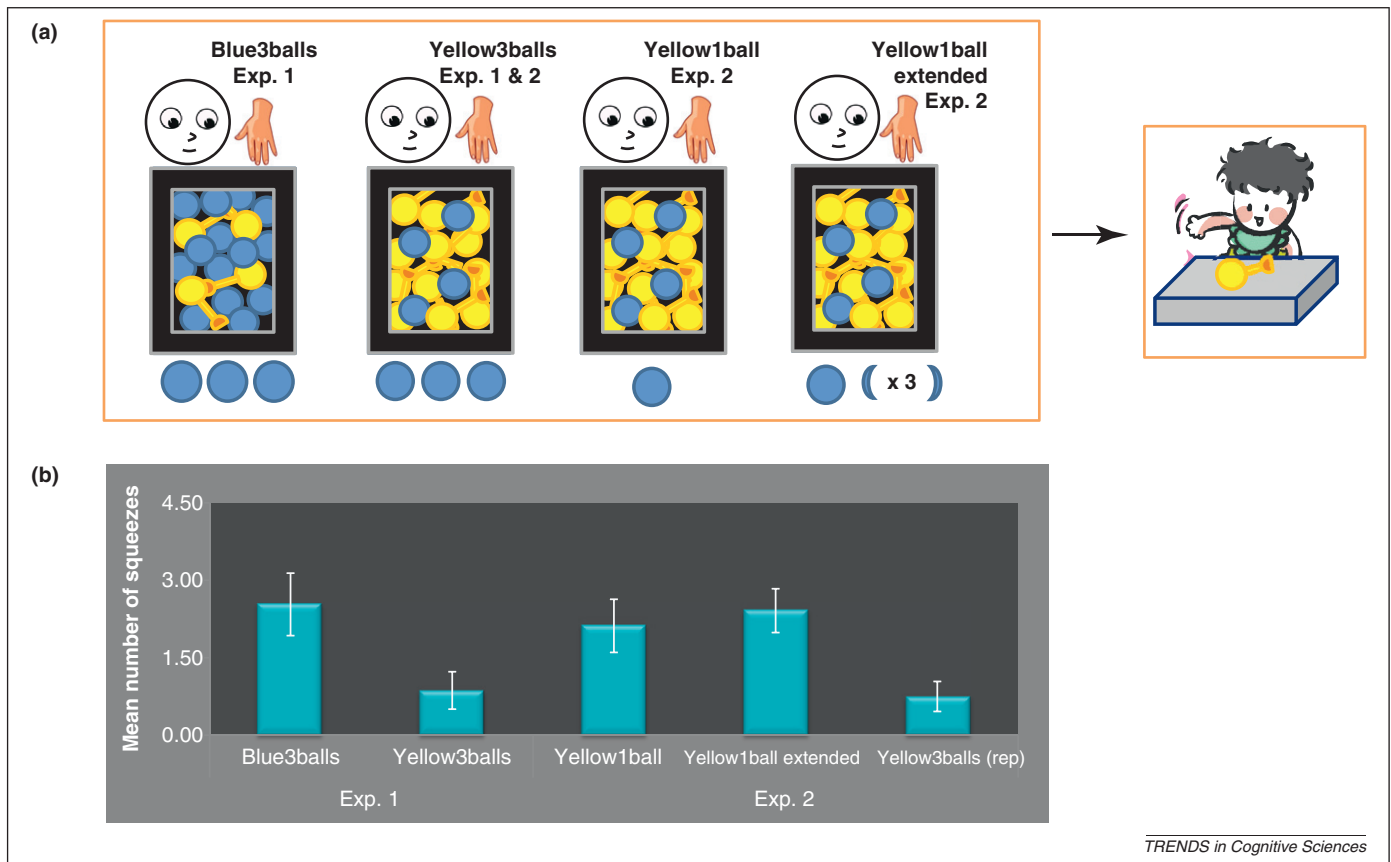


Figure 4. This experiment investigated whether infants would constrain their generalizations of object properties depending on how objects are sampled. (a) Infants (mean: 15 months, 15 days; range 13-18 months) were shown one, two, or three blue balls that squeaked, drawn either from a box containing mostly blue balls or from a box containing mostly yellow balls. Infants were then given an inert yellow ball. The prediction was that infants would be more likely to generalize the property (and thus persistently try to squeak the yellow ball) when they had observed evidence apparently randomly sampled from the population as a whole. In Experiment 1, the experimenter drew three blue balls from the mostly blue box (a sample that might have been randomly drawn from the whole box). She squeezed each ball and it squeaked. This was compared with a condition in which the experimenter did the same actions, but drew the three blue balls from a box containing mostly yellow balls (a sample unlikely to have been drawn randomly from the whole box) (b). Infants squeezed the inert yellow ball less often when the blue balls had apparently been sampled selectively (0.87 vs. 2.53 times, $t(28) = 2.45$, $p < 0.05$). Arguably, infants might be willing to generalize from majority objects to minority objects but not vice versa so in Experiment 2, only a single blue ball was drawn from the mostly yellow box and squeezed either once (Yellow1ball), or three times (Yellow1ball Extended). This condition was compared with a replication of the condition in which three blue balls were drawn from the box containing mostly yellow balls (Yellow3balls(rep)). Critically, drawing only a single yellow ball from a mostly blue box is not an improbable sample. If infants are sensitive to the sampling process, then even though they had seen more balls squeezed in the three-ball condition than the one-ball conditions, they should think that the evidence in the one-ball conditions is more likely to have been randomly sampled from the population. Thus, the prediction was that infants themselves should squeeze the ball more often in the one-ball conditions. This is what was found (0.75 squeezes in the Yellow3ball replication vs. 2.12 in the Yellow1ball condition and 2.41 in the Yellow1ball extended condition, $p < .05$ by t-test throughout) (b). In subsequent experiments (not depicted), infants' tendency to squeeze the ball was graded with respect to the data. When two blue balls were drawn from the mostly yellow box, children showed a response intermediate between the one ball and three balls condition. When behavioral cues (i.e., by turning the box upside down and letting balls fall out versus fishing around deliberately for the balls) clearly specified that the sampling process was random or selective, infants made different responses to identical data. Reproduced, with permission, from [50].

one of three non-pedagogical conditions: an interrupted condition (identical to the pedagogical condition, except that the experimenter was interrupted immediately after the demonstration), an accidental demonstration ('Whoops, look at that.'), and at baseline. In the pedagogical condition, preschoolers spent most of their free play exploring only the demonstrated function; consequently, they failed to learn other properties of the toy. By contrast, in the non-pedagogical conditions, preschoolers explored broadly [57]. Instruction is thus a 'double-edged sword': teaching promotes efficient learning by constraining the hypotheses learners consider; however, this means that learners in pedagogical contexts are less likely than learners in non-pedagogical contexts to discover uninstructed information. Similar trade-offs affect scientific inquiry. Because prior knowledge constrains the hypothesis space, experts will be less likely than novice researchers

to investigate some hypotheses (some of which may turn out to be true).

Towards a computational account of learning in early childhood

I began by noting that we can characterize epistemic practices characteristic of both science and early childhood cognition formally as well as informally (Box 1.). Recent advances in computational modeling have provided formal accounts of how abstract knowledge can both be learned and support learning from sparse data across content domains [60–65]. Some of children's inferential abilities can be captured with these models (see, e.g., [61,66], for reviews and analysis). Critically, however, there is nothing magic or exhaustive about this list of epistemic practices, precisely because there is no computational model that can generate all and only those processes necessary for

Box 2. Outstanding questions

- Computational models of many of the individual epistemic practices critical to inquiry and discovery are available, ranging from how individuals generalize from samples of data to how they optimize information gain by isolating variables. However, it is not yet possible to formally specify any process that generates all and only those epistemic practices fundamental to inquiry and discovery. Can a unified theory of exploration and learning be advanced?
- Because abilities critical to inquiry emerge in infancy and early childhood, these abilities may be universal. However, these abilities have been studied almost exclusively in populations with substantial exposure to modern technology. To what extent are the epistemic practices that characterize scientific inquiry present in the behavior of children cross-culturally?
- Most of the research on inductive inference in infancy and early childhood has emphasized continuity through development. How might well-established changes in children's information processing abilities and their world knowledge (e.g., [32,35]) be distinguished from changes in their core inferential abilities? What, if any, changes occur in inductive reasoning abilities through infancy and early childhood?
- Current approaches to hypothesis generation combine randomized trial-and-error searches with quite general principles for proposing changes to existing hypotheses (i.e., allowing any changes expressible within the grammar of a theory) [69,70]. Intuitively however, people seem to have many abilities that could constrain hypothesis generation that are not accounted for in current approaches to learning: the ability to distinguish 'good' wrong ideas from 'bad' ones (even when both ideas are wrong), the feeling of knowing when we are 'on the right track', and the ability to recognize a promising possible solution well before we know whether the solution in fact better predicts the data. These abilities might be explained by additional constraints on hypothesis generation; in particular, the ability to constrain our hypotheses based on an abstract representation of the kind of errors made by our current hypotheses [67]. Such constraints might also help learners decide which problems are potentially tractable (and thus worth devoting time to) and which problems are not. As yet, there is little empirical evidence for the idea that learners have these intuitions and no formal account of learning that takes advantage of the possibility of such higher-order constraints on hypothesis generation. Although speculative, such ideas might be promising areas for future investigation.

knowledge acquisition. Much as we want simple, elegant, unified principles of learning – Hebb's rule, Rescorla-Wagner, Bayes' law – none of them does justice to what children can do. There is as yet no algorithm for this kind of learning (Box 2). Developing a unified theory of hypothesis generation, inquiry and discovery remains a hard problem of cognitive science [67]. What we can do is provide empirical evidence that children engage in this kind of learning, setting the standard to which the next generation of theories of learning must aspire.

References

- 1 Emery, N.J. and Clayton, N.S. (2004) The mentality of crows. Convergent evolution of intelligence in corvids and apes. *Science* 306, 1903–1907
- 2 Weir, A.A.S. et al. (2002) Shaping of hooks in New Caledonian crows. *Science* 297, 981
- 3 Whiten, A. et al. (2005) Conformity to cultural norms of tool use in chimpanzees. *Nature* 437, 737–740
- 4 Gopnik, A. et al. (2001) Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Dev. Psychol.* 37, 620–629
- 5 Kushnir, T. and Gopnik, A. (2005) Young children infer causal strength from probabilities and interventions. *Psychol. Sci.* 16, 678–683
- 6 Kushnir, T. and Gopnik, A. (2007) Conditional probability versus spatial contiguity in causal learning: preschoolers use new contingency evidence to overcome prior spatial assumptions. *Dev. Psychol.* 44, 186–196
- 7 Schulz, L.E. and Gopnik, A. (2004) Causal learning across domains. *Dev. Psychol.* 40, 162–176
- 8 Sobel, D.M. and Kirkham, N.Z. (2006) Blickets and babies: the development of causal reasoning in toddlers and infants. *Dev. Psychol.* 42, 1103–1115
- 9 Sobel, D.M. et al. (2007) The blicket within: Preschoolers' inferences about insides and causes. *J. Cogn. Dev.* 8, 159–182
- 10 Gopnik, A. et al. (2004) A theory of causal learning in children: causal maps and Bayes nets. *Psychol. Rev.* 111, 1–31
- 11 Carey, S. (1985) *Conceptual Change in Childhood*, MIT Press
- 12 Gopnik, A. and Meltzoff, A.N. (1997) *Words, Thoughts, and Theories*, MIT Press
- 13 Keil, F. (1989) *Concepts, Kinds, and Cognitive Development*, MIT Press
- 14 Wellman, H.M. and Gelman, S.A. (1992) Cognitive development: foundational theories of core domains. *Annu. Rev. Psychol.* 43, 337–375
- 15 Tenenbaum, J.T. and Griffiths, T.L. (2009) Theory-based causal induction. *Psychol. Rev.* 116, 661–716
- 16 Flavell, J.H. et al. (1995) Young children's knowledge about thinking. *Monogr. Soc. Res. Child Dev.* 60, (1, Serial No. 243)
- 17 Inagaki, K. and Hatano, G. (1993) Young children's understanding of the mind-body distinction. *Child Dev.* 64, 1534–1549
- 18 Kalish, C. (1996) Preschoolers' understanding of germs as invisible mechanisms. *Cogn. Dev.* 11, 83–106
- 19 Perner, J. (1991) *Understanding the Representational Mind. Learning, Development, and Conceptual Change*, MIT Press
- 20 Shultz, T.R. (1982) Rules of causal attribution. *Monogr. Soc. Res. Child Dev.* 47, (Serial No. 194)
- 21 Spelke, E.S. et al. (1992) Origins of knowledge. *Psychol. Rev.* 99, 605–632
- 22 Carey, S. (2009) *The Origin of Concepts*, Oxford University Press
- 23 Harris, P.L. et al. (1996) Children's use of counterfactual thinking in causal reasoning. *Cognition* 61, 233–259
- 24 Sobel, D.M. and Munro, S.E. (2009) Domain generality and specificity in children's causal inference about ambiguous data. *Dev. Psychol.* 45, 511–524
- 25 Karmiloff-Smith, A. and Inhelder, B. (1975) If you want to get ahead, get a theory. *Cognition* 3, 195–212
- 26 Kohn, A.S. (1993) Preschoolers' reasoning about density: Will it float? *Child Dev.* 64, 1637–1650
- 27 Kuhn, D. et al. (1988) *The Development of Scientific Thinking Skills. Developmental Psychology Series*, Academic Press (San Diego)
- 28 Alloy, L.B. and Tabachnik, N. (1984) Assessment of covariation by humans and animals: the joint influence of prior expectations and current situational information. *Psychol. Rev.* 91, 112–149
- 29 Fischhoff, B. and Beyth-Marom, R. (1983) Hypothesis evaluation from a Bayesian perspective. *Psychol. Rev.* 90, 239–260
- 30 Koslowski, B. (1996) *Theory and evidence: the development of scientific reasoning*, MIT Press
- 31 Murphy, G.L. and Medin, D.L. (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316
- 32 Sobel, D.M. and Munro, S.A. (2009) Domain generality and specificity in children's causal inferences about ambiguous data. *Dev. Psychol.* 45, 511–524
- 33 Sobel, D.M. et al. (2004) Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cogn. Sci.* 28, 303–333
- 34 Schulz, L.E. et al. (2007) Can being scared cause tummyaches? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Dev. Psychol.* 43, 1124–1139
- 35 Bonawitz, E.B. et al. (2012) Teaching the Bayesian child: Three-and-a-half-year-olds' reasoning about ambiguous evidence. *J. Cogn. Dev.* <http://dx.doi.org/10.1080/15248372.2011.577701>
- 36 Gweon, H. and Schulz, L.E. (2011) 16-month-olds rationally infer causes of failed actions. *Science* 332, 1524
- 37 Schulz, L.E. and Sommerville, J. (2006) God does not play dice: causal determinism and children's inferences about unobserved causes. *Child Dev.* 77, 427–442

- 38 Schulz, L.E. *et al.* (2008) Judicious imitation: young children imitate deterministic actions exactly, stochastic actions more variably. *Child Dev.* 79, 395–410
- 39 Schulz, L.E. *et al.* (2008) Going beyond the evidence: abstract laws and preschoolers' responses to anomalous data. *Cognition* 109, 211–223
- 40 Bonawitz, E.B. *et al.* (2012) Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychol.* 64, 215–234
- 41 Gelman, S.A. and Markman, E.M. (1986) Categories and induction in young children. *Cognition* 23, 183–209
- 42 Gelman, S.A. and Markman, E.M. (1987) Young children's inductions from natural kind: The role of categories and appearances. *Child Dev.* 58, 1531–1541
- 43 Schulz, L.E. *et al.* (2008) Word, thought, and deed: The role of object labels in children's inductive inferences and exploratory play. *Dev. Psychol.* 44, 1266–1276
- 44 Butler, L.P. and Markman, E.M. (2012) Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Dev.* <http://dx.doi.org/10.1111/j.1467-8624.2012.01775.x>
- 45 Legare, C.H. (2012) Exploring exploration: explaining inconsistent information guides hypothesis-testing behavior in young children. *Child Dev.* 83, 173–185
- 46 Legare, C.H. *et al.* (2010) Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Dev.* 81, 929–944
- 47 Schulz, L.E. and Bonawitz, E.B. (2007) Serious fun: preschoolers play more when evidence is confounded. *Dev. Psychol.* 43, 1045–1050
- 48 Cook, C. *et al.* (2011) Where science starts: spontaneous experiments in preschoolers' exploratory play. *Cognition* 120, 341–349
- 49 Sodian, B. *et al.* (1991) Young children's differentiation of hypothetical beliefs from evidence. *Child Dev.* 62, 753–766
- 50 Gweon, H. *et al.* (2010) Infants consider both the sample and the sampling process in inductive generalization. *Proc. Natl. Acad. Sci. U.S.A.* 107, 9066–9071
- 51 Denison, S. *et al.* (2012) The emergence of probabilistic reasoning in very young infants: Evidence from 4.5- and 6-month-old infants. *Dev. Psychol.* <http://dx.doi.org/10.1037/a0028278>
- 52 Xu, F. and Garcia, V. (2008) Intuitive statistics by 8-month-old infants. *Proc. Natl. Acad. Sci. U.S.A.* 105, 5012–5015
- 53 Xu, F. and Denison, S. (2009) Statistical inference and sensitivity to sampling in 11-month-old infants. *Cognition* 112, 97–104
- 54 Kushnir, T. *et al.* (2010) Young children use statistical sampling to infer the preferences of others. *Psychol. Sci.* 21, 1134–1140
- 55 Ma, L. and Xu, F. (2011) Young children's use of statistical sampling evidence to infer the subjectivity of preferences. *Cognition* 120, 403–411
- 56 Buchsbaum, D. *et al.* (2011) Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition* 120, 331–340
- 57 Bonawitz, E.B. *et al.* (2011) The double-edged sword of pedagogy: teaching limits children's spontaneous exploration and discovery. *Cognition* 120, 322–330
- 58 Kushnir, T. *et al.* (2008) The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition* 107, 1084–1092
- 59 Shafto, P. and Goodman, N. (2008) Teaching games: statistical sampling assumptions for pedagogical situations. In *Proceedings of the Cognitive Science Society* (Love, B.C., McRae, K. and Sloutsky, V.M., eds), pp. 1632–1637, Cognitive Science Society
- 60 Tenenbaum, J.B. *et al.* (2006) Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318
- 61 Tenenbaum, J.B. *et al.* (2011) How to grow a mind: statistics, structure, and abstraction. *Science* 331, 1279–1285
- 62 Kemp, C. and Tenenbaum, J.B. (2008) The discovery of structural form. *Proc. Natl. Acad. Sci. U.S.A.* 105, 10687–10692
- 63 Kemp, C. *et al.* (2007) Learning overhypotheses with hierarchical Bayesian models. *Dev. Sci.* 10, 307–321
- 64 Goodman, N.G. *et al.* (2011) Learning a theory of causality. *Psychol. Rev.* 118, 110–119
- 65 Dewar, K.M. and Xu, F. (2010) Induction, overhypothesis, and the origin of abstract knowledge. *Psychol. Sci.* 21, 1871–1877
- 66 Gopnik, A. and Wellman, H. (2012) Reconstructing constructivism: causal models, Bayesian learning mechanisms and the theory theory. *Psychol. Bull.* <http://dx.doi.org/10.1037/a0028044>
- 67 Schulz, L.E. (2012) Finding new facts; thinking new thoughts. In *Rational constructivism, Advances in Child Development and Behavior* 42 (Xu, F. and Kushnir, T., eds), Elsevier
- 68 Goodman, N. (1946) A query on confirmation. *J. Philos.* 43, 383–385
- 69 Bonawitz, E.B. *et al.* (2012) Rational randomness: the role of sampling in an algorithmic account of preschooler's causal learning. In *Rational constructivism, Advances in Child Development and Behavior* 42 (Xu, F. and Kushnir, T., eds), Elsevier
- 70 Ullman, T.D. *et al.* (2010) Theory acquisition as stochastic search. In *Proceedings of the Cognitive Science Society* (Camtrabone, R. and Ohlsson, S., eds), pp. 2840–2845, Cognitive Science Society