# 3 Beyond Rational Constructivism

We argued in chapter 2 that young children can integrate their existing knowledge with new data in order to learn about the world, developing and changing representations of a set of naïve theories in ways analogous to the process of science. We also argued that this process of learning and theory formation can be described by the causal graphical model framework, because this process is essentially one of making a series of causal inferences. Finally, we suggested that Bayesian inference serves as a good computational-level description of how children (and potentially adults) engage in causal reasoning.

These theoretical constructs were introduced to solve a problem with the way that theory theory was initially formulated. Theory theory described cognitive development by stating that children are born with impoverished theories, which they refine through observation and interaction with the world. This description, though generally accurate (we believe), was vague. And while we left a great many of the computational details out of chapter 2 (so maybe *we're* vague), the computational framework of Bayesian inference over causal graphical models is not a vague description of how learning works (see Gopnik & Schulz, 2007; Gopnik & Tenenbaum, 2007; Tenenbaum et al., 2006; Woodward, 2003, for many more computationally oriented examples).

However, the rational constructivist framework is not without its challenges. In this chapter, we identify several of them and attempt to respond to them. It is important to note that our goal in this chapter is not to present a full description and defense of this framework, or a full accounting of its faults and flaws. The theoretical dissatisfactions we want to highlight here are ones primarily related to the main focus of this book: how children's causal reasoning underpins their scientific thinking. For that reason, we will not exhaustively cover every objection to the causal graphical modeling framework or to rational constructivism more generally. Rather, we outline a few key problems as a way of beginning to build toward a reconciliation of this work (which has primarily been influential within the field of cognitive development) with work on children's scientific thinking. These analyses serve as background for the empirical studies presented in part II of this book, which provides several case studies for how this reconciliation could proceed.

## Nonindependence

One objection to the causal graphical model framework as a description of children's causal inference is that most of the studies in this area (including Gopnik et al., 2001, and many of the other studies that we mentioned in chapter 2) only tested children's inferences about one kind of causal structure: a *common effect* structure in which two objects could independently be potential causes of the machine's activation. But this is not the only way to test the Markov assumption. One could also use a *causal chain* structure (as illustrated by the example of the backwards doG of Thunder in chapter 2) or a *common cause* structure. Interestingly, when presented with some of these structures, adults do not always reason according to the Markov assumption (e.g., Rehder & Burnett, 2005; Walsh & Sloman, 2004). But if adults do not reason according to this assumption in simple paradigms where they can be asked fine-grained questions, it seems even less likely that children do reason according to this assumption. In that case, why should we believe that the causal graphical model framework serves as a good starting point for describing children's causal reasoning?

At the outset of our response to this challenge, we want to say that some investigations with preschoolers have shown that they reason according to the Markov assumption for all three of these different kinds of causal structures, not just common effect structures (e.g., Schulz et al., 2007; Sobel & Sommerville, 2009).[1] But the broader challenge from the adult literature has important implications for our understanding of children's causal reasoning, so it is worth discussing in detail.

Consider, for example, a common cause model, such as the one shown in figure 3.1. In this model, event A causes two other events (B and C).

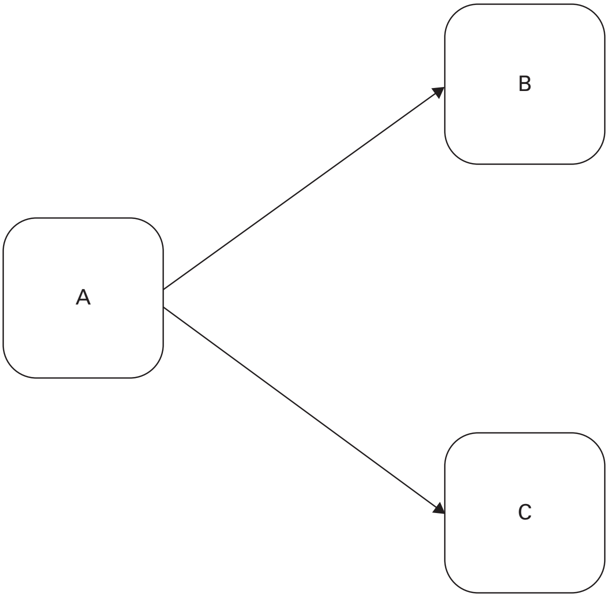

**Figure 3.1**
A causal graphical model with a common cause structure.

There is a question about how people make predictions when told about this model. Specifically, if we observe event A occur, but then event B fails to occur, should this change our judgment about whether event C occurred? Is our judgment in this case different from the case where we observe only event A occurring with no knowledge of event B? Under the Markov assumption, causes produce their effects independently of each other. More precisely, the Markov assumption says that the probability that C happens given that A happened is the same as the probability that C happens given that both A and B happened. This means that the answer to the questions above should be that these cases are the same; we should make judgements about C's occurrence based only on A, not on B.

However, that prediction does not match how many adults think about these two situations. Several experiments indicate that adults show these robust *nonindependence* effects: Adults change their judgments about whether C should occur based on what they know about B (e.g., Rehder & Burnett, 2005). Nonindependence has been interpreted either as a fallacy of the modeling framework or as evidence that the modeling framework does not provide a good description of causal reasoning.

Leaving aside the possibility that adults and children might engage in causal reasoning differently (see e.g., Liquin & Gopnik, 2022; Lucas et al., 2014, for good evidence for this claim), there is another explanation for nonindependence effects. When researchers tell adult participants about a causal model, either by describing it verbally or by instantiating it in the real world, these participants might not be reasoning about the specific causal model that the researchers intend. Because these models are descriptions of the way children and adults might represent causal knowledge, they are subject to other domain-general processes that could influence those representations. One such process, which we discuss in detail in chapter 10, is imagination. Simply put, when told about or shown causal evidence, reasoners might represent not only those events, but many other possible ways in which those events could be related. This relates to the issue of mechanisms that we discussed in chapter 2: Because any given model could be compatible with many different mechanisms, more information is required to narrow down the set of possibilities, even though both adults and children understand that there must be some mechanism or other by which these events are related.

So let's consider again the common cause model introduced above (figure 3.1). We can instantiate this model with an example (adapted from Buchanan & Sobel, 2014): Alice is using her cell phone to try to call Bob and Charlie. Alice dials Charlie. What's the likelihood that she connects with Charlie? Cellular technology being what it is, the answer is probably not 100%. But it is also not 0%, and (hopefully) it is closer to the former than the latter (otherwise Alice should really switch her service provider). Now what if you know that Alice tried dialing Bob first but could not connect. Do you revise your estimate about Alice's chances of successfully calling Charlie?

You probably said yes. If Alice can't get through to Bob, then there is probably something wrong, and whatever is wrong might affect Alice's ability to connect with Charlie, or possibly anyone. Maybe she just misdialed the numbers. But maybe the cellular signal is weak where Alice is, or her battery is dead, or someone is jamming her signal. It is also possible that something is wrong with Bob's phone—he could be in a dead zone, or his battery died, or he's the target of jammers, or his phone is turned off.

Some of these lines of thinking could lead you (and participants in the experiments cited above) to violate the Markov assumption and to consider events that should be independent as not. That is, maybe people don't reason according to this principle after all.

We can respond to this objection by noticing that, once you start imagining these *somethings* that could be affecting Alice's or Bob's phones, you are no longer reasoning about the simple common cause model depicted in figure 3.1. Instead, you're reasoning about a different model, one with a placeholder for the *somethings* that you assumed were happening with Alice's and Bob's phones. For example, some of the situations that we described would be better represented with a model like that in figure 3.2.

If this is the model that you are using to think about the situation, then you are not necessarily violating the Markov assumption. You have just constructed a different model. In a model that specifies an extra "something" (i.e., another node or set of nodes) between event A and events B and C, the probability of C given A and B is not the same as the probability of C given just A. This point, made by a number of researchers (e.g., Rehder & Burnett, 2005), provides a good way of thinking about how to solve the problem

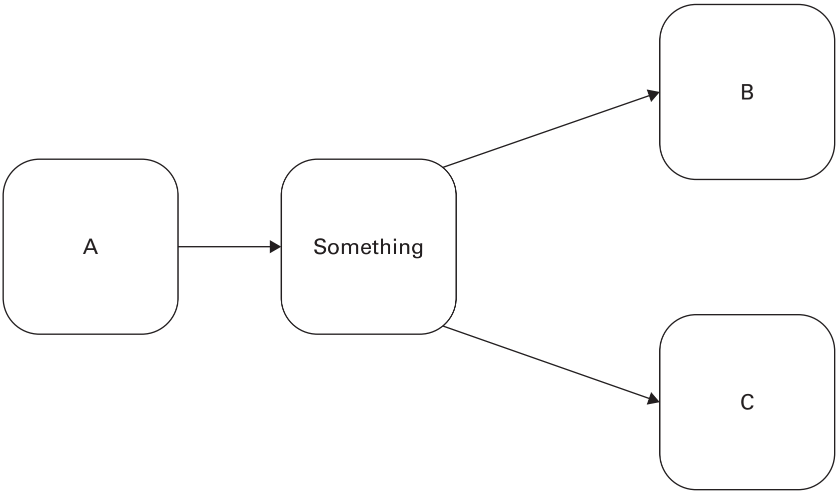

**Figure 3.2**
Constructing a common cause model with an additional node can explain why people sometimes appear to violate the Markov assumption.

of nonindependence. On the other side of the same coin, Park and Sloman (2013) showed that adult reasoners show nonindependence effects only in certain circumstances. If adults have a good representation of a potential hidden mechanism that might relate observed variables together, they are much more likely to reason according to the Markov assumption.

To address these issues, we extended the causal graphical modeling framework to account for the possibility that children and adults might imagine possible mechanisms for the relations they observe (Buchanan & Sobel, 2014). We called this extension *edge replacement*. In the edge replacement framework, people reason about a set of causal graphical models that posit different representations of mechanisms between each causal relation. That is, for each graph, reasoners simulate a set of possible mechanisms and try to match those mechanisms to the data. In this view, nonindependence is a feature, not a bug. It suggests that people are not reasoning about a simple model that contains only Alice, Bob, and Charlie (or at least, only their phones). Rather, people are reasoning about the mechanisms through which these causal relations might function (i.e., how their phones work). Critically, edge replacement does not say that there's necessarily a single "something" that stands for causal mechanisms. Rather, it provides a rational way that people could generate and reason about a hypothesis space of possible mechanisms—possible "somethings," which may relate to each other in complex ways. An important facet of this framework is that it allows a learner to posit nonobvious causes or side effects that are not part of the initial representation of events. This is also a way to scale up reasoning from the toy examples that we (and others) have often provided to more real-world contexts. We refer the interested reader to Buchanan and Sobel (2014; see also Buchanan et al., 2010) for the computational details and to Buchanan and Sobel (2011), Erb, Buchanan, and Sobel (2013), and Ahl and Keil (2017) for empirical data consistent with this framework.[2]

The important point for our purposes is that the edge replacement framework provides a computational-level description of how to generate possible hypotheses for how events are causally related to each other and a way of evaluating the likelihood of those possible hypotheses. So although psychologists build incredibly simple causal graphical models in their studies, these models could stand for much more complex and more contextualized systems. This framework also addresses the concern that we as human beings almost never construct models that precisely represent the actual causal structure of the world. For example, we tend to like simple models, like the one that states that having a cold causes sneezing. But the actual causal structure that describes this relation is *really* complicated. The edge replacement framework provides a way of explaining how the shorthand of Cold → Sneezing could stand for this actual causal structure.

Besides being of theoretical interest as a potential objection to the causal graphical model framework, discussions of nonindependence are important because they force us to consider the context in which people are reasoning about causal relations. As mentioned in chapter 2, and as illustrated by the cellphone example, context may affect how both adults and children reason about the data they are given. This is relevant to our discussions of the development of scientific thinking because one of the big differences between measures of young children's causal reasoning abilities and measures of their scientific thinking abilities is their use of context (or their *contextualization*). As we discuss in more detail in chapter 4, causal reasoning studies often use blicket detectors or other machines or vignettes in which the mechanisms are stripped down and in which relatively few hidden causes are assumed or implied. Scientific thinking studies, in contrast, can have many hidden causes or can imply many different types of hidden structures. We believe that these differences may explain some of the perceived conflicts between the results using these different types of tasks.

As an aside, the idea that reasoners are constantly thinking about potential mechanisms, particularly when they observe probabilistic events, could lead to interesting explanations for various cognitive phenomena. One example of such a case is the *gambler's fallacy* (Tune, 1964; Tverksy & Kahneman, 1971), in which reasoners misunderstand the independence of random events. People tend to think that, because a certain number has not hit on a roulette wheel for a while, it's "due," hence more likely to come up on the next spin. Similarly, if you hear that a friend won the lottery, you assume it's the friend who has been playing each week for ten years as opposed to the friend who played for the first time. In both cases, as in the cellphone case above, people posit (faulty, incorrect, improper) mechanisms that lead to these inferences.

Another example is the *illusion of explanatory depth* (Rozenblit & Keil, 2002), in which adults believe that they understand a concept better than they actually do. In studies of this phenomenon, adults are asked how well they understand how common objects work, like bicycles or toilets. Most people rate their knowledge fairly high. But these ratings drop precipitously after they're asked to actually explain how a bicycle or a toilet works, which most adults are unable to do. Second graders also show this effect (using modified procedures, see Mills & Keil, 2004). The tendency to posit extra "somethings" in causal models as stand-ins for our mechanistic knowledge can also explain these results, because our constant thinking about mechanisms might lead us to believe that we understand causal structures more than we actually do (see Alter et al., 2010, who make a similar argument, although not relying on this computational framework).

## Where Does a Concept of "Cause" Come From?

The previous section was concerned about objections to the causal graphical model framework from a line of research in adult cognitive psychology. We suggested that there are ways of integrating these objections into rational constructivism, which uses this framework to formalize its representation of children's learning. An equally important concern comes from studies on the causal reasoning capacities of very young infants, which raise the question of where a concept of "cause" comes from. Are infants born perceiving the world in terms of causal structure? Or are the reasoning mechanisms available to preschoolers and school-age children different from the ones available to infants? And if the latter view is correct, then how should we describe the nature of this discontinuity?

Rational constructivism answers these questions by suggesting a developmental continuity. In this view, children are born with some aspects of theories—representations of causal structure—which means that the inherently causal nature of these representations is available to children from the earliest ages. Some versions of this theory have suggested that, while the nature of these representations changes with development, the basic concept of what a cause is stays the same. This aspect of this theoretical approach is similar to contemporary nativist theories, which suggest that causality is piece of core cognition (e.g., Carey, 2009).[3]

The evidence for this position comes from studies on the inferences that infants make about causal relations among events. Babies around 6 or 7 months old can recognize various aspects of contact causality, including registering certain configurations of perceptual features as causal (e.g., Leslie & Keeble, 1987; Oakes & Cohen, 1990) and understanding the relations between the speed and force of objects (e.g., Kominsky et al., 2017). They can also recognize that there are differences between self-moving agents and inanimate objects, namely that the former can cause the latter to move but not the reverse (e.g., Saxe et al., 2005; see also Muentener & Carey, 2010, and Newman et al., 2010, for related findings on slightly older infants).

In summarizing these findings, Carey (2009) writes that the fact that "infants take into account the causally relevant properties of the participants in events in their representation of these events provides evidence that they are making causal attributions" (p. 488). And we agree. By about 6 months of age, infants have a variety of sophisticated causal reasoning capacities. But where we disagree is that the concept of a cause is innate in Carey's sense of the word: shared with other animals,[4] not learned, and unchanging throughout the process of development.

We would like to suggest a different possibility. We question whether infants at very early ages (i.e., newborns) have the capacity to engage in causal inferences, as opposed to merely registering statistical regularities among events (following Haith et al., 1993; Kirkham et al., 2002). As we suggested in chapter 2, one of the big challenges for a description of causal reasoning is that learning the associative relations among events gets you pretty far.[5] That is, it is possible that infants represent causal relations first by recognizing and using associative information. In the second half of the first year of life, they move to appreciating those statistical regularities in terms of more genuinely causal (not merely associative) structure. Those representations now support the interventions that are part of the causal graphical model framework.

This hypothesis is consistent with a multitude of findings that suggest that, while infant causal perception undergoes development related to information processing (e.g., Cohen et al., 1999), it is mostly in place by the second half of the first year of life (Leslie & Keeble, 1987; Oakes & Cohen, 1990; Saxe et al., 2005, 2007). Indeed, most of the citations that Carey (2009) uses to argue that causality is an innate concept come from infants in the second half of the first year of life.

What kind of evidence would differentiate between the claim that infants register relations among events in terms of their statistical regularity and the claim that they register events in terms of causality? One is their use of the Markov assumption. Some work that we did on infants' causal inference suggests that 8-month-olds respond to the backward blocking sequences we described in chapter 2 in the same manner as 4-year-olds (Sobel & Kirkham, 2006), but 5-month-olds treat these sequences in terms of their associative relations (Sobel & Kirkham, 2007).

Another way to differentiate between these claims is with reference to the fact that the causal graphical model framework supports interventions. In this framework, interventions force variables in the network to take on a particular value. Psychologically, interventions are typically actions on the world that bring about a situation. It is important to consider what role the development of action abilities might play on the emergence of a causal reasoning system that places such an emphasis on actions.

During the second half of the first year of life, infants develop the capacity to coordinate their actions in complex ways. For instance, at 7 months old, infants can reach for an object directly in front of them, but they struggle when they have to coordinate their actions with barriers or obstacles in the environment to make that reach. As an example, Diamond and Gilbert (1989) showed that infants develop the capacity to reach for objects in a coordinated fashion between the ages of 7 and 11 months. As part of this study, they placed a LEGO brick in an open, transparent box. When the brick was placed in the middle of the box, infants across all of the ages that they tested could reach for it without difficulty. Similarly, when the brick was placed outside of the box, adjacent to it, infants at all ages could reach for it without difficulty. However, when the brick was inside the box, adjacent to the side of the box closest to the infants, the youngest infants struggled with their reaching. There was a clear developmental trajectory in their ability to reach for the brick in this configuration between 7 and 11 months. The trouble with this particular setup is that infants have to coordinate their movement—they have to reach into the box from the top and then change the direction of their reach to get to the brick, which requires inhibitory demands that are developing during this time (Diamond, 1991a, 1991b). Critically, even the youngest infants had no difficulty when the brick was in the same position, but the box was tilted, such that they did not have to change the direction of their reach.

These data suggest that, during the second half of the first year of life, infants are developing domain-general capacities, which facilitate their ability to coordinate their actions to accomplish a goal. We speculate that, in addition to this developing inhibitory control, the concurrent maturation of the supplemental motor cortex and dorsolateral prefrontal cortex (the neural systems that Diamond, 1991a, hypothesized were responsible for pattern of results described above) might serve to integrate children's representation of the structure of their own actions with their understanding of statistical regularity, allowing them to construct causal representations. That is, prior to infants being able to coordinate the ways in which they intervene on the world (e.g., by reaching and grasping), they may see the world in terms of statistical regularity. Being able to generate actions on the world may then afford the infant more of a fully causal interpretation (i.e., I can change the state of the world through my actions).

An analogy for this idea comes from an interpretation of a body of research on children's capacity to understand that actions are goal-directed. As an example of this kind of research, a study by Woodward (1998) habituated 5- to 9-month-old infants to a display in which a hand reached to one of two objects. After habituation, the locations of the two objects were switched. Infants saw the hand reach for either the same object (via a new path) or the other object (via the same path). Overall, infants looked longer when the hand reached for the other object (that is, when it reached to the same location as before). Woodward interpreted these results as demonstrating infants' abilities to encode the actions of others as being goal-directed: The agent (represented by the hand) intended their action to reach a particular object and not simply to reach toward a particular point in space.

What is interesting about this capacity is that younger infants—3-month-olds—do not interpret this experiment in the same manner. Instead, they look equally to the same object (new path) and same path (new object) events. This suggests that they do not encode the relation between the actor and the goal of the action (see Sommerville et al., 2005, footnote 1). But Sommerville et al. (2005) argued that these very young infants could recognize this relation, they just needed the experience to do so. Their argument is similar to the one we made above regarding the link between self-generated actions and an understanding of causality. Specifically, 5-month-olds have enough coordinated motor development to reach for objects and grasp them themselves. Their reaching ability is not adult-like, but they can begin to manipulate objects in their environment. Three-month-olds? Not so much. They can flail their arms about, but they cannot make a coordinated reach for an object. This difference in motor abilities could potentially explain why 5-month-olds in Woodward's experiment interpret the hand's reach as goal-driven, but 3-month-olds do not: They themselves cannot coordinate reaches for objects in a goal-directed way, so they do not appreciate goal-directedness in others.

This argument implies that, if 3-month-olds were given experience with manipulating objects, they would become able to register that reaches are goal-directed. That is, the infants' own reaching actions could help them to appreciate that others' actions have goal-directed structure. To test this, Sommerville et al. (2005) sewed Velcro strips onto teeny-tiny mittens and gave the 3-month-olds two minutes of experience manipulating a Velcro block world with the mittens on. When using these "sticky mittens," as mentioned in chapter 1, these babies' flailing attempts to pick up objects would actually successfully allow them to manipulate the block—or at least to move it around in space. After the mittens experience, the infants were presented with a version of Woodward's experiment described above, in which the actor was wearing the same kind of Velcro mitten as the infants had worn and reached for a Velcro block that looked similar to the one that the infant participants had themselves played with. These infants picked up on the goal-directedness of the reach, and looked longer after habituation when the actor reached for the new object than for the same object on the new path—just like the 5-month-olds did.

One can interpret these data as suggesting that the motor experience of being able to reach for objects provides infants with the ability to register that those actions are goal-directed, and thus to infer that similar actions by others are goal-directed. While that understanding naturally develops between the ages of 3 and 5 months as infants learn to reach and grasp objects, it can be accelerated artificially through experience.[6] In turn, this suggests that a concept of causality might develop out of infants' growing understanding that their actions can affect the world in a variety of ways.[7] Given these findings, we suggest that infants are not necessarily born with the concept of "cause"; they develop this concept during the first year of life, based on their observations of the world and especially on their experiences with acting on the world.

Although we think that the issue of where a concept of "cause" comes from is interesting in its own right, we highlighted this issue here for another reason. In chapter 4, we suggest that the causal reasoning system that children have during the preschool years develops during the early elementary-school years; it changes in fundamental ways that might facilitate success on some scientific thinking measures. As such, it is important to conceptualize children's causal reasoning as a system that itself undergoes significant change throughout the lifespan, starting in infancy.

## Active Learning

An interesting set of findings in from the psychological literature shows that people overemphasize data that they themselves generate. Much like the infants in the sticky mitten experiment described above, both adults (Lagnado & Sloman, 2004; Sobel & Kushnir, 2006; Steyvers et al., 2003) and children (Sobel & Sommerville, 2010) benefit from engaging in actions themselves when learning causal structure, particularly when compared to cases in which they simply observe others' actions. But young children, in particular, can be overinfluenced by their own actions; they have been shown to rely on a single data point that they themselves have generated as opposed to considering the aggregate data that they observed (Kushnir & Gopnik, 2005; see also Sobel & Letourneau, 2018). Moreover, numerous studies suggest that various aspects of young children's cognition are influenced by their movement and actions (see Lillard, 2004, for a review).

This tendency to favor data that one has generated oneself poses a problem for the causal graphical model framework as a description of children's inferences. As we suggested in the previous section, an important aspect of this framework is intervention—the idea that actions on the world can support different inferences than simply observing data. Interventions are important because data from interventions can provide conditional independence information that may be difficult to detect by observation alone. For example, we can observe the correlation between a rooster crowing and the sun rising, but making the rooster crow at 2 am will not cause the sun to rise (example taken from a 1913 manuscript of *Theory of Knowledge* by Bertrand Russell, as cited in Pearl, 2000). However, the framework does not distinguish between data that come from one's own interventions and data that come from someone else's. Given that learners do seem to process these different types of data differently, this tendency could undermine the ability of the framework to serve as a model for how learning works.

Some hints for how to address this problem comes from work on adults, who—although they do tend to be influenced by self-generated data during causal learning—might be doing so for rational reasons. For example, Sobel and Kushnir (2006) suggested that the reason adults learn better from their own actions than from observing the same data generated by others is that, when one generates actions, one can test one's own hypotheses. This study showed that adults were better learners when they generated their own interventions on a causal system than when they observed another adult generate the same interventions, in support of the work cited above. But this study also showed that the number of critical interventions that learners made (interventions that generated the relevant conditional probability information necessary to learn the particular causal structure) predicted adults' learning, but only when they generated their own actions. When adult participants observed the same sequence of data, there was no relation between the quality of the data sequence and adults' learning. These results suggest that, while adults may appear to simply favor data generated by their own interventions, these data may be genuinely more informative for them, hence it may not be an error to rely more heavily on them.

To test this hypothesis with children, Sobel and Sommerville (2010) presented 4-year-olds with a situation in which they could learn causal structures from their own actions as opposed to observing an adult play with the machine. Like the adults in Sobel and Kushnir (2006), children were better at describing the causal structure in the former condition, and were especially better when those actions discovered novel information. We'll discuss findings like this more in chapter 4, but for now we want to note that there is also neuroscientific evidence that suggests that self-generated action recruits different neural systems for learning in children. James and Swain (2011) introduced 5-year-olds to novel actions on novel objects, which the children either performed themselves or observed another person perform. Each action was given a novel label. When these children were exposed to the stimuli and novel labels, they showed greater recruitment of motor regions for the labels that named actions that they had performed themselves as opposed to the ones they had observed. These data suggest that children conceptualize their own actions differently than the actions of others (see also Sommerville & Hammond, 2007).

This work has close ties to a body of literature examining the role of play in development, much of which asks a similar question: whether children learn better when adults provide them with direct instruction or when they are allowed to explore the world on their own (Klahr, 2000; Klahr & Nigam, 2004; Kuhn, 1989; Kuhn & Dean, 2005). With respect to this question, there is a general notion that *guided play* is a good compromise between these two approaches to learning (Mayer, 2004; Weisberg, Hirsh-Pasek, et al., 2013, Weisberg et al., 2016). For example, in a large-scale meta-analysis of learning outcomes, Alfieri et al. (2011) argued that unguided ("free") play typically resulted in worse learning outcomes than direct instruction, although adding some kind of adult guidance to the play improved learning.

Similarly, research on *scaffolding* (Wood et al., 1976) suggests that parents (or adults more generally) can teach children to solve problems by presenting them with the parts of the task that they believe children are capable of doing successfully. The adults then build them up to engage in more and more complex behaviors (e.g., Connor & Cross, 2003; Meins, 1997; Pratt et al., 1988; Wood & Middleton, 1975; see Mermelshtine, 2017, for a review). Of particular importance is the contingent interaction that adults generate during these interactions. For scaffolding to be effective, adults must react to children's actions to facilitate learning. Scaffolding allows children to take an active role in completing aspects of the problem that they understand, while also being exposed to more parts of the problem that they might be able to tackle the next time it is encountered. Research on both guided play and scaffolding assert a crucial role for the child's own self-guided actions in learning.

We investigated this issue as part of a large-scale study of parent-child interaction while families played at a gear exhibit at three children's museums around the United States (Callanan, Legare, Sobel, et al. 2020). After families played together, children were asked to participate in a set of activities that measured their causal knowledge about gears, while parents completed demographic questionnaires and other measures about their beliefs toward science and play.

While this work examined many facets of children's interaction and learning, here we want to highlight one set of findings that is pertinent to this discussion: How did children explore when they played? To examine children's actions in detail, we looked for a set of behaviors that we thought a priori would be important to learning. We called these behaviors *systematic exploration*. These behaviors primarily included building a gear machine and then testing that machine by spinning it. Perhaps unsurprisingly, we found that the proportion of systematic exploration behaviors that children engaged in increased with age. Independent of age and of many other social factors like household income and parental education levels, this proportion was related to children's persistence in problem-solving when things went wrong during their play. This proportion was also related to children's abilities to successfully engage in causal reasoning about gears, as measured by our post-exploration activities. Systematic exploration generated by children themselves therefore seems to be an important set of behaviors for children's learning.

On the surface, there is a clear interpretation of this result (and of many of the results that show a benefit for children's learning from their own interventions) that fits with the rational constructivist framework: How children explore the exhibit relates to the way in which they interpret the exhibit content. Because rational constructivism is committed to children's learning (i.e., updating the causal representations of their theories), it should not be problematic that children can do so through their active engagement with the world, or even that such exploration is an effective way for them to learn. So although it may be difficult for the causal graphical model framework to take this into account, we believe that this framework could be amended to include information about where the intervention data is coming from and to reflect more accurately children's (and adults') biases for their own actions.

But this is only half of the story. The idea that self-generated action is an effective way to learn could potentially be accommodated by rational constructivism, but it is more difficult for rational constructivism to account for the ways in which parents' and children's interactions at this exhibit affected the children's learning. Specifically, our analyses focused on the explanations that parents provided while their children were playing. The overall proportion of parents' causal language had no bearing on how children played; it didn't relate to the proportion of systematic exploration that children engaged in. Nor did parents' proportion of causal language relate to children's causal reasoning when they were tested on their understanding of gears. But what did relate to children's systematic exploration was the *timing* of when parents generated causal language. Specifically, if parents began their causal utterance while children were connecting gears together in their exploration, children were more likely to complete this systematic exploration sequence than if parents began the causal utterance after the gears were already connected, or before or after this entire sequence. Although parental language in the aggregate might not relate directly to children's causal reasoning, parents do influence whether children engage in behaviors that might be related to their causal learning.

To accommodate findings like these, rational constructivism needs to be integrated with the idea that children are not just internally constructing meaning, representations, and concepts through their actions on the world. Rather, learning from self-generated action occurs within the context of an interactive social framework (e.g., Vygotsky, 1978; Wood et al., 1976). We suspect that the best way to reconcile these findings with the tenets of rational constructivism is to suggest that children have internal mechanisms for learning, but that those mechanisms are affected by various different kinds of social and cultural interactions.

Unfortunately for this suggestion, sociocultural theories have typically been cast in opposition to constructivist approaches (see e.g., Cobb, 1994). There are two general reasons for this. The first is theoretical; the two theories posit different mechanisms for how learning is accomplished. Specifically, there is disagreement as to whether knowledge is represented internally by the mind or co-constructed by the interaction (see e.g., Rogoff, 1990); we discuss this idea in more detail below. The second reason for the opposition between these theories is more practical, in that the views often lead to different educational approaches, which emphasize different cultural practices (e.g., Cobb, 1994; Van Oers, 1990). Nevertheless, because both theoretical approaches have important insights to offer, we believe that constructivist approaches can—and should—be integrated with sociocultural approaches to cognitive development (see Callanan & Valle, 2008; Legare et al., 2017). We begin to take up this challenge in the next section.

As an important note, we think it is unlikely that the parents in our museum study explicitly timed their causal utterances to match children's behaviors during their interaction. This study is correlational and does not imply that parents controlled when they generated causal language in an attempt to facilitate children's causal learning. No demographic factor that we measured related to whether parents generated causal language at that particular point in time. In contrast, parental education level related to whether they generated causal language *after* children completed a systematic exploration behavior. This suggests that some parents might notice the behavior itself and react to it with causal language, but this reaction did not relate to children's exploration. Instead, parental interaction with children during play might create opportunities for parents to help children generate co-constructed interpretations of their own actions.

## The Social Nature of Learning

As noted in the previous section, another source of challenges to rational constructivism comes from a large literature on the *social* nature of learning, noting that learning is largely (perhaps essentially) a social process. As such, blicket detector experiments and similar bodies of work, on their own, do not provide a complete picture of children's learning, hence the causal graphical model framework may not be able to fully account for this process. Here, we review two bodies of work that examine children's learning in social contexts. In both cases, we attempt to synthesize these results together with rational constructivism as well as try to explain how rational constructivism needs to evolve to take these results into account.

### Trust in Testimony

For the past fifteen years, in response to research on children's causal learning, a body of work within cognitive development has argued that learning from intervention and observation is not particularly efficient. While one can (re)discover causal concepts from observation and intervention, many of those discoveries take quite a bit of time and data. As noted in chapter 1, preschoolers have a great deal of knowledge about the physical, psychological, and biological worlds, and it is unlikely that they learn all of this information from observation and intervention alone. More to the point, children possess many beliefs for which they have no evidence. Children have never seen a germ or a ghost, and yet they tend to believe that both of these entities are real. Further, they have different kinds of beliefs about these two particular entities, and they have differential beliefs about scientific and religious entities in general (e.g., Harris & Koenig, 2006).

The thrust of this argument is that children cannot gather enough evidence about the world simply through observation or solitary play (Harris et al., 2006; see also Keil, 2010); they are not "stubborn autodidacts" (Faucher et al., 2002, p. 341; Harris, 2002). Rather, they need social transmission of knowledge for learning to occur. Specifically, children's beliefs about the world are acquired not just from observation and interaction with the environment, but from social transmission and communication (e.g., Csibra & Gergely, 2009; Harris, 2002, 2012; Vygotsky, 1978). This process is often referred to as children's "trust in testimony" (Harris, 2002), and studies in this area have shown how children learn directly from others' explicit verbal statements (see Sobel & Finiasz, 2020, and Tong et al., 2020, for two meta-analyses of this literature).

Children are also influenced by the pedagogical nature of the environment they are in, and they learn more when others are teaching them directly and presenting them with data in a way that ensures their attention (e.g., Butler & Markman, 2012, 2014). Such pedagogical inferences can have a cost, however. In particular, children assume that teachers tell learners all they need to know, so when teachers commit sins of omission in their teaching, children's exploration is negatively affected (Bonawitz et al., 2011), as are their beliefs about the teacher's credibility (Gweon et al., 2014). These data generally suggest that, while children can learn from observation and interaction with the world, they are also learning from the social interactions that they have (Harris et al., 2006).

These studies illustrate how this research tradition not only emphasizes communicative norms, like word meanings, but also constraints on the acquisition of social knowledge, like cultural conventions and norms. Knowing from whom to learn and under what circumstances to learn, and then actually learning from others, is the basis of cultural knowledge and transmission (e.g., Bergstrom et al., 2006; Harris & Koenig, 2006; Kline, 2015; Mascaro & Sperber, 2009).

Given this, an important question is whether children simply believe what others tell them, or whether they learn from others in a more judicial manner. That is, are children credulous or skeptical, and does their credulity or skepticism change with development and experience with informants? Traditionally, young children have been thought to be overly credulous, trusting what others tell them (e.g., Piaget, 1930). This assumption, however, has not been widely supported. Across many experiments, it has been shown that preschoolers can track the reliability of other informants and use that information to make judgments about their epistemic competence and about whether to learn from them. For example, preschoolers will not learn a label for a new object if that label is said by a speaker who has previously labeled objects inaccurately (e.g., Clément et al., 2004; Koenig et al., 2004; Koenig & Harris, 2005). *Selective learning* capacities are an important part of what makes children such good learners.

But how do children's selective learning capacities develop? One major theory about children's trust in testimony posits that children are credulous early in development—they have what Jaswal and colleagues describe as a “default bias to trust” (Jaswal et al., 2010, p. 1541). In this view, what develops is the capacity to discount inaccurate testimony in favor of other sources of information, such as one’s own direct observations or testimony from a more reliable source. That is, learning from others requires assessing whether another person is providing accurate information and discounting that information when one judges the other person’s statements to be false. Children’s developing cognitive control thus underlies their capacity to inhibit others’ inaccurate testimony (Harris et al., 2018; Jaswal et al., 2014). A similar—although separate—proposal was put forward by Mills (2013), who suggested that the developing capacity to track others’ reliability involves the development of the metacognitive capacity to register skepticism. This latter proposal also suggests that children have an initial default bias to trust, but what develops is a broader and more metacognitive conception of skepticism—that others’ statements can be false, regardless of their intention.

Both of these proposals have a great deal of merit. While 3- and 4-year-olds are capable of discerning between reliable and unreliable sources of knowledge, a study by Krogh-Jespersen and Echols (2012) suggested that 2-year-olds trusted an accurate, inaccurate, and ignorant informant equally. The 2-year-olds did not discount the inaccurate informants’ information, although older children did. This study is important because it supports the argument that children have a default bias to trust.

But it is also important because it used a somewhat different method for studying this topic. So far, the work we have discussed that examines whether children can learn selectively introduced children to multiple informants who differed in some way. In some cases, they differed on epistemic competence (e.g., one was an accurate source of knowledge and the other was not), and in other cases they differed on a social or physical cue that could indicate differences in epistemic knowledge (e.g., one was an adult while the other was a child, or one spoke in the child’s native language with a native accent while the other spoke with a nonnative accent). But this is a not a particularly ecologically valid paradigm. How often do children observe two people generating different labels for the same object, particularly for familiar objects (that’s a stapler, no, that’s a cup)? A more ecologically valid method might be to introduce children to a single informant and vary the epistemic or social information that that informant generated between groups of children.

That is the method used in Krogh-Jespersen and Echols (2012). Children in that study were asked about the testimony of only one informant, who in the past had been accurate (or inaccurate or ignorant). These researchers found that 2-year-olds tended to use the labels generated by an interlocutor, regardless of that individual's epistemic status. Sabbagh and Baldwin (2001) also used a single-informant paradigm and found that the 3- and 4-year-olds they studied learned word meanings better from a knowledgeable than from an ignorant informant. So 2-year-olds might have more of a default bias to trust, but there is development away from that default bias during the preschool years.

That very young children can learn selectively from conflicting informants, but not from a single informant, speaks to the possibility that selective learning is indeed governed by inhibitory control factors. It is more difficult for young children to integrate past information with an informant's current statement than it is to judge between two informants who provide this information at roughly the same time. In support of this argument, some studies suggested that children were more likely to respond selectively when they observed two informants disagree about the objects' labels, as opposed to when they observed one informant simply generate inaccurate labels (Koenig & Woodward, 2010; Vanderbilt et al., 2014). More directly, Jaswal et al. (2014) showed that preschoolers' capacity to discount inaccurate information—particularly when it conflicted with their own observations of a situation—correlated with their inhibitory control capacity, measured by a separate battery of tasks. Inhibitory control (or cognitive control more generally) does seem necessary to learn from others selectively.

We agree that inhibitory control or other executive function abilities are necessary to make accurate judgments about others' epistemic competence; in fact, this makes a nice parallel to our arguments above about motor skills being necessary prerequisites to infants' abilities to understand others' goal-directed actions. However, it is questionable to us whether inhibitory control is *sufficient* for explaining these findings. In particular, all of these studies ask children to make a choice, for instance between the labels provided by two speakers, or between two novel words as the label for a novel object. These measures seem reasonable for older preschoolers, but they might introduce demand characteristics, particularly for children at the younger end of the age range, like the 2-year-olds discussed above.

There are other techniques for eliciting responses from children this young, many of which rely on measuring the amount of time children spend looking at different displays (called looking time paradigms). To illustrate how these methods can be used to determine how toddlers decide whom to trust, we (Luchkina, Sobel & Morgan, 2018) presented 18-month-olds with speakers who were either accurate or inaccurate. Then, one of those speakers labeled a set of novel objects with novel labels. This speaker then showed the child one of those labeled novel objects together with another novel object, which the child had not yet seen. Children were asked "Where's the [label]?" where the label used was the novel label that the speaker had previously generated for the first object in the test set. When that speaker had been accurate, the 18-month-olds stared more at that object as opposed to the other one, suggesting that they had accepted this label. When that speaker had been inaccurate, they stared equally at the two objects, suggesting that they did not know which object the speaker was labeling. That is, children seem able to track the reliability of others at very early ages. They are not just biased to trust, but rather can be evaluative. However, they may have difficulty displaying these abilities in all cases, illustrating that children's developing cognitive control abilities still have an important role to play.

There is an important control in this study, which further illustrates these children's abilities to think critically about the information they receive. In this control condition, we ran essentially the same procedure with a new group of 18-month-olds but changed one aspect of the design. Usually, in measures of selective learning, children are shown familiar objects, and the informant says "That's an X," where X is either an accurate or inaccurate label. That's what we did in the study described above. But in our second experiment, each informant asked "Is that an X?," where the label that the informant generated was either accurate or inaccurate. The critical idea is that questions do not convey information about epistemic competence, so using an inaccurate label in a question is not a sign of incompetence, like it is in a statement. Children registered this, and they did not learn from either informant in this case. This is important because, when the speaker generated the accurate label in question form, they have generated the same associative information as in the statement case; the speaker has said the correct label in both cases. The fact that children learn differently from statements than from questions suggests that even very young children are not simply using the associative properties of the language they hear. Rather, they undergo a more sophisticated process of evaluating the pragmatics of the situation in order to judge a speaker's accuracy.

This conclusion is related to a set of findings generated by Sabbagh and Shafman (2009), using a single-speaker paradigm. The 4- and 5-year-olds in this study could both remember a knowledgeable speaker's label for a novel object and also use that information to make an inference about meaning; they could state that the object had that label. When a speaker expressed more ignorance about the label, the children in this study could remember the speaker's label, but did not necessarily endorse the object as having that label. As in Luchkina et al. (2018), these children did not simply use a speaker-object-label association to learn the meaning of words (see also Mangardich & Sabbagh, 2018). In addition to inhibitory control, which plays a role in helping children consider a variety of information in their decisions about whom to trust, there are also important situational constraints on learning. Causal knowledge is necessary for judging whether to use the information that others generate.

To appreciate this more integrative perspective, consider work by Kushnir et al. (2008), who varied whether an informant was knowledgeable or ignorant about a novel toy (an epistemic constraint) and also whether the source was permitted to use that knowledge in performing an action (a situational constraint). Specifically, they showed 3- and 4-year-olds two puppets and a machine (a blicket detector, although it was never called that) and a set of blocks. Children were told that one of the puppets knew which blocks activated the machine and the other puppets did not. Across their experiments, children observed the puppets either intentionally pick blocks to make the machine go, observed the puppets pick blocks while blindfolded, or picked two blocks themselves for the puppets to put on the machine. In all three cases, the two blocks were placed on the machine together and the machine activated. Children were asked which block was more likely to be efficacious. Only when the puppets picked the blocks intentionally did children say that the knowledgeable puppet's block was the efficacious one; in the other two cases, they responded at chance.

This study demonstrates that children do not seem to be relying on a single cognitive capacity to make decisions about the value of information. Rather they are integrating both situational and epistemic knowledge when tracking and making inferences about others' accuracy. In turn, this implies that both developing cognitive control abilities and developing knowledge about epistemic states (in this case, the knowledge that one has to have perceptual access in order to intentionally choose an efficacious object) are necessary to explain children's behavior.

Similarly, other work has shown that preschoolers still have access to low-level associative learning mechanisms underlying certain aspects of their ability to track others' reliability (Luchkina et al., 2020), which can be described by dual process accounts of selective learning (e.g., Hermes et al., 2018). Just as children have early-emerging capacities for statistical learning and for appreciating regularity in data, children also have early-emerging capacities to track others' accuracy. Those capacities become integrated with new information and can be constrained by other domain-general cognitive processes like children's developing cognitive control. As children's knowledge changes, it has cascading effects on how they evaluate and learn from the information generated by others.

In general, then, this body of work on children's understanding of testimony and the relative accuracy of different speakers illustrates one way in which social factors influence children's learning. Although children do learn from playing alone, they also learn from playing with others (e.g., Gauvain & Rogoff, 1989). They learn from observing others (e.g., Williamson et al., 2010), but they also ask questions and seek explanations (e.g., Callanan & Oakes, 1992, Chouinard, 2007; Frazier et al., 2009; Mills et al., 2010; Mills et al., 2011). And while children learn some ideas, concepts, and skills readily from direct instruction (e.g., Klahr & Nigam, 2004; Markson & Bloom, 1997), other ideas, concepts, and skills take significantly more time, particularly when that instruction runs contrary to direct observations that children themselves make (e.g., DiSessa, 1993; Vosniadou & Brewer, 1992).

It is unlikely that such learning is governed by two separate processes—one for (relatively solitary) causal learning and the other for social learning. Understanding children's learning thus requires a framework that can encompass both of these situations. Importantly, such an account should make systematic predictions regarding when and from whom children learn and when and from whom they do not. Moreover, such an account should predict what happens when multiple sources of information interact or conflict.

We believe that rational constructivism could serve as such an account because it can be extended to capture children's selective social learning

(Sobel & Kushnir, 2013). Much like children have causal reasoning systems dedicated to integrating their existing knowledge and beliefs with observed data, children also have domain-general capacities to track statistical information generated by other people (e.g., accuracy information). They can integrate that accuracy information into what they know about others and about the situation to make inferences about the reliability of a person as a source of novel information. Rational constructivism should thus be expanded to consider not only the kind of information that children receive, but also the social contexts in which that information is presented to them.

## Integrating Cultural and Sociocultural Approaches with Causal Learning

An appeal of rational constructivism is that it emphasizes how learning mechanisms are part of children's cognitive development: Children observe, interact with, and interpret the world around them. Children also process social information as they observe and interact with the world. In this theoretical perspective, learning is often reduced to information processing, following metaphors from symbolic AI that have dominated cognitive science (e.g., the mind is a computer; Miller, 1956; Turing, 1950). Turing actually considered this, suggested that "instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child? . . . Presumably the child-brain is something like a note-book as one buys it from the stationers. Rather little mechanism, and lots of blank sheets" (1950, p. 456). Turing did not use causal graphical models or Bayesian inference (although some suggest that Turing used Bayesian methods in his codebreaking, see Good, 2000), and the algorithms that researchers use to model learning today are different. But this idea—that we can model learning as the internal workings of the mind, particularly of a child's mind—has not fundamentally changed, at least in certain parts of cognitive science.

In contrast, sociocultural theories of cognitive development suggest that children's learning is not mainly about information processing. Rather, learning should be understood as being contextualized in social activities and cultural practices (Gauvain & Perez, 2015; Vygotsky, 1978), meaning that there is not a specific algorithm that results in particular outputs when given particular inputs. For example, Rogoff (2003) describes children as developing "within cultural communities. . . . Their development can be understood only in light of the cultural practices" (pp. 3–4). In this account, developmental processes are dynamic (both nonlinear and multidimensional)[8] and are dependent on cultural practices and norms, and therefore may not be able to be captured by the modeling frameworks that we have been espousing.

Given this, a fundamental challenge for rational constructivism is that children's learning is embedded in their culture, which may bring about profound individual differences based on what Gutiérrez and Rogoff (2003) call a "repertoire of practice." For instance, we argued above that children and adults learned better from their own actions. But the studies that we used to draw that conclusion used samples of convenience composed mainly of white, upper-middle-class families in so-called Western cultures. Learning from self-generated actions might have been encouraged by those participants' repertoires or by the researchers' beliefs that their own cultural practices are the norm (Medin et al., 2010).

We do not yet know whether these findings generalize across other practices that do not emphasize self-generated action. Rogoff and colleagues, for example, studied Mexican communities in which children are included in a range of activities from which children in European-American communities would be excluded. The cultural practices of such communities might promote different observational practices, attentional practices, or information-seeking practices on the part of children, which are all related to their learning (e.g., Mejía-Arauz et al., 2005; Rogoff et al., 2003; see also Correa-Chávez & Rogoff, 2009).

In general, work on children's causal reasoning and scientific thinking has a markedly WEIRD focus, tending to study cultures that are Western, Educated, Industrialized, Rich, and Democratic (Henrich et al., 2010). One exception is work by Shultz (1982), who tested a group of children in Mali on two of the same experiments he presented to children in Montreal—though we note that it is potentially problematic to assume that a task developed within one cultural context will straightforwardly translate to a different context (see Gutiérrez & Rogoff, 2003). Shultz's experiments tested children's understanding of what he called "generative transmission," the idea that causal relations require some form of mechanistic force to transfer from cause to effect. This work—and the broader idea that there should not be many cross-cultural differences in certain kinds of causal reasoning—was formed as a response to archaic beliefs advocated particularly by Levy-Bruhl (1926) about the absence of causal reasoning capacities within various non-Western cultures.[9] Critically, in direct opposition to this view, Shultz (1982) found numerous similarities between the two samples and emphasized the importance of this similarity: "This provides considerable support for the notion of psychic unity as applied to causal reasoning, and is consistent with the hypothesis that causal reasoning is a very basic mental function not easily influenced by even substantial cultural variations" (p. 39).

Wente et al. (2019) used a similar approach to support this conclusion. These researchers tested a group of children and college-age adults in Peru on whether they could infer the causal structure of a system that presented either conjunctive or disjunctive causality (similar measures as used in Lucas et al., 2014). They found similar results in their Peruvian sample as in their Western sample: Children and adults in both cultural contexts made similar inferences about the disjunctive causal relations, but both groups of children were more likely than both groups of adults to succeed at making inferences about the conjunctive ones.

Although this study and the one done by Shultz (1982) show some cross-cultural consistency in causal reasoning, we should be cautious about drawing that conclusion broadly. Reasoning tasks developed within a particular cultural frame in order to assess a particular group of individuals might be understood differently in a different frame and by different groups of individuals (see Novaes, 2013; Scribner, 1975).

Other research studies that document the effects of cultural practices on children's learning have investigated science learning in the classroom (e.g., Bang et al., 2012; Hudicourt-Barnes, 2003), classroom-based practices more generally (Chavajay, 2006; Davis et al., 2021; Gurven et al., 2017; Heyneman & Loxley, 1983), and the ways that parents and children interact with each other based on individual differences within a culture. For example, Solis and Callanan (2016) examined two groups of Mexican-heritage families, where the parents differed in their level of formal schooling. Parents and children engaged in a science-related activity about floating and sinking. The children of the parents with higher levels of formal schooling asked more questions about the procedure, while the children of parents with lower levels of formal schooling asked more questions about the conceptual structure of the task. These researchers' explanation of this difference is that the parents with higher levels of schooling took on a "teacher-like role and focused on asking children known-answer questions and evaluating children's performance. . . . In contrast, the parents with basic schooling seemed to engage in the task as co-learners with their children" (Callanan, Solis, et al., 2020, p. 84). Parents had different approaches to this task, which in turn focused their children on different facets of what they could learn.

Although one can imagine computational models that might begin to account for this kind of causal learning across cultures in a manner consistent with rational constructivism (e.g., Werchan et al., 2016[10]), some sociocultural approaches to cognitive development reject such information-processing accounts of development in favor of representing knowledge more within social frameworks (e.g., Rogoff, 2003). We want to suggest a compromise between these positions. Cultural factors change dynamically, which relate to how children learn. And because we generally agree that learning can be fruitfully described within the context of an information-processing approach, rational constructivism must endeavor to integrate more deeply with sociocultural perspectives and to take seriously the ways in which cultural and subcultural contexts may not straightforwardly translate to information processing.

For example, ojalehto and Medin (2015) suggest that different cultural frameworks can synthesize together to create unique causal models of explanation. To illustrate this process, they describe studies of adults in South Africa who were asked to provide explanations about the cause of AIDS. Within these explanations, different explanatory frameworks overlapped (e.g., supernatural entities like witches can use natural causes like viruses and other germs to make one sick; Legare & Gelman, 2008; see also Legare et al., 2012).

Similarly, in the work on gear exhibits in museums described previously, we found that parents' causal talk, particularly the timing of such talk, influenced children's causal actions (Callanan, Legare, Sobel, et al., 2020). How parents feel about children's play and their beliefs about how children might learn through play might relate to the ways that children learn through play as well as the ways they interact with their children during play.[11] For instance, Lancy writes, "Unlike the Euro-American cultural model of childhood where parent-child play may be considered 'essential,' elsewhere adults do not play with children (Lancy, 2007), in large part because it violates the child's independence and takes the adult away from more important activity" (2016, p. 658). That is, different cultures might value learning from play differently, which in turn might affect how children learn from play, or what beliefs they have about learning from play, or the process of learning from parent-child interaction and scaffolding.

Even within the WEIRD culture, in which we personally work, most studies of children's causal reasoning involve children participating in a novel activity setting in a social context like a lab or a museum, which is not necessarily a reflection of their authentic experiences. Children are often brought to a strange place, where a strange adult asks them strange questions, or shows them strange stories, or places them in strange situations.[12] Different children (and their parents) construct different meanings from those social contexts, and these differences should be investigated and documented more systematically. In the absence of a robust body of empirical work on this issue, we support the conclusion drawn by ojalehto and Medin (2015): "Investigating causal cognition from the perspective of complex systems may afford new insights into conceptual behavior. Given that humans are surrounded by complex systems (ecologies, societies, consciousness), this is an area of study that deserves more investigation" (p. 265).

These arguments apply directly to the central theme of the book: the development of children's scientific thinking abilities. Scientific thinking, about both the content and the process of science, is often done in a social context, such as a classroom or a museum. Children may inherently possess some causal reasoning and scientific thinking abilities, but these abilities may not exist independently of the social nature of the world in which they acquired them or of the social nature of the tasks in which they are being asked to express them. For example, Packer and Goicoechea (2000) argued that "any social context—a classroom, for example—is itself the product of human language and social practice, not fixed but dynamic, changing over time, in what we call history" (p. 232). How children are taught science, and what children believe that science is, likely influences the way they learn it, and both of these factors are embedded within children's social and cultural contexts. To be a productive theory of children's scientific thinking, rational constructivism must take these factors into account.

## Questions of Explanations

Questions about cultural differences and about different kinds of adult-child interactions illustrate one set of ways in which the contexts in which children develop affect how they learn and how they interpret their experiences. But these contexts are not simply provided to children, nor are they static; developmental contexts are actively and continually co-constructed by children, their parents, and others in the environment. One place in which this process can be observed in children's construction of and requests for explanations—their active seeking for meaning within their environments and their abilities to talk about their causal knowledge.

Prior work has shown that young children are able to generate explanations about a number of different concepts, and these explanations are often argued to reflect their causal knowledge (e.g., Hickling & Wellman, 2001; Legare et al., 2009; Schult & Wellman, 1997; Sobel, 2004b; Wellman et al., 1997). Children also generate various kinds of causal utterances at relatively early ages (e.g., Hood & Bloom 1979). These utterances are often quite sophisticated, particularly in contrast to their predictive inferences. For example, Bartsch and Wellman (1995) documented that young children learning English in the US generated what they called *false belief contrastives*—utterances in which children indicated that they had a false belief about a state of events—during their third year, well before they solved explicit measures of representational change (e.g., Gopnik & Astington, 1988; Perner et al., 1987; see also Sabbagh & Callanan, 1998, for similar findings).[13]

But children's explanations and causal utterances do more than reveal what they already know; they allow children to interpret information within a social context. Based on research on school-age children's problem-solving (e.g., Chi et al., 1994), several studies suggest that encouraging preschoolers to generate explanations of data they observe helps them make causal inferences (e.g., Bonawitz & Lombrozo, 2012; Macris & Sobel, 2017; Walker et al., 2017). Indeed, in their meta-analysis of play-based approaches to learning, Alfieri et al. (2011) showed that the clearest benefit to children's learning outcomes was when children generated their own explanations of the situation. Explanations also allow young children to learn novel pieces of information, which in turn allow them to use systematic strategies in requesting explanations from others.

For example, children begin to ask "why" questions around 30 months old (Hood & Bloom, 1979) and 3- to 5-year-olds use such questions spontaneously (Callanan & Oakes, 1992). Similar work found that preschoolers used "why" questions—and questions more generally—in systematic ways to gain information (Frazier et al., 2009). Children in this study continued to ask the same kind of question when given a nonexplanatory reply, but shifted to a new topic or a different question when they received an explanation.[14] This strategy clearly develops between the ages of 3 and 5; for example, Callanan et al. (1995) showed that children shifted from simply asking single-word "why" questions to elicit explanations to jointly constructing such explanations with their parents.

Various studies also examine children's likelihood to seek or generate explanations in order to address gaps in their knowledge—an expression of their curiosity (Jirout & Klahr, 2012). Legare (2012), for example, showed that children are more likely to generate novel explanations when they observe outcomes that are inconsistent with their existing knowledge. Stahl and Feigenson (2015) demonstrated a similar finding in infants' exploratory behavior: Infants explore objects for a longer amount of time when these objects violate their expectations. Exploratory behavior and the search for explanations potentially share information-seeking mechanisms (see e.g., Legare et al., 2017).

These information-seeking mechanisms might be in place before children's linguistic capacities allow them to articulate a "drive" for explanations (Gopnik, 1998), meaning that some of the increased complexity of young children's abilities to generate explanations over the course of the preschool years reflects their more sophisticated language capacities. This is another case in which children's performance on a task is constrained by development in their domain-general abilities. For example, Ruggeri et al. (2017) showed that a group of 5-year-olds were more likely to choose which of two questions would correctly provide them with more information than a group of 3-year-olds. When children were given the opportunity to ask questions that they designed themselves to learn about the cause of an event, there was a shift between what Ruggeri and Lombrozo (2015) called "hypothesis scanning" questions, which reduced the hypothesis space by considering one particular hypothesis out of many possible ones, to "constraint seeking" questions, which eliminated large sets of possibilities. Seven- and 8-year-olds in this study tended to ask many hypothesis scanning questions, but the frequency of these questions decreased as children got older. The younger children could choose which of two questions (designed by another person) provided them with more information, but there is a developmental lag before children can design and ask those questions themselves in this specific context.

Of importance is the trajectory of children's metacognitive development with respect to understanding how best to seek information. For example, children not only use explanations for learning, they also evaluate others' explanations and learn judiciously. Corriveau and Kurkul (2014) found that 5-year-olds preferred to learn from informants who generated noncircular explanations rather than circular ones, while the younger children they investigated did not (although see Mercier et al., 2014, for findings that suggest that even 3-year-olds are sensitive to others' circular explanations). However, the ability to probe further when given a weak explanation in a more challenging context seems to have a longer developmental trajectory, particularly between the ages of 7 and 10 years (e.g., Danovitch et al., 2021; Mills et al., 2019). Attending to unexpected events and inconsistencies often leads children to seek explanations, which in turn may help them to progress along the path toward knowledge acquisition or hypothesis revision, particularly as they develop the understanding that the weaker explanations that result from their own or others' gaps in knowledge may motivate further exploration.

With respect to the theme of this book, there is potentially an important connection between the drive to generate explanations and scientific thinking. Scientific thinking does not involve ignoring or rejecting inconsistent evidence, nor does it involve declaring inconsistencies or mysteries to be beyond the scope of the investigation (Chinn & Brewer, 1993). Engaging in explanation and exploration thus may serve as a critical mechanism within a cultural context for integrating and reconciling discordant or ambiguous information with existing theories. Further, this may reduce engagement in heuristics like confirmation bias, which preserve children's existing causal knowledge and prevent learning. Frankly, this kind of critical thinking and curious seeking after explanations could help adults as well.

# Constructing Science

## Connecting Causal Reasoning to Scientific Thinking in Young Children

## By: Deena Skolnick Weisberg, David M. Sobel

## Citation:

*Constructing Science: Connecting Causal Reasoning to Scientific Thinking in Young Children*
**By:** Deena Skolnick Weisberg, David M. Sobel
**DOI:** 10.7551/mitpress/11939.001.0001
**ISBN (electronic):** 9780262370615
**Publisher:** The MIT Press
**Published:** 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Weisberg, Deena Skolnick, author. | Sobel, David M., author.
Title: Constructing science : connecting causal reasoning to scientific thinking
in young children / Deena Skolnick Weisberg and David M. Sobel.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes
bibliographical references and index.
Identifiers: LCCN 2021045987 | ISBN 9780262044684 (paperback)
Subjects: LCSH: Science—Methodology. | Reasoning in children. | Scientific ability. |
Science—Study and teaching—Psychological aspects. | Constructivism (Education)
Classification: LCC Q175.32.R45 W45 2022 | DDC 501—dc23/eng/20211214
LC record available at https://lccn.loc.gov/2021045987