

Scientific Thinking in Young Children: Theoretical Advances, Empirical Research, and Policy Implications

Alison Gopnik

New theoretical ideas and empirical research show that very young children's learning and thinking are strikingly similar to much learning and thinking in science. Preschoolers test hypotheses against data and make causal inferences; they learn from statistics and informal experimentation, and from watching and listening to others. The mathematical framework of probabilistic models and Bayesian inference can describe this learning in precise ways. These discoveries have implications for early childhood education and policy. In particular, they suggest both that early childhood experience is extremely important and that the trend toward more structured and academic early childhood programs is misguided.

Thirty years ago, the idea that 2-year-olds think like scientists would have seemed absurd. Jean Piaget, the great pioneer of cognitive development, claimed that preschoolers' thinking was just the opposite of scientific thinking. Preschoolers were irrational, illogical, "pre-causal," and limited to the here and now (1). These ideas informed both education and policy.

These claims have turned out to be wrong. Several waves of empirical work have shown that even infants and very young children have intuitive theories of the world around them. More recently, mathematical models of learning have been developed. Empirical research informed by those models shows that early learning is also remarkably similar to scientific induction (Fig. 1).

During the 1980s and 1990s, researchers discovered that very young children have abstract, structured, coherent, causal representations of the world around them—representations that are similar to scientific theories. They use those representations to make wide-ranging new predictions. These representations appear to be in place even in infancy, but it is particularly clear that preschoolers have intuitive theories of the physical, biological, psychological, and social world (2–4).

New methods led to this first revolution in our understanding of development. The advent of video recording, and striking experimental ingenuity, led to a flood of results that showed sophisticated knowledge in even the youngest infants. By studying what babies looked at, reached for, or imitated, researchers could show that even infants understand both physical objects and other people [e.g., (3, 5)]. Piaget tried to assess preschool children's knowledge by asking them open-ended questions about hypothetical scenarios. But preschoolers show much more

sophisticated knowledge when they respond to focused questions about real-life examples (2–4).

This research showed that young children's knowledge is structurally similar to scientific theories, but not necessarily that children learn like scientists. It could be that much of this knowledge is innate rather than learned—evolutionarily determined rather than inferred from experience. Moreover, until quite recently, there were few theoretical accounts that could encompass learning mechanisms in both childhood and science, or empirical results that showed those mechanisms were similar. In fact, the predominant theories of learning emphasized complex associations be-

tween stimuli. Associative learning appears to be very different from the hypothesis testing and experimentation of science.

In the past 10 years, however, theoretical and empirical research has begun to show that children's learning mechanisms do indeed resemble the basic inductive processes of science. We now have a more precise and formal theory of children's learning mechanisms, derived from ideas about probabilistic models and Bayesian learning methods that originated in computer science, statistics, and philosophy of science.

Probabilistic Models

Philosophy of science, artificial intelligence, and developmental psychology all face the same fundamental dilemma. As adults, we seem to have highly structured, abstract, coherent knowledge of the world around us. This knowledge allows us to make wide-ranging predictions and inferences. But we also seem to learn that highly structured knowledge from the contingent, concrete, probabilistic evidence of our senses. How can this be? Traditionally, philosophers and psychologists have responded to this dilemma in two ways. "Nativists" have argued that this abstract structure must be in place innately because it could not possibly be learned. "Empiricists" have argued that this abstract structure is illusory; in reality there are only specific learned associations between particular pieces of evidence.

The probabilistic models approach [e.g., (6–9)] addresses this dilemma in a new way. Imagine that there is some real structure in the world—a spatial configuration, a grammar, or a network of causal relationships. That structure gives rise



Fig. 1. Child's play is science. ["Playing Doctors" by Frederick Daniel Hardy (1827–1911); image: Stapleton Collection/Corbis]

Department of Psychology, University of California, Berkeley, CA 94720, USA. E-mail: gopnik@socrates.berkeley.edu

to some patterns of evidence rather than others—a particular set of visual images, or spoken sentences, or statistical contingencies between events. That spatial or grammatical or causal structure can be represented mathematically by a generative model, such as a map or tree structure or a graphical network. The representation is a hypothesis about what the actual structure is like. This hypothesis can be precisely described in formal ways. The representation is generative, which means that it will allow you to mathematically compute the patterns of evidence that follow from that structure and then make new inferences accordingly: A particular map will let you predict how to reach a location by a new route; a particular grammatical tree will let you predict whether a new sentence will be acceptable; a particular causal graph will let you predict whether a new event will be followed by other events. If the hypothesis is correct, then these predictions will turn out to be right.

These generative models, then, can describe representations of the world and explain how those representations allow us to make a wide range of new inferences. Critically, the systematic link between structure and evidence in these models also allows you to reverse the process and to make inferences about the nature of the structure from the evidence it generates. It lets you decide which map or tree or causal graph best accounts for the evidence, and so leads you to adopt the most likely hypothesis.

The idea that mental models of the structure of the world generate predictions, and that we can invert that process to learn the structure from evidence, is not itself new. The big advance has been integrating ideas about probability into that basic framework. Typically, a great many hypotheses are, in principle, compatible with any pattern of evidence, so how can we decide on the best one? Integrating probability theory makes this learning problem more tractable. Although many hypotheses may be compatible with the evidence, some hypotheses will be more or less likely to have generated the evidence than others.

One of the most powerful and general ways to solve the learning problem is to use Bayesian inference. If we know the prior probability of a hypothesis, and a generative model tells us the likelihood of the evidence given the hypothesis, then when we observe a new pattern of evidence, we can use Bayes' rule to determine the probability that the hypothesis is true given that evidence. Rather than simply generating a yes-or-no decision about whether a particular hypothesis is true, the probabilistic Bayesian learning algorithms consider multiple hypotheses and assign probabilities to those hypotheses. Bayesian methods let you determine the probability of possibilities.

Bayesian ideas have been successfully applied to a wide range of problems, including vision and motor control (10, 11). This kind of perceptual and motor learning may not appear to resemble scientific learning. But probabilistic models have also been applied to precisely the kinds of knowl-

edge that we see in scientific and intuitive theories. In particular, causal knowledge is central to both kinds of theories. Causal graphical models or "Bayes nets," developed in the philosophy of science and computer science, provide a particularly powerful and successful account of causal knowledge and learning (6, 12, 13). Algorithms that use Bayes nets allow computers to actually do some kinds of science, such as discovering the causal structure of weather systems, gene expression, or brain function from data.

Most recently, this work has been expanded to allow for formal representations of more abstract

higher-order causal structure—for example, the general framework principles that prevail in a scientific paradigm and that shape particular causal hypotheses (14). There has also been work on an even more general "probabilistic logic" that can encode a much wider range of relationships, including spatial and logical as well as causal ones. At least in principle, this logic allows a wide range of generative models to be learned from probabilistic data (15).

Unlike traditional nativism, the probabilistic models approach gives us a way to actually infer abstract hierarchical structure from data, at least




Fig. 2. Schematic representation of the ping-pong ball experiment. The experimenter showed the infants a box full of white and red balls. Then she closed her eyes and randomly took some balls from the box and put them in another small bin. If the sample was truly random, then the distribution of balls in the bin should match the distribution of the balls in the box. Infants saw a sample that either matched or did not match the distribution, and they looked longer at the nonmatching sample. In a control condition, infants saw just the same sequence of events, but the experimenter took the balls out of her pocket rather than taking them from the box, and the looking-time difference disappeared.

in principle. If children learned in this way, they could drastically revise their representations of the world on the basis of their experience, as scientists do. They would not be limited to making small adjustments to innately determined representations. Unlike traditional empiricism, the approach proposes that children, also like scientists, never start from a completely blank slate or with completely pure data. Instead, from the very beginning, they would be testing hypotheses and assessing the data in the light of those hypotheses.

Children as Scientific Learners

Do children actually learn about the world in this way? Over the past 10 years, researchers have systematically given young children patterns of evidence about the world and then observed the conclusions that they draw [e.g., (7) and articles in (16, 17); for an extensive review and tutorial, see (18)]. To a striking extent, children use data to formulate and test hypotheses and theories in much the same way that scientists do. Scientists learn about the world in three ways: They analyze statistical patterns in the data, they do experiments, and they learn from the data and ideas of other scientists. The recent studies show that children also learn in these ways and that they often resemble ideal Bayesian learners. Probabilistic models make accurate and detailed predictions about children's learning.

Statistics. Anyone who has ever taught a methods course knows that adults have a hard time

explicitly understanding statistics. It may be surprising, then, that even very young infants can implicitly reason statistically. The first wave of these experiments showed that even young infants are sensitive to statistical patterns [e.g., (19)]. More recently, researchers have shown that infants and young children not only detect statistical patterns, they use those patterns to test causal hypotheses about people and things.

For example, Xu and Garcia (20) demonstrated that 8-month-olds were sensitive to statistical sampling patterns. They used a "looking-time" technique that has been extensively used to study infant cognition. It depends on the fact that infants look longer at unexpected events. When the experimenter took a sample of mostly red ping-pong balls from a box of mostly white balls, infants looked longer than when she took a sample of mostly red balls from a box of mostly red balls (Fig. 2).

Note that the unlikely events in this experiment were not impossible; you could, after all, pull mostly red balls from a box of mostly white balls. The events were merely improbable if your causal model of the event assumed that the balls in the bin were a random sample. It's as if the infants said to themselves, "Aha! Less than 0.05 probability that this occurred by chance!" But would the surprising evidence drive the children to a new causal model?

Kushnir *et al.* (21) found that it would. In fact, children as young as 20 months interpreted non-

random sampling psychologically. An experimenter took frogs from a box of all frogs or she took frogs from a box of almost all ducks. Then she left the room and another experimenter gave the child a small bowl of frogs and a separate bowl of ducks. When the original experimenter returned, she extended her hand ambiguously between the bowls. The children could give her either a frog or a duck. When she had taken frogs from a box of all frogs, children were equally likely to give her a frog or a duck. When she had taken frogs out of the box that was almost all ducks, children gave her a frog. In the first case, the children concluded that she had merely drawn a random sample from the box, but in the second case they concluded that she had displayed a preference for frogs. Thus, children less than 2 years old had inferred an underlying mental state—a preference—from a statistical pattern.

In another line of research, my colleagues and I designed a simple test to see whether young children would appropriately infer physical causal relationships from statistical evidence about covariation (7, 22). We showed children a "blicket detector"—a box that plays music when you put some objects on it but not others—and then showed them various patterns of statistical dependence between the objects and the effect. Then we asked children to make the machine go or turn it off. Figure 3 shows one such experiment. Based on the child's prior knowledge about the machine, it could have any of the causal structures represented in Fig. 3. We found that 2-, 3-, and 4-year-olds could use the pattern of covariation between the blocks and the machine's activation to infer which of these causal structures was correct, and so to make the machine go or stop. Other studies show that toddlers as young as 24 months can make these inferences even when the statistical pattern is more complicated (23). One recent study shows that even 16-month-olds can use covariation to infer causation in this way (24).

Schulz *et al.* (25) showed that 4-year-old children could also use statistical dependencies to infer more complex causal structures. Children saw a simple machine with a switch on one side and two disks that spun on top. Even this simple machine could work in many different ways (the switch could make the blue disk go, which could make the yellow disk go; the switch could make both disks go; etc.). Preschoolers used evidence correctly to distinguish between causal chain structures (the switch makes the blue disk go, which makes the yellow disk go), common cause structures (the switch makes both disks go), and conjunctive cause structures (the switch and the blue disk are both necessary to make the yellow disk go).

Bayesian inference considers both new evidence and the prior probability of hypotheses. This gives Bayesian learning a characteristic combination of stability and flexibility. In science, we hold on to well-confirmed hypotheses, but enough new evidence can eventually overturn even the




Fig. 3. The blicket detector experiment. Children saw that the machine did not activate when B alone was placed on it, but did activate when A was placed on it and continued to do so when B was added to A. Then they were asked to make the machine stop. Given this evidence, the correct causal interpretation is that A alone activates the machine, and the children should act on A and not B.

most cherished idea. Several recent studies show that children integrate prior knowledge and new evidence, too. For example, 4-year-olds begin by thinking that psychological causes (e.g., being anxious) are unlikely to cause physical effects (e.g., having a stomach ache) and reject evidence to the contrary. But if you give them accumulating evidence in favor of this “psychosomatic” hypothesis, they gradually become more and more likely to accept that initially unlikely idea (26), and a Bayesian model can predict this change quite precisely.

Children also use statistics to infer the existence of unobserved causes—hidden “theoretical entities.” Gopnik *et al.* (7) found that when the observed variables couldn’t explain the evidence, children would look for unobserved variables instead, in a way that could be predicted by Bayes nets. Schulz and Sommerville (27) found that when children saw a “blicket detector” that went off only 2 of 6 times, they inferred that some hidden variable was responsible for the failures.

Finally, we can ask whether children are restricted to making specific inferences about particular causal relationships or whether, like scientists, they can also make inferences about broader “framework principles”—general theoretical ideas or “paradigms.” An exciting development in the computational world has been the discovery that more abstract theoretical laws can actually sometimes be learned more quickly than the specific causal hypotheses they subsume (28). Two recent studies also suggest that preschoolers can make these broader generalizations swiftly and appropriately (29, 30).

Experiments. Anyone who watches young children has seen how they ceaselessly fiddle with things and observe the results. Children’s play can look like experimentation. Recent research by Schulz and colleagues shows that children’s exploratory play does indeed involve a kind of intuitive experimentation. Children’s play is not as structured as the ideal experiments of institutional science; even adults can have a hard time designing ideally controlled experiments [e.g., (31)]. However, recent formal work in the philosophy of science has shown that much less systematic experimentation can yield a remarkable amount of causal knowledge (32). The empirical research shows that play is sufficiently systematic to help children discover causal relationships.

For example, Cook *et al.* (33) performed a variant of the “blicket detector” experiments using “pop-beads,” small plastic beads that could be hooked together to make larger units. First, the experimenter put individual beads on the machine. One group of 4-year-olds saw that some of the beads made the machine go and some didn’t. A second group saw that all the beads made the machine go. Then, the experimenter simply gave the children the machine and two new beads that were hooked together and let them play.

The “some beads” condition sets up a causal problem for the children: Which beads make the machine go? To solve that problem, you need to

test each bead by itself. The “all beads” condition does not; children can assume that both beads will make the machine go. Sure enough, children spontaneously pulled the beads apart and tested them separately in their play in the “some beads” condition but not in the otherwise identical “all beads” condition.

Similarly, Legare (34) showed 4-year-olds that red blocks made a “blicket detector” machine go and then showed them an anomaly—a red block that failed. She asked them “Why did that happen?” and let them play with the machine. Children systematically played with the machine in ways that tested the hypotheses that they expressed in their explanations. Another recent study shows that pretend play is also closely related to counterfactual reasoning—a particularly sophisticated type of causal inference (35).

These results indicate that when young children face a causal puzzle, they try to solve that puzzle in their spontaneous play. Children’s actions ensure that they receive causally relevant and informative evidence. Once that evidence is generated through play, children can use it to make the correct causal inferences.

Learning from Others

The picture of the child as a “little scientist” has sometimes been taken to imply that children are solitary learners. Of course, real science is a highly social endeavor, and scientists must constantly interpret the demonstrations and reports of other people. Children can also learn about causal relationships by watching what other people do and what happens as a result. In our lab (36), 4-year-olds saw an experimenter perform five different sequences of three actions on a toy, which activated or did not activate on each trial. A statistical analysis of the data would suggest that only the last two actions were necessary to activate the toy. When children got the toy, they often produced just the two relevant actions, rather than imitating everything that the experimenter did.

Moreover, very young children and even infants are sensitive to the intentions of others, particularly their intention to teach, and may draw different conclusions from the evidence that teachers give them than from the evidence they gather themselves. Bayesian models can incorporate pedagogical information by assuming that teachers provide a different and more informative sample of evidence than one would get from a random sample (37).

We (36) did exactly the same statistical imitation experiment but now included pedagogical information: The experimenter said “Here’s my toy, I’m going to show you how it works.” In this condition, children were much more likely to assume that everything the adult did was causally effective and to imitate all her actions. A Bayesian model made quite precise quantitative predictions about what the children would do in the pedagogical and nonpedagogical context.

Similarly, Bonawitz *et al.* (38) gave children a complicated toy to explore. The toy had four

tubes, each of which did something different (one lit up, one made a squeaking sound, etc.). In one condition, children saw the experimenter accidentally bump against the toy, setting off one of the squeaky tubes. Then she simply left the child alone to play with the toy. The children imitated the squeak but also discovered all the other things that the toy could do. In another condition, the experimenter introduced the toy by saying “Here is my toy” and then made it squeak. Like the children in the imitation experiment, children in this condition simply repeated what the experimenter did, and didn’t explore the machine’s other possibilities.

These new studies, and many similar ones, suggest that children as well as scientists learn in ways that are well described by probabilistic models. This research also raises myriad new problems and exciting directions for further work. How are these abstract computations actually implemented in detail by limited human minds and, ultimately, how are they instantiated in human brains? Children and scientists often seem to develop radically new hypotheses. Where do these hypotheses come from? Are children simply learners with less experience—in Bayesian terms, do they simply have a different prior? Or do they learn in ways that are Bayesian but are qualitatively different from adult learning? For example, do they search a wider space of hypotheses than adults do (29)? How is this learning influenced by the development of explicit symbol systems, from language itself to the complex mathematical notation of physics? What is the ideal balance between individual discovery and learning from others, and does this balance shift in different educational and scientific contexts? The new theoretical ideas and experimental methods give us a framework for asking and answering these questions.

Implications for Policy

So new theoretical work lets us describe both scientific learning and children’s learning in a newly systematic and rigorous way. New empirical work shows that young children learn from statistics, experiments (i.e., play) and from the actions of others in much the same way that scientists do. What does all this mean for education and policy?

First, this work provides an explanatory foundation for the demonstrable impact of high-quality preschool and care for children on later life (39). Very young children are spontaneously and pervasively learning from experience. Moreover, the Bayesian picture provides a better model for these effects than thinking about early childhood as an irreversible “critical period.” Without the right sequence of evidence, theoretical advances will be delayed or may never emerge at all. Conversely, when particular hypotheses are especially well confirmed early in life (hypotheses, for example, that expressing distress causes caregivers to turn away, or that threat leads to violence), it may be much more difficult for them to be revised later on. Even the most entrenched dysfunctional

models can, however, be overturned with sufficient evidence.

This work also should raise serious red flags about recent pressure, both from parents and policy-makers, to make preschools more structured and academic—more like schools. This research is new, so it's not surprising that early childhood policy-makers still routinely hold an outdated view of development. According to this view, early childhood is about "socioemotional" development, in contrast to "cognitive skills," which are identified with later scholastic abilities such as reading and calculating. Policy-makers may acknowledge the importance of the "socioemotional" aspect, but they systematically underestimate the intellectual capabilities of preschoolers. The new research shows that even very young children are deeply engaged in such profoundly cognitive work as hypothesis testing and causal inference. This work is more cognitively challenging, in fact, than much school work.

Moreover, the research has begun to demonstrate scientifically what most preschool teachers feel intuitively. Children's spontaneous exploratory and pretend play is designed to help them learn. And pedagogy can be a mixed blessing. Even preschoolers know when they are being taught, and quickly take on information from teachers. But explicit teaching can also narrow the range of hypotheses that children are willing to consider. Activities such as encouraging play, presenting anomalies, and asking for explanations prompt scientific thinking more effectively than direct instruction.

Finally, the new research suggests that our everyday thinking and learning is strikingly continuous with scientific thinking and learning. Of course, formal scientific thinking involves a level of self-conscious reflection, including reflection on the very process of science itself. We don't see this reflection in very young children: The preschoolers see probabilistic evidence and revise hypotheses, but they don't necessarily know that that is what they are doing—nor indeed do ordinary adults. Nonetheless, we should be able to exploit the fact that very young children are natural scientists in action to help them under-

stand the principles of formal science. Even university students can understand statistics much more effectively when the material is presented in the context of everyday causal inference (40). Ordinary adults might also learn scientific concepts more effectively through play, experimentation, and observation than through pedagogy.

The new work, then, provides a scientific foundation for a long tradition of "inquiry-based" science education. But our new understanding of children's intuitive science ought to help us go beyond just a general emphasis on active inquiry. Instead, it could lead us to much more specific and scientifically supported proposals for education. Science itself could help turn young children's natural curiosity and brilliance into better science teaching and learning.

References and Notes

- J. Piaget, *The Child's Conception of the World* (Routledge, London, 1929).
- S. Carey, *The Origin of Concepts* (Oxford Univ. Press, New York, 2009).
- A. Gopnik, A. N. Meltzoff, *Words, Thoughts, and Theories* (MIT Press, Cambridge, MA, 1997).
- H. M. Wellman, S. A. Gelman, *Annu. Rev. Psychol.* **43**, 337 (1992).
- E. S. Spelke, K. Breinlinger, J. Macomber, K. Jacobson, *Psychol. Rev.* **99**, 605 (1992).
- C. N. Glymour, *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology* (MIT Press, Cambridge, MA, ed. 1, 2001).
- A. Gopnik et al., *Psychol. Rev.* **111**, 3 (2004).
- A. Gopnik, L. E. Schulz, Eds., *Causal Learning: Psychology, Philosophy, and Computation* (Oxford Univ. Press, Oxford, 2007).
- J. B. Tenenbaum, C. Kemp, T. L. Griffiths, N. D. Goodman, *Science* **331**, 1279 (2011).
- D. Kersten, P. Mamassian, A. Yuille, *Annu. Rev. Psychol.* **55**, 271 (2004).
- D. M. Wolpert, *Hum. Mov. Sci.* **26**, 511 (2007).
- J. Pearl, *Causality* (Cambridge Univ. Press, Cambridge, 2000).
- P. Spirtes, C. Glymour, R. Scheines, *Causation, Prediction, and Search* (MIT Press, Cambridge, MA, ed. 2, 2000).
- T. L. Griffiths, J. B. Tenenbaum, *Psychol. Rev.* **116**, 661 (2009).
- N. D. Goodman, V. K. Mansinghka, D. M. Roy, K. Bonawitz, J. B. Tenenbaum, in *Proceedings of the 24th Conference in UAI*, D. A. McAllester, P. Myllymäki, Eds. (AUAI Press, Helsinki, 2008), pp. 220–229.
- A. Gopnik, J. B. Tenenbaum, *Dev. Sci.* **10**, 281 (2007).
- F. Xu, T. Griffiths, Eds., special issue on Probabilistic Models of Cognitive Development. *Cognition* **120**, 299 (2011).
- A. Gopnik, H. Wellman, *Psychol. Bull.*, 10.1037/a0028044 (2012).
- J. R. Saffran, R. N. Aslin, E. L. Newport, *Science* **274**, 1926 (1996).
- F. Xu, V. Garcia, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 5012 (2008).
- T. Kushnir, F. Xu, H. M. Wellman, *Psychol. Sci.* **21**, 1134 (2010).
- A. Gopnik, D. M. Sobel, L. E. Schulz, C. Glymour, *Dev. Psychol.* **37**, 620 (2001).
- D. M. Sobel, N. Z. Kirkham, *Dev. Psychol.* **42**, 1103 (2006).
- H. Gweon, L. E. Schulz, *Science* **332**, 1524 (2011).
- L. E. Schulz, A. Gopnik, C. Glymour, *Dev. Sci.* **10**, 322 (2007).
- L. E. Schulz, E. B. Bonawitz, T. L. Griffiths, *Dev. Psychol.* **43**, 1124 (2007).
- L. E. Schulz, J. Sommerville, *Child Dev.* **77**, 427 (2006).
- N. D. Goodman, T. D. Ullman, J. B. Tenenbaum, *Psychol. Rev.* **118**, 110 (2011).
- C. G. Lucas, A. Gopnik, T. L. Griffiths, in *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, S. Ohlsson, R. Catrambone, Eds. (Cognitive Science Society, Austin, TX, 2010), pp. 28–52.
- L. E. Schulz, N. D. Goodman, J. B. Tenenbaum, A. C. Jenkins, *Cognition* **109**, 211 (2008).
- L. Kuhn, *Education for Thinking* (Harvard Univ. Press, Cambridge, MA, 2005).
- F. Eberhardt, R. Scheines, *Philos. Sci.* **74**, 981 (2007).
- C. Cook, N. D. Goodman, L. E. Schulz, *Cognition* **120**, 341 (2011).
- C. H. Legare, *Child Dev.* **83**, 173 (2012).
- D. Buchsbaum, S. Bridgers, D. Skolnick Weisberg, A. Gopnik, *Philos. Trans. R. Soc. London Ser. B* **367**, 2202 (2012).
- D. Buchsbaum, A. Gopnik, T. L. Griffiths, P. Shafto, *Cognition* **120**, 331 (2011).
- P. Shafto, N. D. Goodman, M. C. Frank, *Perspect. Psychol. Sci.* **7**, 341 (2012).
- E. Bonawitz et al., *Cognition* **120**, 322 (2011).
- J. J. Heckman, *Science* **312**, 1900 (2006).
- R. Scheines, M. Easterday, D. Danks, in *Causal Learning: Psychology, Philosophy, and Computation*, A. Gopnik, L. Schulz, Eds. (Oxford Univ. Press, New York, 2007), pp. 119–138.

Acknowledgments: Supported by NSF grant BCS-1023875 and by the James S. McDonnell Causal Learning Collaborative, which also supported almost all the ideas and research reported here, both by myself and others. I acknowledge and deeply thank the other members of the Collaborative and the James S. McDonnell Foundation itself.

10.1126/science.1223416