Advanced Review

Causal learning in children

David M. Sobel[1*] and Cristine H. Legare[2]

How do children learn the causal structure of the environment? We first summarize a set of theories from the adult literature on causal learning, including associative models, parameter estimation theories, and causal structure learning accounts, as applicable to developmental science. We focus on *causal graphical models* as a description of children's causal knowledge, and the implications of this computational description for children's causal learning. We then examine the contributions of explanation and exploration to causal learning from a computational standpoint. Finally, we examine how children might learn causal knowledge from others and how computational and constructivist accounts of causal learning can be integrated. © 2014 John Wiley & Sons, Ltd.

How to cite this article:
WIREs Cogn Sci 2014. doi: 10.1002/wcs.1291

INTRODUCTION

Children are remarkable causal learners. Despite the fact that traditional cognitive development research has suggested that young children are 'precausal',[1] contemporary accounts of cognitive development have demonstrated that young children have sophisticated domain-specific causal reasoning abilities. Infants register particular aspects of physical causality.[2–5] Toddlers recognize various causal relations in the psychological domain, especially about others' desires and intentions.[6] Preschoolers understand that biological and psychological events can rely on nonobvious, hidden causal relations.[7,8]

More generally, preschoolers also display a variety of domain-general causal reasoning abilities. Young children recognize the importance of Hume's principles—temporal priority, spatial priority, and contingency—in making judgments about causal relations.[9,10] Preschoolers also possess predictive,[11] explanatory,[12,13] and counterfactual reasoning abilities.[14,15] But what explains the process of causal learning? How do children take the specific and often sparse data they observe and construct abstract representations of causal knowledge?

*Correspondence to: Dave_Sobel@brown.edu

[1]Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI, USA

[2]Department of Psychology, The University of Texas at Austin, Austin, TX, USA

Conflict of interest: The authors have declared no conflicts of interest for this article.

We will attempt to answer this question in three ways. The first will be to summarize a set of general theories of causal learning, which have been developed in adult cognitive psychology. This section will conclude with a discussion of the application of these theories to developmental science. The second will be to consider a particular set of issues concerning the relation between the explanations children hear and generate and their exploration of the environment. This section will conclude with a set of discussion questions attempting to integrate how children explore and explain the environment with descriptions of the way they may engage in causal learning. Third, we will examine how children learn causal knowledge from others, particularly focused on the idea that such a learning process might be best explained by general theories of causal learning.

THEORIES OF CAUSAL LEARNING AND THEIR APPLICATION TO COGNITIVE DEVELOPMENT

To describe how children engage in causal inference, we must consider both the existing domain-specific knowledge they possess about how causal relations work (called 'substantive principles' for causal learning, see Ref 11) and the more domain-general mechanisms by which children acquire new causal knowledge from information in the environment ('formal principles' for causal learning, see Ref 11). Children's substantive principles take the form of

© 2014 John Wiley & Sons, Ltd.

content information that applies to the causal information they are learning. Such knowledge could also be quite general regarding the inferences it licenses (e.g., a broad piece of substantive knowledge is that children understand is temporal priority—that causes precede their effects). Other pieces of substantive knowledge might specify only certain kind of causal inferences (e.g., children might know that plants grow without knowing a more general conceptual structure about plants, such as that they are alive). This content knowledge is usually acquired through the accretion of facts and information, and could differ across domains and situations (i.e., temporal priority might be more general than inferences about plant behavior). Critically, substantive knowledge often allows children to make better causal inferences when mechanisms are familiar to them (see e.g., Ref 16).

Describing how children learn causal knowledge (including the substantive principles they possess for any domain) usually involves describing the formal principles, or domain-general mechanisms for causal learning. In the next section, we outline several descriptions of such formal mechanisms for causal learning. Critically, no formal learning mechanism is independent of the existing knowledge (i.e., the substantive principles) that the child possesses. Our goal is to review how such formal mechanisms might interact with children's existing knowledge to formulate a description of causal learning.

Theories Based on Associative Strength and Parameter Estimation

One possible way children engage in causal inference is to simply associate causes and effects, in the same way that animals associate conditioned and unconditioned stimuli in classical conditioning.[17,18] On this view, there is nothing to understanding causality beyond recognizing associations. These descriptions of causal learning assume that candidate causes and effects have been identified (typically based on relatively low-level substantive information such as temporal priority), and output the strength of each cause–effect association.

As such models only output associative information, it is hard to determine how they make predictions about the ways in which people use causal knowledge or act on the world. Thus, contemporary advocates of this approach have suggested that causal learning and inference takes place by translating associative information into a measure of causal structure. That measure of causal strength might then be combined with substantive information to make causal inferences or generate new interventions.[19,20] For instance, building on the model described in Ref,[18] there are various accounts of causal learning that calculate the causal strength of known causal relations from associative information.[21–23] More recently, similar models of semantic cognition (including causality) have been proposed based on neural network architectures.[24,25]

A related approach to causal learning relies on estimating causal parameters based on the frequency with which events co-occur. Two prominent proposals in this category are the ΔP model[26–28] and the Power PC model[29] and extensions for interactive causality (i.e., when two causes must combine to produce an effect, as opposed to being just additive, see Ref 30). These models calculate an estimate of the maximal likelihood value of the strength of a presumed causal relation given a set of data.[31,32a] How these presumed causal relations are determined is typically a function of prior substantive knowledge (e.g., 'focal sets', see Ref 34).

Are such mechanisms plausible as developmental accounts? Associative learning mechanisms are available to children at very early ages in the form of their statistical learning capacities.[35,36] Such statistical learning is implicated in infants' processing of causal data,[37] social knowledge,[38] and linguistic information.[39,40]

Statistical learning capacities, in turn, are related to infants' ability to generalize. Eight-month-olds register the appropriate generalizations from population to sample and from samples to populations.[41] Such statistical learning mechanisms can lead to broader inferences. For instance, nonrandom sampling leads 18–24-month-old infants to infer that the individual has a subjective preference for those objects.[42,43] Similarly, by age 2, children can use the regularity in other people's choices to generalize their preferences to categories of objects.[44] Although it is possible that causal knowledge can be learned from registering associative or frequency information among events, it is not clear *how* such statistical knowledge is integrated with children's existing substantive knowledge. For instance, by their first birthday, infants are clearly integrating their existing physical knowledge[45] or social knowledge[46,47] into their statistical inferences. Similarly, children's causal inferences are influenced by their mechanistic understanding of the domain[48,49] and their familiarity with that domain.[16]

Theories Based on Learning Causal Structure

Children's ability to integrate their existing substantive knowledge with formal mechanisms for causal learning have led many researchers to propose that children are learning an abstract 'causal model'[31] or 'causal

© 2014 John Wiley & Sons, Ltd.

map'.[50b] Causal graphical models (CGMs)[31,44,50,51] can define causal relations over a variety of domains, such as physics, biology, and psychology, but they should not be taken as a domain-general representation of knowledge. An individual causal model is domain-specific (i.e., an individual model could represent a particular piece of substantive knowledge). The framework itself is more general (i.e., it can represent knowledge across domains). These models have been posited as a computational description of children's naiuml;ve theories.[50,52c] Because of the potential importance of these models, below we briefly introduce and summarize research using this framework as a description of causal inference and learning.

To begin, a graphical model is a representation of a joint probability distribution—a list of all possible combinations of events under consideration and the probability that each combination occurs. Conditional probability information can be extracted from this list. In this formalism, events or objects are represented as nodes, and vertices represent particular types of dependencies between such objects or events. Interpreting these models as representations of causal knowledge involves making three assumptions about the underlying structure of the connections between nodes (events or properties/features of objects) and vertices (dependency relations that indicate causal structure): *Mechanisms*, the *Markov Assumption*, and *Faithfulness*.

Assumption 1: Mechanisms

The first assumption is that any vertex represents a causal relation between the two nodes, specifically in the form of a mechanism that can be either observed or unobserved. That is, given a particular relation $X \rightarrow Y$, the arrow indicates that there is some mechanism whereby changing the probability of X directly affects the probability that Y will occur. A causal graph is consistent with an infinite set of probabilistic models that specify how the variables are related. A single representation of that causal structure is made by *parameterizing* the graph: defining the probability distribution for each variable conditioned on its parents. Critically, a graph's parameterization reflects assumptions about the nature of the mechanism by which causes produce effects. For instance, consider the hypothetical model of the weather shown in Figure 1. When we draw an arrow from 'rains yesterday' to 'rains today', we are positing that there is some (hidden) mechanism that relates the probability of raining yesterday to the probability that it rains today. This mechanism might be remarkably simple or complex, but critically, we reason as if such a mechanism is

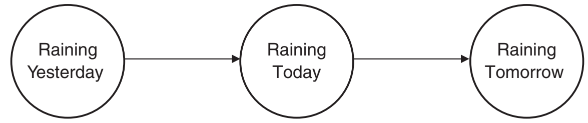

FIGURE 1 | A toy model of the weather.

present (and regardless of the complexity of the mechanism, such inferences might be available even to very young children, see Ref 53). Positing the presence of such causal mechanisms allows for a 'calculus of intervention' (see Ref 44, p. 85, see also Ref 54)—a way of interpreting how an intervention on one part of the causal system affects the system as a whole.

Intuitions about causal mechanisms, and by extension, the 'calculus of intervention' have been examined in several psychological studies. Some researchers have suggested that causal transmission was inherent in particular perceptual features of a display.[55] Given those features, one could not help to see particular sequences of events as causal. Moreover, adults reason about causal relations by virtue of the 'do' operator is described in Ref 44 (see also Ref 56); causal learning from interventions is superior to just observing the same data.[57–59]

Children engage in similar inferences. There have been various replications and extensions of the Michotte paradigm, suggesting that from very early ages, infants register certain configurations of perceptual features as causal.[3,60] Shultz argued that children understood causal relations in terms of 'generative transmission'[10] and demonstrated that preschoolers treated mechanism information as more important for judging whether a causal relation was present than correlational information (see also Ref 61). Regarding interventions, children reason differently about intervening on a causal system than when simply observing that system.[50] Four-year-olds also infer the presence of hidden causes when shown stochastic data, suggesting that they interpret probabilistic events as indicating the presence of hidden mechanisms.[62] Although it is not clear that children (or adults for that matter, see e.g., Ref 63) can articulate the mechanism underlying any particular causal relation, it does seem clear that young children can reason as if such mechanisms are present.[49,64]

Assumption 2: The Markov Assumption

The Markov assumption translates conditional probability information into causal knowledge. It states that the value of an event (i.e., a node in the graph) is independent of all other events except its children (i.e., its direct effects) conditional on its parents (i.e., its direct causes). Let's return to the model of the weather shown in Figure 1. According to this model, whether it rained

© 2014 John Wiley & Sons, Ltd.

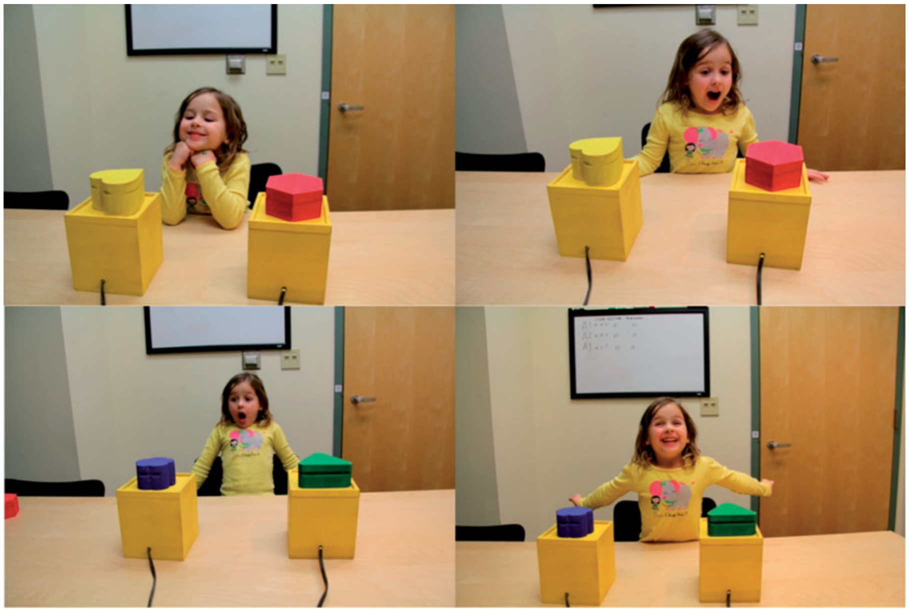

FIGURE 2 | Various reactions from a child interacting with two 'Blicket Detectors' (originally used by Gopnik and Sobel[65]), from the protocol described in Legare.[66]

yesterday and whether it rains today are statistically related, as is whether it rains today and whether it will rain tomorrow. Given those dependencies, it is also true that raining yesterday and raining tomorrow are dependent. The Markov Assumption states that raining yesterday and tomorrow are independent given the knowledge of whether it rains today. The only influence raining yesterday has on raining tomorrow is through whether it rains today.

How could we examine whether children are reasoning about the relations among events using the Markov assumption? One difficulty in answering this question is that we need a method that allows us to test whether children recognize the conditional independence relations among events separately from their prior knowledge about how these events are related. One such experimental paradigm was developed by Gopnik and Sobel,[65] who introduced children to a 'blicket detector' (see Figure 2), a machine that lights up and plays music when certain objects are placed upon it. The detector presents a novel, nonobvious property of each object: its activation potential. (The machine is controlled through an 'enabling' switch. When the switch is on, any object will activate the detector. When it is off, no object will activate the detector). Because the machine is novel, children have few expectations about what kinds of objects activate it.

Using this paradigm, researchers have found that children treat objects that activate the detector by themselves differently from objects that only activate the detector dependent on the presence of another object—that is they examined whether children obey the Markov assumption. Three- and 4-year-olds were trained to know that objects that activated the detector were called 'blickets'. Then, children observed a set of trials in which objects either independently activated the machine, or did so only dependent on the presence of another object. Specifically, on the *one cause* trials, children were shown two objects. One object (*A*) activated the detector by itself. The other object (*B*) did not. Children then saw objects *A* and *B* activate the detector together (twice). Children labeled only object *A* as a blicket. Even though object *B* activated the detector 2 out of the 3 times it was placed on it, it only did so dependent on the presence of object *A*. If children reasoned according to the Markov assumption, they would not use the positive association between object *B* and the machine's activation to infer the efficacy of object *B*, but rather recognize that such efficacy is conditionally dependent on the presence of object *A*. Remove object *A* from the equation, and object *B* lacks efficacy. Children reasoned in this manner and stated that object *B* was not a blicket. In contrast, in an analogous *two cause* condition, in which an object activated the detector independently two out of three times, children were likely to label the *B* object a blicket. Here, object *B* activates the machine independent from object *A*, and thus children should use the associations they observe.[11]

Various investigations have extended these findings to younger children,[37] other domains of knowledge,[67] and other kinds of screening-off inferences (i.e., those involving chains or common causes, instead of the common effect structure presented above, see Ref 68). These data all suggest that children robustly adhere to the Markov assumption in their causal inferences. That said, there are a variety of findings in adult cognition suggesting *nonindependence*—that adults' probabilistic inferences are inconsistent with the Markov assumption.[69–72] As an example, suppose you know that smoking causes thick blood vessels and smoking increases risks for cataracts. The Markov assumption states that the probability that a heavy smoker has a risk for cataracts is the same regardless of whether he/she has thick or normal blood vessels. However, if adult participants are given examples like this, the nonindependence effect is their likelihood to judge that the probability a heavy smoker has risk for cataracts is higher if that smoker also has thick blood vessels than normal blood vessels. Because of these results, some of the researchers cited above have suggested that this framework does not describe adults' causal inference well (or at the least provides an incomplete account of causal inference, see also Refs 73 and 74).

One potential explanation for this inconsistency between the adult and developmental findings is that the methods used to test adults' causal reasoning often contextualize the problem in such a way that adults' prior knowledge (i.e., their substantive principles) might influence the causal model they construct. Rehder and Burnet[71] were sensitive to this issue, and argued that adults do reason according to the Markov assumption, but the causal model that they build when they are asked to make such inferences incorporates prior knowledge in the form of mechanism information. That is, in the smoking example above, adult participants assume a mechanism through which smoking causes both increased risk for cataracts and thick blood vessels. In this way, nonindependence is nothing more than adults reasoning according to the Markov assumption, just not representing the simplest possible causal model.

Some have expanded on this hypothesis, suggesting that in order to describe children's and adults' causal reasoning, one underlying principle of the CGM framework should be modified, that is *minimality*. Under minimality, whenever observed data posit a causal relation, the standard instantiation of that relation is a single vertex linking cause to effect.[44] However, one could also specify a distribution of intermediary causal structures between cause and effect, with each structure's prior probability dependent on its complexity. This modification (called *Edge Replacement*) nicely explains the phenomenon of nonindependence, as it posits various kinds of mechanism knowledge that potentially influences causal inferences.[53] In addition, it explains several other novel aspects of children's causal inference.[64,75]

Assumption 3: Faithfulness

Faithfulness specifies that the data a learner observes is actually indicative of the causal structure in the world. Put simply, the faithfulness assumption is that the data children observe indicate the actual causal structure of the world. To illustrate, suppose that children observe two events X and Y. The actual structure of the world is that there is a generative relation between X and Y, such that raising the probability of X would therefore raise the probability that Y occurs. However, it is possible that there is another event (call it Z), which children do not observe, but that also affects Y. Event Z has a preventative relation with Y, such that increasing the probability of Z decreases the probability of Y. Suppose further that the extent to which X affects Y and Z affects Y are exactly the same and that they always affect Y in tandem. Thus, even though raising the probability of X should raise the probability of Y, X and Y would be seen as independent from one another. The faithfulness assumption is that this sequence of coincidental events never occurs. The causal relations among X, Y, and Z will never work out such that X and Z cancel each other's effects on Y exactly.

To our knowledge, there are no direct psychological investigations dedicated primarily to faithfulness. That said, because it essentially involves positing the presence of Cartesian demons, we do feel, however, that it is safe to assume this principle. More psychological investigation is warranted, however, to be certain of this assumption.

Integrating the Two Accounts

At this point in our review, we wish to speculate on a way of integrating accounts of causal learning based on recognizing associations among events and those based on building CGMs. The statistical learning literature suggests that infants have associative learning capacities, potentially even from birth.[35,36] This is supported by classic work demonstrating that very young infants can learn and remember associations they observe among events (see Ref 76, for a review). Thus, one possibility is that such a mechanism can account for both the acquisition of substantive principles of causal knowledge and the way in which children come to recognize any kind of causal relation.

Yet for the same reason that we lecture our introductory statistics students that 'correlation doesn't

© 2014 John Wiley & Sons, Ltd.

equal causality', we believe that such accounts cannot solely describe the formal principles through which children learn causal structure. As mentioned in the section on the Markov Assumption above, any case in which the dependence relation between two events (X and Y) switches given the presence or absence of a third (Z) suggests a causal structure in which a direct causal relation does not exist between them (e.g., $X \rightarrow Z \rightarrow Y$). Children must have a mechanism for recognizing statistical regularities among events, but also parsing out conditional independence and dependence relations, and making causal inferences based on observed data. Thus, another possibility is that children's causal reasoning is explained by a formal mechanism that is described by the CGM framework. That is, from birth, the way in which children learn causal knowledge and make causal inferences is guided by principles from the CGM framework.

One way of integrating these two approaches comes from refining Piaget's descriptions of the development of causal reasoning.[77,78] In the earliest stages of the sensorimotor stage, Piaget described the infant as only experiencing causality as a form of association of experiences: 'there is no causality for the child other than his own actions; the initial universe is not a web of causal sequences but a mere collection of events arising in extension of the activity itself' (see Ref 78, p. 220). As the infant learns to act on the world, he/she might move beyond such associations to recognize deeper relations among events.

Pearl's description of CGMs supporting inferences about *intervention* (see Ref 44, see also Ref 54) nicely fit with this description. When the child realizes that objects themselves can have efficacy on the environment (roughly consistent with Piaget's substage 4 of the sensorimotor stage, or around 8 months old), they might also begin to extract the conditional probability information inherent in associations among events.

There are two lines of support for this hypothesis. The first comes from several investigations generated by Sommerville and colleagues, suggesting that infants' emerging actions predict their causal inferences about others' intentional actions.[79,80] Sommerville et al.,[81] in particular demonstrated that providing infants with the ability to act on the environment changed whether they perceived others' actions as goal-directed.

Second, there are now numerous findings in the infant literature suggesting that infants are capable of sophisticated causal inferences.[82–85] What is interesting about these findings is that they all come from infants in the second half of the first year of life (usually 8-month-olds or older), and rarely are developmental differences investigated.[d] Cohen and colleagues[60,87,88] have suggested that the perception of simple causal relations develops between approximately ages 5–10 months. Sobel and Kirkham[37] found that children's ability to recognize conditional independence and dependence in statistical regularity developed between the ages of 5 and 8 months. Similar findings in infants' statistical learning capacities using a different (but one could argue simpler) paradigm, show analogous development between 4.5 and 6 months.[89] This developmental trajectory is consistent with the constructivist interpretation outlined above.

Two Caveats About CGMs as a Description of Children's Causal Knowledge

Examining how this computational framework describes children's causal inference engages a particular debate within the psychological and computational literatures concerning the level of representational breadth. Thinking about CGMs as a description of children's causal knowledge contrasts with associative accounts of causal reasoning[20] or similar neural network descriptions,[24] which suggest that such a computational description of causal knowledge must be more domain-general. Because causal knowledge can cross or link domains, however, this approach should also be treated differently from modularity or certain starting-state nativist accounts that propose that there are separate domain-specific causal structures, which potentially have neural correlates.[86,90,91]

Thus, one important caveat to the proceeding discussion is that CGMs are a way of representing causal structure, not a specific commitment to how causal knowledge is learned. Most psychological research on learning such causal structures have relied on algorithms based on Bayesian inference (see e.g., Refs 49, 92 and 93). That said, other algorithms do exist to describe such structure learning.[50,51] Although Bayesian approaches have been instrumental in facilitating our understanding of causal learning, they should not be taken as a model for what algorithm or calculations children are computing when faced with causal learning problems. That is, these models provide a description of the causal learning process and have been useful in generating psychological theory. They should not be considered algorithmic or process accounts of how children are exactly make causal inferences. For instance, it is possible that neural network simulations could model many of the psychological phenomena described in this paper (see Ref 25 for example). It is also possible that algorithms that simulate Bayesian updating will provide a good

© 2014 John Wiley & Sons, Ltd.

algorithmic level description for the same inferences (see Refs 92 and 94 for example). We are agnostic as to which modeling architecture will ultimately provide the best and most descriptive algorithmic-level approach. Any computational description must successfully account for extant data, but also make novel predictions concerning its psychological implications, and perhaps the most exciting work in the field comes from this endeavor.

Second, CGMs have the potential to describe how children represent their causal knowledge at multiple levels. One could imagine a CGM model describing a particular event (e.g., Abe dials a number on his phone, which causes Bob's phone to ring). One could also imagine a CGM describing this kind of event (phones can cause each other to ring). Such a distinction between *specific theories* and *framework theories*[6] is critical to thinking about the role the CGM framework plays in describing children's causal knowledge.[93,95] It is likely that learning about specific events using an algorithm from the CGM framework is guided not only by the data that children observe, but also the knowledge they possess regarding what kinds of specific causal models can be built (or that have greater *a priori* probability).

EXPLANATION AND EXPLORATION IN CHILDREN'S CAUSAL LEARNING

Having presented a basic description of computational models of causal learning, we now turn to several programmatic lines of constructivist research examining the relation between explanation, exploration, and the development of causal learning.

Explanation

The search for explanations motivates the causal learning process.[96,97] The tendency to seek and generate explanations is so pervasive and compelling that some psychologists have posited a 'drive to explain'.[98] At very early ages, children generate appropriate domain-specific explanations [13,15,81] and use questions effectively to elicit explanations from others as a means of acquiring new knowledge.[99–101] Given that young children generate and seek out explanations, how might explanations benefit causal learning?

A growing body of research confirms that the process of generating explanations, for others or for oneself, has educational benefits.[102–106] This 'self-explanation effect' has been documented in a variety of learning contexts, ranging from the acquisition of scientific content knowledge[107,108] to conceptual transitions in early childhood.[109,110] Given the intimate relation between explanation and conceptual representation,[111–113] generalization,[114] and learning,[102,115] an understanding of explanation is foundational for causal learning (see also Ref 116).

Even though there are many documented effects of explanation,[103–105,117–119] the process by which explanations benefit learning is underspecified.[12,115] If explanation is a mechanism for learning, children should benefit from providing explanations for events that afford new learning opportunities. Events that are inconsistent with prior knowledge provide just such an opportunity. The ability to explain such events could aid in learning by focusing children on current causal knowledge and provoke causal reasoning. For instance, Legare[12,66,120] examined the triggers that motivate children to construct explanations. Legare and her colleagues found that preschoolers generated more explanations when faced with outcomes that were inconsistent with their prior knowledge. Moreover, these explanations tended to refer to unobserved causal mechanisms and internal causal properties, and not external perceptual appearances. This provides promising evidence that explanation provides children the opportunity to articulate new hypotheses for events that, at first, disconfirm their current knowledge. These data are consistent with the proposal that children's explanations play an active role in the learning process and provide an empirical basis for investigating the mechanisms by which children's explanations function in the service of discovery.[e]

Despite this evidence, merely attending to inconsistency does not always lead to belief revision and theory change.[122–128] Explaining inconsistency may serve as a critical mechanism for integrating and reconciling discordant information with existing theories and reduce engagement in theory-preserving strategies like rejection and postponement. But how might the process of explaining inconsistency generate amended beliefs?

One possibility we endorse is that explaining inconsistency triggers a process of hypothesis generation that encourages learners to formulate and entertain hypotheses they would not have spontaneously considered otherwise. Generating hypotheses in the service of explanation may influence the kinds of hypotheses formulated. Both children and adults have strong intuitions about what makes something a good explanation,[129,130] and explanation-triggered hypothesis generation may promote the production of hypotheses that make for informative explanations.

In this way, there might be some deep, but as of yet unexplored connections between explanations and how children represent their causal knowledge (based on the CGM framework described above). These two

© 2014 John Wiley & Sons, Ltd.

areas of causal learning are not seen as connected; as Wellman and Liu point out, ‘causal Bayes nets seem silent on how to characterize explanations and on what role explanations might play in causal learning’ (Ref 119, p. 261). Although it is not clear how children generate explanations from the way they represent their causal knowledge, one way of interpreting the self-explanation effect is that the child treats the act of explaining in the same manner as observing analogous data. That is, children who generate an explanation might treat that explanation as data, which might affect the existing model of causal structure they possess (i.e., strengthening it if they believe their explanation is good, or weakening it if they believe their explanation is weak or uncertain, see Ref 131, for a version of this argument).

Exploration

We have proposed that the act of explaining serves as data for causal learning or conceptual change. We view the child as an active participant in the construction of causal knowledge, and not a passive viewer of information. Casting the child in this light is a hallmark of many constructive accounts.[50,52] It suggests that children also seek out data when faced with ambiguity or uncertainty, and such exploration can inform the ways in which children learn new causal structure. That is, the weaker the representation of children’s causal knowledge, the more likely they will explore their environment (presumably to strengthen that representation).

Building on classic research on children’s play,[132,133] there is converging evidence that inconsistent or problematic events also trigger exploration.[134,135] Children’s exploratory play is affected by the quality of the evidence that they observe. When multiple candidate causes are available for the same outcome and underlying causal structure is ambiguous, children preferentially explore confounded (as opposed to unconfounded) causal relations, show more variable play behavior when presented with probabilistic (as opposed to deterministic) information,[136] and can spontaneously disambiguate confounded variables.[135] Moreover, they recognize when it is necessary to explore the environment as opposed to seek help from a more knowledgeable source.[137]

Despite the evidence that anomalous or inconsistent information motivates both explanation and exploration, the way in which the two processes may jointly facilitate or drive causal knowledge acquisition has remains underdetermined. Does the process of constructing a causal explanation for inconsistent outcomes inform and constrain children’s exploratory behavior? Do causal explanation and exploratory behavior operate in tandem as hypothesis-generating and hypothesis-testing mechanisms? Is this process associated with tangible learning outcomes?

Bonawitz and colleagues[138] examined this interaction between explanation and exploration. They first assessed children’s understanding of balance events (i.e., do children know that center of mass indicates balance). Next they presented children with a free play environment that provided evidence either for a geometric center-based or center of mass-based account of balance (by, following Ref 139), using stimuli that appeared to support one kind of balance relation, but by virtue of their mass, actually supported a counterintuitive relation). They then assessed how children learned from their explanations and exploration of these objects. They found that older children (6–7-year-olds) could revise their beliefs in light of theory-inconsistent evidence, but also that children would discount such evidence if they explain these events in terms of potential auxiliary hypotheses to their existing theory. Preschoolers, in contrast, struggled to revise their beliefs given this evidence.

Preschoolers, however, do trade-off exploration and explanation in certain ways. When children lack explanatory information, they can learn causal structures from exploration,[140] and this learning is facilitated when their exploration uncovers new knowledge as opposed to confirming information they have already observed.[66,141] Moreover, young children explore novel toys more when given ambiguous data about how the toy works[135] and explore environments in systematic ways to resolve that ambiguity.[142]

Critically, however, explanations and exploration interact when learning.[143] Bonawitz and colleagues[144] demonstrated that children who heard a particular set of instructions regarding a novel object’s function were less likely to explore the object (and discover novel functions) than children who heard incomplete explanations. They suggest that children’s exploration was affected by their understanding of the pedagogical intent of the individual who generated the explanations (i.e., someone who was more knowledgeable about the object than they were). Such effects of pedagogy—taking a teacher’s intentions into account to determine why they are presenting the information they are—are detectable in adults[145] as well as toddlers[146,147] and are potentially part of a natural human communicative process.[148]

Given the trade-off between explanations heard from others and exploration of the environment demonstrated in these experiments, there are a number of compelling reasons to examine the

© 2014 John Wiley & Sons, Ltd.

interaction between explanation and exploration for learning, particularly when faced with the problem of inconsistency. Encouraging children to explain inconsistency confronts children with the inconsistent evidence most likely to foster theory revision, guides the hypothesis-testing process, and promotes learning. These explanatory intuitions may constrain learners to focus on some aspects of what they are trying to explain over others. In particular, explanation may focus learners on *causal mechanisms*[66,106] and on *abstraction*.[115,130] Generating hypotheses in the service of explanation may influence the kinds of hypotheses formulated, as well as their impact on cognition.

That said, how might the process of explaining inconsistency with prior knowledge inform children's exploratory, hypothesis-testing behavior? Legare[66] demonstrated that children's explanations and subsequent exploratory behavior following events that are consistent with their existing knowledge differ from those following inconsistent events. When children observed inconsistent events, the kind of explanation children provided differentially predicted the kind of exploratory behavior they engaged in. The kind of explanations children provided also influence rates of spontaneous, hypothesis-testing exploratory behavior and the tendency to modify existing explanatory hypotheses in the face of disconfirming evidence. For example, children who provided explanations that referred to problems with causal function engaged in more extended and more variable exploratory behavior than children who provided different kinds of explanations (e.g., explanations referring to category membership). Encouraging children to explain inconsistency confronts children with the inconsistent evidence most likely to foster theory revision, guides the hypothesis-testing process, and promotes exploration.[143]

CAUSAL LEARNING FROM OTHERS

Much of the constructivist research we have described so far focuses on the problem of causal learning as being directed by the child. When children encounter novel data or events from which they learn, they do so by integrating that information with their existing knowledge to make novel inferences, generate novel explanations, or engage in specific actions. Such a description might be correct, but it assumes that all the data children use to learn causal structure is directly observable. This is obviously not the case, and the fact that children learn from explanations (both their own and others) provides evidence for the power of learning through testimony from others.

Furthermore, children make inferences about unobservable biological events,[149] psychological events,[119] and even supernatural events.[150] Children also appeal to and easily learn culturally constructed explanations[151] and social conventions.[152,153] All of these events are not directly observable and could not be learned just from interacting with the environment.[f] In order to learn all of this information, children must rely on information generated from others. How do children learn causal knowledge from others? Do the same processes we have described in learning from observation and interaction with the world apply to learning from others?

A significant number of studies now show that children are not simply credulous of others' information (for reviews, see Refs 154–158). Children as young as 2 years are capable of judiciously using different sources by tracking informants' history of past accuracy.[159–162] Children's rapid cultural learning potentially emanates from their ability to learn selectively from others (see e.g., Ref 163).

We propose that how children learn from others is as rational as how children learn from integrating their own knowledge with observed data.[158] Preschoolers' beliefs about whom to trust are influenced by their existing knowledge about people (e.g., adults are more knowledgeable than children[164]), kinds (e.g., speakers knowledgeable about objects' labels should also be knowledgeable about those objects' functions, [161]), and expertise (e.g., speakers with a certain specific knowledge base might not be more knowledgeable overall, just about that base[165,166]). These findings all suggest that young children can integrate what they know about the world with the data they receive about the world from others.

Indeed, consistent with this rational account, some have suggested different computational accounts of the way in which children update their beliefs—including their causal knowledge—given information generated from others.[167–169] Many of these accounts use CGMs as a representation of children's existing causal knowledge, and promote different kinds of rational learning algorithms on this representation to explain how children's knowledge changes. As with our discussion of the CGM framework, such computational accounts should not be considered process models, but rather descriptions of how children might learn from others, which in turn could inform new psychological investigations.

CONCLUSION

To answer the question of how children construct abstract representations of causal knowledge from the

© 2014 John Wiley & Sons, Ltd.

data they observe, we have appealed to computational (i.e., CGMs), constructivist, and social learning frameworks to describe the process of causal learning. Our objective is to illustrate the striking sophistication of young children's causal learning capacities, as well as demonstrating how useful computational modeling can be for making predictions about those capacities. There are many outstanding open questions, such as how to translate between a causal graphical model and a verbal explanation provided by a child, or the role of children's developing cognitive capacities, such as attention or memory, in the process of such learning. The present review, however, suggests an important conclusion, that very much emanates from constructivist theories of cognitive development: The child is an active seeker of information—regardless of what kind of knowledge they are acquiring. Regarding causal knowledge, children begin to generate 'why' questions around the time they themselves offer causal explanations.[101,119,170] But, similar behavior is also seen for children in the naming spurt—the acquisition of many labels often coincides with children generating a lexical item soliciting an object's name.[171] Such active learning allows children to recognize and fill the gaps in their knowledge and construct new representations of the causal structure of the world.

NOTES

[a] Descriptions of causal learning based on calculating causal strength from associative information typically estimate these asymptotic parameters. For example, the results of the Rescorla-Wagner equation can converge to ΔP given infinitely many randomly intermixed trials.[33]

[b] We prefer the 'causal map' designation, as it does not suggest that children's causal knowledge is represented by a graphical model explicitly. Instead, it suggests that the features of such a computational account describe that representation.

[c] Interestingly, there are several domains of knowledge like number, language, or spatial relations in which causal maps potentially have little applicability because those domains typically do not represent knowledge in terms of causal relations. For example, the development of numerical knowledge (such as knowledge of the mapping between the natural numbers and numerals) is not represented in terms of causal relations. How the CGM framework potentially relates to the representation of this knowledge is beyond the scope of this article.

[d] Two points should be made about these findings. First, they are usually interpreted as evidence for causality being part of 'core knowledge' or causal reasoning as an innately specified capacity (see e.g., Ref 86). Second, to our knowledge, all of these findings are consistent with the CGM framework described above, although they are usually not described in that way.

[e] It is also interesting to consider cases in which children learn from explaining information they already know, particularly when teaching another person. There are many demonstrations of preschoolers being able to make sophisticated inferences about teaching,[121] but only few demonstrations that children's actual teaching affects their learning. We believe that this is an important line of research for future investigation.

[f] Critically, children also learn language, which, while not causal in nature, does involve understanding a set of culturally constructed, but arbitrary mappings between phonology and meaning.[154]

ACKNOWLEDGMENTS

The authors were supported by NSF (DLS/REESE-1223777 to D.M.S.) and by an ESRC (Economic and Social Research Council) Large Grant (REF RES-060-25-0085 to C.H.L.) during the writing of this article.

REFERENCES

1. Piaget J. *The Child's Conception of the World*. London: Routledge and Kegan Paul; 1929.
2. Hespos SJ, Baillargeon R. Reasoning about containment events in very young infants. *Cognition* 2001, 78:207–245.
3. Leslie AM, Keeble S. Do six-month-old infants perceive causality? *Cognition* 1987, 25:265–288.
4. Needham A, Baillargeon R. Intuitions about support in 4.5-month-old infants. *Cognition* 1993, 47:121–148.
5. Spelke ES, Breinlinger K, Macomber J, Jacobson K. Origins of knowledge. *Psychol Rev* 1992, 99: 605–632.
6. Wellman HM. *The Child's Theory of Mind*. Cambridge, MA: MIT Press; 1990.

© 2014 John Wiley & Sons, Ltd.

7. Gelman SA, Wellman HM. Insides and essences: early understandings of the non-obvious. *Cognition* 1991, 38:213–244.
8. Gopnik A, Wellman HM. The theory theory. In: Hirschfield L, Gelman S, eds. *Mapping the Mind: Domain Specificity in Cognition and Culture*. New York: Cambridge University Press; 1994, 257–293.
9. Bullock M, Gelman R, Baillargeon R. The development of causal reasoning. In: Friedman WJ, ed. *The Developmental Psychology of Time*. New York: Academic Press; 1982, 209–254.
10. Shultz TR. Rules of causal attribution. *Monogr Soc Res Child Dev* 1982, 47:1–51. Serial No. 194.
11. Gopnik A, Sobel DM, Schulz L, Glymour C. Causal learning mechanisms in very young children: two, three, and four-year-olds infer causal relations from patterns of variation and co-variation. *Dev Psychol* 2001, 37:620–629.
12. Legare CH, Gelman SA, Wellman HM. Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Dev* 2010, 81:929–944.
13. Schult CA, Wellman HM. Explaining human movements and actions: children's understanding of the limits of psychological explanation. *Cognition* 1997, 62:291–324.
14. Harris PL, German T, Mills P. Children's use of counterfactual thinking in causal reasoning. *Cognition* 1996, 61:233–259.
15. Sobel DM. Exploring the coherence of young children's explanatory abilities: evidence from generating counterfactuals. *Br J Dev Psychol* 2004, 22:37–58.
16. Berzonsky MD. The role of familiarity in children's explanations of physical causality. *Child Dev* 1971, 42:705–715.
17. Mackintosh NJ. A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol Rev* 1975, 82:276–298.
18. Rescorla RA, Wagner AR. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black AH, Prokasy WF, eds. *Classical Conditioning II: Current Theory and Research*. New York: Appleton-Century-Crofts; 1972, 64–99.
19. Cramer RE, Weiss RF, William R, Reid S, Nieri L, Manning-Ryan B. Human agency and associative learning: Pavlovian principles govern social process in causal relationship detection. *Q J Exp Psychol* 2002, 55:241–266.
20. Dickinson A, Shanks D. Instrumental action and causal representation. In: Sperber D, Premack D, Premack AJ, eds. *Causal Cognition: A Multidisciplinary Debate*. New York: Clarendon Press/Oxford University Press; 1995, 5–25.
21. Kruschke JK, Blair NJ. Blocking and backward blocking involve learned inattention. *Psychon Bull Rev* 2000, 7:636–645.
22. Van Hamme LJ, Wasserman EA. Cue competition in causality judgments: the role of nonpresentation of compound stimulus elements. *Learn Motiv* 1994, 25:127–151.
23. Wasserman EA, Berglan LR. Backward blocking and recovery from overshadowing in human causal judgement: the role of within-compound associations. *Q J Exp Psychol* 1998, 51:121–138.
24. McClelland JL, Thompson RM. Using domain-general principles to explain children's causal reasoning abilities. *Dev Sci* 2007, 10:333–356.
25. Rogers TT, McClelland JL. *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: MIT press; 2004.
26. Allan LG. A note on measurement of contingency between two binary variables in judgment tasks. *Bull Psychon Soc* 1980, 15:147–149.
27. Jenkins HM, Ward WC. Judgment of contingency between responses and outcomes. *Psychol Monogr* 1965, 79:1.
28. Shanks DR. Is human learning rational? *Q J Exp Psychol* 1995, 48:257–279.
29. Cheng PW. From covariation to causation: a causal power theory. *Psychol Rev* 1997, 104:367–405.
30. Novick LR, Cheng PW. Assessing interactive causal influence. *Psychol Rev* 2004, 111:455–485.
31. Glymour C. *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology*. Cambridge, MA: The MIT Press; 2001.
32. Tenenbaum JB, Griffiths TL. *Structure learning in human causal induction*. In: *Proceedings of the 13th Annual Conference on the Advances in Neural Information Processing Systems*, Vancouver, Canada, 2001.
33. Danks D. Equilibria of the Rescorla-Wagner model. *J Math Psychol* 2003, 47:109–121.
34. Cheng PW, Novick LR. A probabilistic contrast model of causal induction. *J Pers Soc Psychol* 1990, 58:545–567.
35. Haith MM. Future-oriented processes in infancy: the case of visual expectations. In: Granrud C, ed. *Visual Perception and Cognition in Infancy*. Hillsdale, NJ: Erlbaum; 1993, 235–264.
36. Kirkham NZ, Slemmer JA, Johnson SP. Visual statistical learning in infancy. *Cognition* 2002, 83:B35–B42.
37. Sobel DM, Kirkham NZ. Bayes nets and blickets: infants developing representations of causal knowledge. *Dev Sci* 2007, 10:298–306.
38. Johnson S, Slaughter V, Carey S. Whose gaze will infants follow? The elicitation of gaze-following in 12-month-olds. *Dev Sci* 1998, 1:233–238.
39. Goldstein MH, Schwade JA, Bornstein MH. The value of vocalizing: five-month-old infants associate their own noncry vocalizations with responses from caregivers. *Child Dev* 2009, 80:636–644.

© 2014 John Wiley & Sons, Ltd.

40. Saffran JR, Aslin RN, Newport EL. Statistical learning by 8-month-old infants. *Science* 1996, 274:1926–1928.
41. Xu F, Garcia V. Intuitive statistics by 8-month-old infants. *Proc Natl Acad Sci USA* 2008, 105:5012–5015.
42. Kushnir T, Xu F, Wellman HM. Young children use statistical sampling to infer the preferences of others. *Psychol Sci* 2010, 21:1134–1140.
43. Ma L, Xu F. Preverbal infants infer rational agents from the perception of regularity. *Dev Psychol* 2013, 49:1330–1337.
44. Pearl J. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press; 2000.
45. Denison S, Xu F. Integrating physical constraints in statistical inference by 11-month-old infants. *Cogn Sci* 2010, 34:885–908.
46. Wu R, Gopnik A, Richardson DC, Kirkham NZ. Infants learn about objects from statistics and people. *Dev Psychol* 2011, 47:1220–1229.
47. Tummeltshammer KS, Kirkham NZ. Learning to look: probabilistic variation and noise guide infants' eye movements. *Dev Sci* 2013, 16:760–771.
48. Schulz LE, Bonawitz EB, Griffiths T. Can being scared make your tummyache? Naiuml;ve theories, ambiguous evidence and preschoolers' causal inferences. *Dev Psychol* 2007, 43:1124–1139.
49. Sobel DM, Munro SA. Domain generality and specificity in children's causal inferences about ambiguous data. *Dev Psychol* 2009, 45:511–524.
50. Gopnik A, Glymour C, Sobel DM, Schulz LE, Kushnir T, Danks D. A theory of causal learning in children: causal maps and Bayes nets. *Psychol Rev* 2004, 111:3–32.
51. Spirtes P, Glymour C, Scheines R. *Causation, Prediction, and Search (Springer Lecture Notes in Statistics*. 2ndrevised ed. Cambridge, MA: MIT Press; 2001.
52. Gopnik A, Wellman HM. Reconstructing constructivism: causal models, Bayesian learning mechanisms, and the theory theory. *Psychol Bull* 2012, 108:1085–1108.
53. Buchanan DW, Tenenbaum JB, Sobel DM. Edge replacement and nonindependence in causation. In: *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*, Portland, OR, 2010.
54. Woodward J. *Making Things Happen*. New York: Oxford; 2003.
55. Michotte A. *Causalite, Permanence Et Realite Phenomenales*. Oxford: Publications Universitaires; 1962.
56. Sloman SA, Lagnado D. Do we "do"? *Cognit Sci* 2005, 29:5–39.
57. Lagnado DA, Sloman S. The advantage of timely intervention. *J Exp Psychol Learn Mem Cogn* 2004, 30:856–876.
58. Steyvers M, Tenenbaum JB, Wagenmakers EJ, Blum B. Inferring causal networks from observations and interventions. *Cognit Sci* 2003, 27:453–489.
59. Waldmann MR, Hagmayer Y. Seeing versus doing: two modes of accessing causal knowledge. *J Exp Psychol Learn Mem Cogn* 2005, 31:216–227.
60. Oakes LM, Cohen LB. Infant perception of a causal event. *Cognit Dev* 1990, 5:193–207.
61. Madole KL, Cohen LB. The role of object parts in infants' attention to form-function correlations. *Dev Psychol* 1995, 31:637–648.
62. Schulz LE, Sommerville J. God does not play dice: causal determinism and children's inferences about unobserved causes. *Child Dev* 2006, 77:427–442.
63. Rozenblit L, Keil F. The misunderstood limits of folk science: an illusion of explanatory depth. *Cognit Sci* 2002, 26:521–562.
64. Buchanan DW, Sobel DM. Mechanism-based causal reasoning in young children. *Child Dev* 2011, 82:2053–2066.
65. Gopnik A, Sobel DM. Detecting blickets: how young children use information about causal powers in categorization and induction. *Child Dev* 2000, 71:1205–1222.
66. Legare CH. Exploring explanation: explaining inconsistent information guides hypothesis-testing behavior in young children. *Child Dev* 2012, 83:173–185.
67. Schulz LE, Gopnik A. Causal learning across domains. *Dev Psychol* 2004, 40:162–176.
68. Sobel DM, Sommerville JA. Rationales in children's causal learning from others' actions. *Cognit Dev* 2009, 24:70–79.
69. Mayrhofer R, Hagmayer Y, Waldmann MR. Agents and causes: a Bayesian error attribution model of causal reasoning. In: *Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society*, Portland, OR, 2010.
70. Park J, Sloman S. Mechanistic beliefs determine adherence to the Markov property in causal reasoning. *Cognit Psychol* 2013, 67:186–216.
71. Rehder B, Burnett R. Feature inference and the causal structure of categories. *Cognit Psychol* 2005, 50:264–314.
72. Walsh CR, Sloman SA. Revising causal beliefs. In: Forbus K, Gentner D, Regier T, eds. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum; 2004, 1423–1427.
73. Bes B, Sloman S, Lucas CG, Raufaste E. Non-Bayesian inference: causal structure trumps correlation. *Cognit Sci* 2012, 36:1178–1203.
74. Walsh C, Sloman SA. Updating beliefs with causal models: violations of screening off. In: Gluck MA, Anderson JR, Kosslyn SM, eds. *Memory and Mind: A Festschrift for Gordon H. Bower*. Mahwah, NJ: Lawrence Erlbaum Associates; 2008, 345–357.

© 2014 John Wiley & Sons, Ltd.

75. Erb C, Buchanan DW, Sobel DM. Variability indicates complexity in children's inference about causal mechanisms. *Cognition* 2013, 129:494–500.
76. Rovee-Collier C. The development of infant memory. *Curr Dir Psychol Sci* 1999, 8:80–85.
77. Piaget J. *The Origins of Intelligence in Children*. Madison, WI: International Universities Press; 1952.
78. Piaget J. *The Construction of Reality in the Child*. New York: Basic Books; 1955.
79. Sommerville JA, Woodward AL. Pulling out the intentional structure of action: the relation between action processing and action production in infancy. *Cognition* 2005, 95:1–30.
80. Sommerville JA, Woodward AL, Needham A. Action experience alters 3-month-old infants' perception of others' actions. *Cognition* 2005, 96:B1–B11.
81. Shultz T, Wells D, Sarda M. Development of the ability to distinguish intended actions from mistakes, reflexes, and passive movements. *Br J Soc Clin Psychol* 1980, 19:301–310.
82. Kosugi D, Ishida H, Fujita K. 10-Month-old infants' inference of invisible agent: distinction in causality between object motion and human action. *Jpn Psychol Res* 2003, 45:15–24.
83. Muentener P, Carey S. Infants' causal representations of state change events. *Cognit Psychol* 2010, 61:63–86.
84. Saxe R, Tenenbaum JB, Carey S. Secret agents: inferences about hidden causes by 10- and 12-month-old infants. *Psychol Sci* 2005, 16:995–1001.
85. Teglas E, Vul E, Girotto V, Gonzalez M, Tenenbaum JB, Bonatti LL. Pure reasoning in 12-month-old infants as probabilistic inference. *Science* 2011, 332:1054–1059.
86. Carey S. *The Origin of Concepts*. Oxford: New York; 2009.
87. Cohen LB, Amsel G. Precursors to infants' perception of the causality of a simple event. *Infant Behav Dev* 1998, 21:713–732.
88. Cohen LB, Oakes LM. How infants perceive simple causality. *Dev Psychol* 1993, 29:421–433.
89. Denison S, Reed C, Xu F. The emergence of probabilistic reasoning in very young infants: evidence from 4.5- and 6-month-olds. *Dev Psychol* 2013, 49:243–249.
90. Atran S. *Cognitive Foundations of Natural History: Towards an Anthropology of Science*. New York: Cambridge University Press; 1990.
91. Leslie AM, Roth D. What autism teaches us about metarepresentation. In: Baron-Cohen S, Tager-Flusberg H, Cohen D, eds. *Understanding Other Minds: Perspectives from Autism*. Oxford: Oxford University Press; 1993, 83–111.
92. Bonawitz EB, Denison S, Chen A, Gopnik A, Griffiths TL. A simple sequential algorithm for approximating Bayesian inference. In: *Proceedings of the 33*rd *Annual Cognitive Science Society*, Boston, MA, 2011.
93. Kemp C, Goodman ND, Tenenbaum JB. Learning to learn causal models. *Cognit Sci* 2010, 34:1185–1243.
94. Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. Theory acquisition as stochastic search. In: *Proceedings of the 32*nd *Annual Cognitive Science Society*, Boston, MA, 2011.
95. Griffiths TL, Tenenbaum JB. Two proposals for causal grammars. In: Gopnik A, Schulz L, eds. *Causal Learning: Psychology, Philosophy, and Computation*. Oxford: Oxford University Press; 2007.
96. Brewer WF, Chinn CA, Samarapungavan A. Explanation in scientists and children. *Minds Mach* 1998, 8:119–136.
97. Keil F, Wilson R. *Explanation and Cognition*. Cambridge, MA: MIT Press; 2000.
98. Gopnik A. Explanation as orgasm. *Minds Mach* 1998, 8:101–118.
99. Chouinard MM. Children's questions: a mechanism for cognitive development. *Monogr Soc Res Child Dev* 2007, 72:vii.
100. Legare CH, Mills CM, Souza AL, Plummer LE, Yasskin R. The use of questions as problem-solving strategies in early childhood. *J Exp Child Psychol* 2012, 114:63–76.
101. Hood L, Bloom L. What, when, and how about why: a longitudinal study of early expressions of causality. *Monogr Soc Res Child Dev* 1979, 44:1–47. Serial No. 181.
102. Chi MTH. Self-explaining expository texts: the dual processes of generating inferences and repairing mental models. In: Glaser R, ed. *Advances in Instructional Psychology*. Hillsdale, NJ: Erlbaum; 2000, 161–238.
103. Crowley K, Siegler RS. Explanation and generalization in young children's strategy learning. *Child Dev* 1999, 70:304–316.
104. de Leeuw N, Chi MTH. The role of self-explanation in conceptual change learning. In: Sinatra G, Pintrich P, eds. *Intentional Conceptual Change*. Hillsdale, NJ: Erlbaum; 2003, 55–78.
105. Lombrozo T. The structure and function of explanations. *Trends Cognit Sci* 2006, 10:464–470.
106. Legare CH, Lombrozo T. Selective benefits of explanation on learning in early childhood. *J Exp Child Psychol*. In press.
107. Chi MTH, Bassok M, Lewis M, Reimann P, Glaser R. Self-explanations: how students study and use examples in learning to solve problems. *Cognit Sci* 1989, 13:145–182.
108. Chi MT, De Leeuw N, Chiu MH, LaVancher C. Eliciting self-explanations improves understanding. *Cognit Sci* 1994, 18:439–477.

© 2014 John Wiley & Sons, Ltd.

109. Amsterlaw J, Wellman HM. Theories of mind in transition: a microgenetic study of the development of false belief understanding. *J Cogn Dev* 2006, 7:139–172.
110. Siegler RS. How does change occur: a microgenetic study of number conservation. *Cogn Psychol* 1995, 25:225–273.
111. Carey S. *Conceptual Change in Childhood*. Cambridge, MA: MIT Press/Bradford Books; 1985.
112. Lombrozo T. Explanation and categorization: how "why?" informs "what?". *Cognition* 2009, 110: 248–253.
113. Murphy GL, Medin DL. The role of theories in conceptual coherence. *Psychol Rev* 1985, 92:289–316.
114. Sloman SA. When explanations compete: the role of explanatory coherence on judgments of likelihood. *Cognition* 1994, 52:1–21.
115. Williams JJ, Lombrozo T. The role of explanation in discovery and generalization: evidence from category learning. *Cognit Sci* 2010, 34:776–806.
116. Koslowski B. *Theory and evidence: the development of scientific reasoning*. Cambridge, MA: MIT Press; 1996.
117. Rittle-Johnson B, Saylor M, Swygert K. Learning from explaining: does it matter if mom is listening? *J Exp Child Psychol* 2007, 100:215–224.
118. Roy M, Chi MTH. Self-explanation in a multi-media context. In: Mayer R, ed. *Cambridge Handbook of Multimedia Learning*. Cambridge, MA: Cambridge University Press; 2005, 271–286.
119. Wellman HM, Liu D. Causal reasoning as informed by the early development of explanations. In: Gopnik A, Schulz LE, eds. *Causal Learning: Psychology, Philosophy, and Computation*. New York: Oxford University Press; 2007, 261–279.
120. Legare CH, Gelman SA. Examining explanatory biases in young children's biological reasoning. *J Cogn Dev* 2014. doi:10.1080/15248372.2012.749480.
121. Ziv M, Frye D. Children's understanding of teaching: the role of knowledge and belief. *Cogn Dev* 2004, 19:457–477.
122. Bindra D, Clarke KA, Shultz TR. Understanding predictive relations of necessity and sufficiency in formally equivalent "causal" and "logical" problems. *J Exp Psychol Gen* 1980, 109:422–443.
123. Chinn CA, Brewer WF. The role of anomalous data in knowledge acquisition: a theoretical framework and implications for science instruction. *Rev Educ Res* 1993, 63:1–49.
124. Klahr D, Dunbar K. Dual space search during scientific reasoning. *Cognit Sci* 1988, 12:1–48.
125. Fay AL, Klahr D. Knowing about guessing and guessing about knowing: preschoolers' understanding of indeterminacy. *Child Dev* 1996, 67:689–716.
126. Kuhn D. Children and adults as intuitive scientists. *Psychol Rev* 1989, 96:674–689.
127. Vosniadou S, Brewer WF. Mental models of the earth: a study of conceptual change in childhood. *Cogn Psychol* 1992, 24:535–585.
128. Vosniadou S, Brewer WF. Mental models of the day/night cycle. *Cognit Sci* 1994, 18:123–184.
129. Frazier BN, Gelman SA, Wellman HM. Preschoolers' search for explanatory information within adult-child conversation. *Child Dev* 2009, 80:1592–1611.
130. Lombrozo T. Simplicity and probability in causal explanation. *Cogn Psychol* 2007, 55:232–257.
131. Williams JJ, Lombrozo T, Redher B. The hazards of explanation: overgeneralization in the face of exceptions. *J Exp Psychol Gen* 2013, 142:1006–1114.
132. Hutt C, Bhavnani R. Predictions from play. *Nature* 1972, 237:171–172.
133. Switzky HN, Haywood HC, Isett R. Exploration, curiosity, and play behavior in young children: the effects of stimulus complexity. *Dev Psychol* 1974, 10:321–329.
134. Baldwin DA, Markman EM, Melartin RL. Infants' ability to draw inferences about nonobvious object properties: evidence from exploratory play. *Child Dev* 1993, 64:711–728.
135. Schulz LE, Bonawitz E. Serious fun: preschoolers play more when evidence is confounded. *Dev Psychol* 2007, 43:1045–1050.
136. Schulz LE, Hooppell K, Jenkins A. Judicious imitation: young children imitate deterministic actions exactly, stochastic actions more variably. *Child Dev* 2008, 79:395–410.
137. Gweon H, Schulz LE. 16-Month-olds rationally infer causes of failed actions. *Science* 2011, 332:1524.
138. Bonawitz EB, van Schijndel T, Friel D, Schulz LE. Balancing theories and evidence in children's exploration, explanations, and learning. *Cognit Psychol* 2012, 64:215–234.
139. Karmiloff-Smith A, Inhelder B. If you want to get ahead, get a theory. *Cognition* 1974, 3:195–212.
140. Schulz LE, Gopnik A, Glymour C. Preschool children learn about causal structure from conditional interventions. *Dev Sci* 2007, 10:322–332.
141. Sobel DM, Sommerville JA. The importance of discovery in children's causal learning from interventions. *Front Dev Psychol* 2010, 1:176–183.
142. Cook C, Goodman ND, Schulz LE. Where science starts: spontaneous experiments in preschoolers' exploratory play. *Cognition* 2011, 120:341–349.
143. Legare CH. The contributions of explanation and exploration to children's scientific reasoning. *Child Dev Perspect*. In press.
144. Bonawitz EB, Shafto P, Gweon H, Goodman N, Spelke E, Schulz LE. The double edged sword of pedagogy: teaching limits children's spontaneous exploration and discovery. *Cognition* 2011, 120:322–330.

© 2014 John Wiley & Sons, Ltd.

145. Shafto P, Goodman ND. Teaching games: statistical sampling assumptions for learning in pedagogical situations. In: *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, Washington, DC, 2008.
146. Gweon H, Tenenbaum JB, Schulz LE. Infants consider both the sample and the sampling process in inductive generalization. *Proc Natl Acad Sci USA* 2010, 107:9066.
147. Yang DJ, Bushnell EW, Buchanan DW, Sobel DM. Infants' use of contextual cues in imitation and generalization of effective causal actions. *J Exp Child Psychol* 2013, 116:510–531.
148. Csibra G, Gergely G. Natural pedagogy. *Trends Cogn Sci* 2009, 13:148–153.
149. Kalish C. Causes and symptoms in preschoolers' conceptions of illness. *Child Dev* 1996, 67:1647–1670.
150. Woolley JD. Thinking about fantasy: are children fundamentally different thinkers and believers from adults? *Child Dev* 1997, 68:991–1011.
151. Legare CH, Evans EM, Rosengren KS, Harris PL. The coexistence of natural and supernatural explanations across cultures and development. *Child Dev* 2012, 83:779–793.
152. Herrmann PA, Legare CH, Harris PL, Whitehouse H. Stick to the script: The effect of witnessing multiple actors on children's imitation. *Cognition* 2013, 129:536–543.
153. Rakoczy H, Warneken F, Tomasello M. The sources of normativity: young children's awareness of the normative structure of games. *Dev Psychol* 2008, 44:875–881.
154. Bloom P. *How Children Learn the Meaning of Words*. Cambridge, MA: MIT Press; 2000.
155. Harris PL. *Trusting What You're Told: How Children Learn from Others*. Cambridge, MA: Harvard University Press; 2012.
156. Heyman GD, Legare CH. Social cognitive development: learning from others. In: Carlston DE, ed. *The Oxford Handbook of Social Cognition*. New York: Oxford University Press; 2013, 749–766.
157. Mills CM. Knowing when to doubt: developing a critical stance when learning from others. *Dev Psychol* 2013, 49:404–418.
158. Sobel DM, Kushnir T. Knowledge matters: how children evaluate the reliability of testimony as a process of rational inference. *Psychol Rev* 2013, 120:779–797.
159. Clement F, Koenig M, Harris P. The ontogenesis of trust. *Mind Lang* 2004, 19:360–379.
160. Koenig MA, Clement F, Harris PL. Trust in testimony: children's use of true and false statements. *Psychol Sci* 2004, 15:694–698.
161. Koenig MA, Harris PL. Preschoolers mistrust ignorant and inaccurate speakers. *Child Dev* 2005, 76:1261–1277.
162. Koenig MA, Woodward AL. Sensitivity of 24-month-olds to the prior inaccuracy of the source: possible mechanisms. *Dev Psychol* 2010, 46:815–826.
163. Harris PL, Koenig MA. Trust in testimony: how children learn about science and religion. *Child Dev* 2006, 77:505–524.
164. Jaswal VK, Neely LA. Adults don't always know best preschoolers use past reliability over age when learning new words. *Psychol Sci* 2006, 17:757–758.
165. Koenig MA, Jaswal VK. Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Dev* 2011, 82:1634–1647.
166. Sobel DM, Corriveau KH. Children monitor individuals' expertise for word learning. *Child Dev* 2010, 81:669–679.
167. Butterfield J, Jenkins OC, Sobel DM, Schwertfeger J. Modeling aspects of theory of mind with Markov random fields. *Int J Soc Robot* 2009, 1:41–51.
168. Shafto P, Eaves B, Navarro DJ, Perfors A. Epistemic trust: modeling children's reasoning about others' knowledge and intent. *Dev Sci* 2012, 15: 436–447.
169. Shafto P, Goodman ND, Frank MC. Learning from others: the consequences of psychological reasoning for human learning. *Perspect Psychol Sci* 2012, 7:341–351.
170. Hickling A, Wellman H. The emergence of children's causal explanations and theories: evidence from everyday conversation. *Dev Psychol* 2001, 37: 668–683.
171. Nelson K. Structure and strategy in learning to talk. *Monogr Soc Res Child Dev* 1973, 38:136.

© 2014 John Wiley & Sons, Ltd.