

8 Infants' Physical Reasoning and the Cognitive Architecture that Supports It

Yi Lin, Maayan Stavans, and Renée Baillargeon

8.1 Introduction

Traditionally, research on early physical reasoning has focused on the simple types of physical events our distant human ancestors routinely observed and produced as they interacted with objects. These types include, for example, occlusion, containment, support, and collision events. Over the first two years of life, infants become increasingly sophisticated at reasoning about these events. How is this sophistication achieved? In this chapter, we describe three successive waves of infancy research that each brought to light critical components of the cognitive architecture that supports early physical reasoning and its development.

8.2 First Wave: Core Knowledge and Information-Processing Capabilities

The study of early physical reasoning began with Piaget (1952, 1954), who was the first researcher to systematically investigate the development of infants' physical knowledge. He examined their responses in various action tasks and concluded that infants initially possess little knowledge about the physical world. For example, after observing that infants under eight or nine months of age (henceforth young infants) do not search for objects they have observed being hidden, Piaget proposed that young infants lack a concept of *object permanence* and do not yet understand that

objects are objective, permanent entities that continue to exist when out of sight. Piaget's conclusion that young infants understand very little about physical events was generally accepted until the 1980s, when researchers became concerned that his exclusive reliance on action tasks (the only ones available to him at the time) might have led him to underestimate infants' physical knowledge.

This concern led investigators to seek alternative methods for exploring young infants' physical reasoning. One method that proved particularly helpful in revealing hitherto unsuspected competencies was the *violation-of-expectation* (VOE) method (Baillargeon et al., 1985). This method takes advantage of infants' natural tendency to look longer at events that violate, as opposed to confirm, their expectations. In recent years, several variations of the VOE method have been developed. For example, researchers have found that infants spend more time exploring objects featured in unexpected as opposed to expected events (Stahl & Feigenson, 2015; Zhang & Wang, 2019) and select unexpected over expected events when allowed to choose what they see next (Jin et al., 2018). All of these VOE methods depend on infants' propensity to use their mental model of the world to predict how events will unfold; when an event does not unfold as expected, infants inspect it to glean information for revising their model, so as to better predict outcomes in the future.

8.2.1 Object Permanence Revisited

Over time, numerous VOE experiments on object permanence revealed that, contrary to what Piaget (1952, 1954) had claimed, even very young infants realize that objects continue to exist when out of sight (Baillargeon, 1993). For example, infants aged two and a half to five months detected a violation when an object was hidden behind a screen that then rotated through the space occupied by the object (Baillargeon, 1987; Baillargeon et al., 1985); when an object moved through an obstacle behind a screen (Baillargeon & DeVos, 1991; Spelke et al., 1992); when an object was hidden in one location and retrieved from a different location (Newcombe et al., 1999; Wilcox et al., 1996); when an object moved behind one screen and reappeared from behind a different screen without appearing in the gap between them (Aguiar & Baillargeon, 1999); when an object was hidden in a container that was then slid forward and to the side to reveal the object standing in the container's initial position (Hespos & Baillargeon, 2001b); and when an object disappeared from behind a screen (Wynn, 1992) or from under a cover (Wang et al., 2005).

These and other similar results provided converging evidence that from a very young age, infants can represent and reason about hidden objects. By the same token, these results also called into question the Piagetian view, prevalent during most of the twentieth century, that infants are limited sensorimotor processors incapable of representation or thought (Piaget, 1952, 1954). As might be expected, fierce controversies ensued as researchers steeped in the Piagetian tradition questioned these new VOE tasks and offered deflationary accounts for their findings. According to many of these accounts, infants looked longer at the unexpected than at the

expected test event in each task because the familiarization events used to introduce the task inadvertently induced a transient and superficial preference for the unexpected event (Bogartz et al., 1997; Cashon & Cohen, 2000; Haith, 1998; Thelen & Smith, 1994). However, empirical tests of these alternative accounts provided little support for them: Even when given a VOE object-permanence task with no familiarization trials, only test trials, young infants still looked significantly longer at the unexpected than at the expected event, suggesting that they did possess a concept of object permanence (Wang et al., 2004).

Today, there is general agreement among developmental researchers that young infants can represent objects that go out of sight. Indeed, researchers often take advantage of this capacity to explore other facets of early cognition, such as infants' ability to track others' beliefs (Hyde et al., 2018; Kampis et al., 2015; Kovács et al., 2010; Southgate & Vernetti, 2014). For example, in a study using functional near-infrared spectroscopy (Hyde et al., 2018), seven-month-olds watched videotaped scenarios in which an agent saw a toy being hidden in one of two containers. Next, the agent either faced away while the toy was transferred to the other container (*false-belief* scenario) or witnessed this transfer (*true-belief* scenario). In each scenario, activation in the temporal-parietal junction (a brain region involved in the tracking of others' beliefs) was measured prior to the agent's search for the toy. Like adults (Hyde et al., 2015), infants showed more activation during the false- than the true-belief scenario, suggesting that they were tracking what information was available to the agent about the location of the hidden toy (see Chapters 12 and 13). This conclusion presumes, of course, that infants could represent the continued existence of the hidden toy.

8.2.2 Further Physical Expectations

The findings from VOE object-permanence tasks did not only demonstrate that young infants expect an object to continue to exist when out of sight: In many cases, due to the specific events shown, the findings provided evidence of additional physical expectations. In particular, they indicated that young infants already understand that an object cannot pass through space occupied by another object (Baillargeon et al., 1985; Spelke et al., 1992), cannot follow a discontinuous path through space (Aguiar & Baillargeon, 1999; Spelke et al., 1995a), and cannot exert a force on another object without contact (Kotovsky & Baillargeon, 2000; Spelke et al., 1995b).

Encouraged by these findings, investigators began exploring other aspects of infants' physical world, adapting the VOE method as needed for the purpose. It soon became clear that while young infants held the expectations listed above for both inert and self-propelled objects (including humans; Baillargeon et al., 1990; Saxe et al., 2006), the same was not true of other expectations, which differed for the two types of objects. For example, young infants detected a violation if an inert object suddenly began to move on its own, if it spontaneously reversed direction after being set into motion, if it failed to move when forcibly hit or pulled, and if it failed to fall when released in midair (Luo et al., 2009; Needham & Baillargeon, 1993; Saxe et al., 2007; Spelke et al., 1995b). Strikingly, all of these expectations differed for self-propelled objects: Young infants did not find it unexpected if a self-propelled object reversed direction on its own, if it failed to move when forcibly hit or pulled, and if it failed to fall when released in midair (Baillargeon et al., 2009b; Leslie & Keeble, 1987; Luo et al., 2009; Spelke et al., 1995b). These and related

results suggested that when a novel object gives sufficient evidence of being self-propelled, young infants endow it with an internal source of energy and understand that it can use this energy to initiate or alter its own motion as well as to resist or exert external forces (Baillargeon et al., 2009b; Gelman, 1990; Leslie, 1994; Luo et al., 2009; Saxe et al., 2007).

8.2.3 Core Knowledge

The results summarized in Sections 8.2.1 and 8.2.2 hinted at remarkably sophisticated physical knowledge in young infants. As such, these results naturally gave rise to the following questions: Where did this knowledge come from? How could we explain its presence in young infants with limited motor skills and scant experience of the world?

An influential proposal, the *core-knowledge hypothesis*, suggested an answer to these questions. This hypothesis holds that infants are born with a skeletal framework of core principles and concepts that guides their reasoning about physical events (Baillargeon, 2008; Baillargeon & Carey, 2012; Carey, 2011; Gelman, 1990; Keil, 1995; Leslie, 1995; Spelke et al., 1992, 1995b; Ullman et al., 2017; Wellman & Gelman, 1992). Descriptions of these principles and concepts differ among researchers, and they have also changed substantially over time as new findings have come to light. Nevertheless, a common assumption is that young infants are capable of sophisticated reasoning about physical events because they are innately prepared by evolution to do so: Their skeletal framework places them in the right ball park, so to speak, to begin reasoning about events in ways that will make possible rapid learning about the physical world.

8.2.3.1 Core Principles

The core principles in the skeletal framework underlying infants' physical reasoning constrain their expectations about the displacements and interactions of objects and other physical entities. To the best of our knowledge, these principles include “persistence,” “inertia,” and “gravity” (these are introduced in quote marks to emphasize that they are only rudimentary versions of the principles used by physicists).

The *persistence* principle states that all other things being equal, objects persist, as they are, in time and space (Baillargeon, 2008; Baillargeon et al., 2009a). This principle has many corollaries, which dictate that an object cannot occupy the same space as another object (solidity) and cannot spontaneously disappear (continuity), break apart (cohesion), fuse with another object (boundedness), or change into a different object (unchangeableness) (Baillargeon, 2008; Baillargeon et al., 2009a; Spelke et al., 1992, 1995b). (Of course, objects can undergo such modifications through causal transformations, but our focus here is on spontaneous, unassisted, physically impossible modifications). The positive findings of the VOE object-permanence tasks reviewed earlier indicate that, from a young age, infants are sensitive to persistence violations (Aguiar & Baillargeon, 1999; Baillargeon, 1987; Hespos & Baillargeon, 2001b; Spelke et al., 1992; Wang et al., 2005; Wilcox et al., 1996; for related findings with non-solid substances, see Anderson et al., 2018; Hespos et al., 2016).

The *inertia* principle states that objects at rest will remain at rest and objects in motion will follow a smooth path without abrupt changes in direction or speed, unless they are acted upon by forces sufficient to alter their rest or motion states (Baillargeon et al., 2009b;

Luo et al., 2009). The evidence reviewed earlier that young infants find it unexpected if an inert object initiates its own motion, spontaneously reverses direction, or remains stationary when forcibly hit or pulled, indicates that, from an early age, infants are sensitive to inertia violations (Baillargeon et al., 2009b; Kosugi & Fujita, 2002; Kotovsky & Baillargeon, 2000; Luo et al., 2009; Saxe et al., 2007; Spelke et al., 1995b).

Finally, the *gravity* principle states that, all other things being equal, objects fall when unsupported (Baillargeon & DeJong, 2017). The evidence reviewed earlier that young infants find it unexpected if an inert object remains suspended in midair indicates that, from an early age, infants are sensitive to gravity violations (Baillargeon et al., 2009b; Luo et al., 2009; Needham & Baillargeon, 1993).

8.2.3.2 Core Concepts

The core concepts in the skeletal framework underlying infants' physical reasoning involve unobservable elements that help explain events' outcomes. Core concepts include “internal energy” and “force.” As we saw in Section 8.2.3.1, when a novel object gives sufficient evidence of being self-propelled (e.g., begins to move on its own), young infants endow it with *internal energy* and recognize that it can use this energy to control its motion and to resist or exert forces (Baillargeon et al., 2009b; Gelman, 1990; Leslie, 1995; Luo et al., 2009; Saxe et al., 2007). When infants see an object hit another object, they represent a *force* – like a directional arrow – being exerted by the first object onto the second one (Kominsky et al., 2017; Kotovsky & Baillargeon, 1994, 2000; Leslie, 1995; Leslie & Keeble, 1987; Mascialzoni et al., 2013). There are no doubt other explanatory concepts that play an important role in infant's physical reasoning. Some of these, like the concept of *cause*, may be

highly abstract and shared with other domains of core knowledge, such as psychological reasoning (Liu et al., 2019).

8.2.3.3 Kinds of Explanations

When watching physical events, infants bring to bear their core knowledge to build explanations for the events and predict how they will unfold. In these explanations, principles and concepts are woven together seamlessly. To illustrate, consider an experiment in which six-month-olds were first introduced to a novel self-propelled box (Luo et al., 2009). Next, the box rested behind a screen that lay flat on the apparatus floor, and infants saw one of two test events. In the *one-screen* event, the screen was lifted and lowered to reveal no box, and then it was lifted and lowered again to reveal the box once more. The *two-screen* event was identical except that a second screen stood upright to the right of the first; when raised, the first screen occluded the left edge of the second screen, making it possible for the box to surreptitiously move behind it.

Infants looked significantly longer if shown the one-screen as opposed to the two-screen event, suggesting that they (1) categorized the box as a self-propelled object, endowed with internal energy, (2) found it unexpected in the one-screen event when the box magically disappeared and reappeared, in violation of the persistence principle, and (3) inferred in the two-screen event that the box used its internal energy to slip behind the second screen when it “disappeared” and to return behind the first screen when it “reappeared.” Control results with an inert box supported this interpretation, as infants then found both events unexpected. Together, these findings nicely illustrate how infants’ core knowledge can support their physical reasoning and help them generate plausible explanations for novel or unfamiliar events.

Of course, the explanations infants build for physical events are typically shallow and lacking in mechanistic detail (Keil, 1995; Wilson & Keil, 2000) (e.g., how did the self-propelled box use its internal energy to move back and forth behind the screens?). As Keil (1995) noted, these are “kinds of explanations” (p. 261), rather than specific, detailed, mechanistic explanations. Nevertheless, infants’ shallow causal understandings are sufficient to support many sophisticated inferences (Aguiar & Baillargeon, 2002; Saxe et al., 2005).

8.2.4 Information-Processing Capabilities

Although infants’ core physical knowledge could explain their success at VOE object-permanence tasks, one important question remained: If infants could represent the continued existence of hidden objects from a very young age, why did they fail manual-search tasks for several months after they learned to reach for objects? The dissociation between the positive findings of VOE object-permanence tasks and the negative findings of manual-search tasks (see also Ahmed & Ruffman, 1998; Daum et al., 2009) remained the focus of heated debate for many years, until a new approach suggested a way of reconciling these divergent findings (Boudreau & Bushnell, 2000; Diamond, 2013; Keen & Berthier, 2004). Proponents of this *processing-load* approach suggested that (1) infants’ information-processing resources are initially limited and improve gradually with age; (2) the processing demands of any action task depend on both the difficulty of the physical reasoning involved and the difficulty of the actions involved; and (3) infants may fail at an action task because the *combined demands* of the task overwhelm their limited resources. From this perspective, the reason why young infants who are able to reach for objects fail at manual-search tasks is

not that they cannot represent a hidden object (they do so in VOE tasks; Baillargeon et al., 1985), and not that they cannot plan and execute means-end actions to retrieve an object (they do so in action tasks with partly visible objects; Shinskey, 2002). Rather, it is that doing *both* of these activities at once (i.e., representing a hidden object *and* planning and executing the actions necessary to retrieve it) overwhelms their limited information-processing capabilities.

The processing-load approach has led researchers to seek action tasks that minimize overall demands when investigating at what age infants first demonstrate specific physical knowledge in their actions. Several VOE findings have now been confirmed using low-demand action tasks, making clear that when task demands are kept at a minimum to avoid taxing infants' limited information-processing capacities, action tasks can reveal the same physical knowledge as VOE tasks (for a review, see Hauf et al., 2012).

8.3 Second Wave: Developments in the Physical-Reasoning System

As investigations continued, it soon became clear that one could not fully account for infants' physical reasoning by considering only their core knowledge and information-processing capabilities. A key difficulty was that when tested with subtle core violations that could not be discerned without attending to the specific properties of objects and their arrangements, infants often failed to detect these violations. When an object passed behind a large screen, for example, infants under three months did not detect a violation if the object failed to appear in a low opening in the screen (Aguiar & Baillargeon, 1999, 2002); infants under three and a half months did not detect a violation if the object failed to appear in a high opening in the screen (Baillargeon &

DeVos, 1991; Luo & Baillargeon, 2005); infants under seven and a half months did not detect a violation if the object surreptitiously changed pattern behind the screen (Wilcox, 1999; Wilcox et al., 2011); and infants under eleven and a half months did not detect a violation if the object surreptitiously changed color behind the screen (Káldy & Leslie, 2003; Wilcox & Chapa, 2004).

These and similar negative results with other events (Baillargeon, 1991; Baillargeon et al., 1992; Kotovsky & Baillargeon, 1998; Newcombe et al., 1999) led to two broad realizations. First, because infants apply their core knowledge not to events in the world but to *mental representations* of these events, they can detect subtle core violations involving specific properties of objects and their arrangements only if they include the relevant information in their event representations (e.g., when an object passes behind a screen, infants can detect a surreptitious change to the color of the object only if they include color information in their representation of the event). Second, the evidence that infants initially detect few violations and come to detect more and more with age indicates that their event representations are at first very sparse and become gradually richer and more detailed. Spurred by these realizations, researchers began to investigate how event representations develop over time (for detailed reviews, see Baillargeon et al., 2009a, 2011).

8.3.1 Event Representations

Research on early event representations has yielded a large body of evidence that we summarize in three sets of findings.

8.3.1.1 Event Categories and Vectors

As infants observe and produce physical events, they form distinct *event categories*,

such as occlusion, containment, support, collision, covering, tube, and burying events (Casasola, 2008; Hespos & Baillargeon, 2001a, 2006; Kotovsky & Baillargeon, 2000; Mou & Luo, 2017; Newcombe et al., 1999; Wang & Baillargeon, 2006; Wang et al., 2005). Each event category represents a type of causal interaction between objects. To predict how events from a category will unfold over time, infants have to learn about multiple facets of the events; we refer to these facets as *vectors*. When an object is lowered into a container, for example, vectors for this containment event include: whether the object will fit into the opening of the container, whether it will protrude above the rim of the container, whether it will remain partly visible through the sidewalls of the container, and, when retrieved from a container large enough to contain multiple objects, whether it is the same individual object or a different object.

8.3.1.2 Rules and Causally Relevant Features

For each vector of an event category, infants acquire *rules* that identify *features* (i.e., properties of objects and their arrangements) that are causally relevant for predicting outcomes (Baillargeon et al., 1992; Hespos & Baillargeon, 2001a; Kotovsky & Baillargeon, 1998; Wang & Baillargeon, 2006; Wang et al., 2003, 2005, 2016; Wilcox, 1999). Once infants have identified a feature as relevant to an event category, from that point on they routinely include information about the feature when representing events from the category.

For some vectors, the rule needed to predict outcomes is fairly straightforward and is acquired without much difficulty. For example, infants as young as four and a half months of age realize that the width of an object relative to that of a container's opening determines whether the object can be lowered into the container (Goldman & Wang, 2019;

Wang et al., 2004). For other vectors, however, the rule needed to predict outcomes is more complex or multi-faceted, and infants acquire a series of rules that gradually approximate the correct rule, with each new rule revising or elaborating the one(s) before it. In the case of support events involving inert objects, for example, one key vector is whether an object will remain stable or fall when released in contact with another object (henceforth *base*). Initially, young infants have no particular expectation about the outcomes of these events; between about four and a half and thirteen and a half months of age, however, they identify a series of rules that help them predict these outcomes more and more accurately.

Thus, by about four and a half to five and a half months, infants acquire a *type-of-contact* rule: An object remains stable if released on top of a base, but not if released against or under it (Baillargeon, 1995; Hespos & Baillargeon, 2008; Merced-Nieves et al., 2020). By about six and a half months, they acquire a *proportion-of-contact* rule: An object on a base remains stable as long as 50 percent or more of its bottom surface is supported (Baillargeon et al., 1992; Hespos & Baillargeon, 2008; Luo et al., 2009). By about eight months, they acquire a *position-of-contact* rule: An object on a base can remain stable with less than 50 percent support as long as the middle of its bottom surface is supported (Dan et al., 2000; Huettel & Needham, 2000; Wang et al., 2016). Finally, by about thirteen months, infants acquire a *proportional-distribution* rule: When released with one end on a base, an object remains stable as long as the proportion of the entire object (not just its bottom surface) on the base is greater than that off the base (Baillargeon & DeJong, 2017). This last rule allows infants to correctly predict the outcomes of support events involving asymmetrical as well as symmetrical objects (e.g., an L-shaped box

released with the rightmost 50 percent of its bottom surface on a base will fall, because the proportion of the entire box that is off the base is greater than that on the base).

8.3.1.3 Errors of Omission and Commission

By knowing what rules infants have acquired, researchers can predict the types of errors infants will produce in VOE tasks (Baillargeon & DeJong, 2017; Luo & Baillargeon, 2005; Wang et al., 2016; Zhang & Wang, 2019). *Errors of omission* occur when infants see a physically impossible event (e.g., in the laboratory) and view it as expected because it happens to be consistent with their faulty rule. An example of an error of omission is a five-month-old who fails to detect a violation when an object remains stable with only the leftmost 15 percent of its bottom surface supported, because this event is consistent with her type-of-contact rule (Baillargeon et al., 1992; Hespos & Baillargeon, 2008). In contrast, *errors of commission* occur when infants see a physically possible event (e.g., in the laboratory or in daily life) and find it unexpected because it happens to contradict their faulty rule. An example of an error of commission is a seven-month-old who detects a violation when an object remains stable with only the middle 33 percent of its bottom surface supported, because this event is inconsistent with her proportion-of-contact rule (Wang et al., 2016; Zhang & Wang, 2019).

8.3.2 Explanation-based Learning

How do infants acquire and revise their physical rules? There is growing evidence that *explanation-based learning* (EBL) is one of the key processes that enable them to do so (Baillargeon & DeJong, 2017; Wang, 2019; Wang & Baillargeon, 2008a; Wang & Kohne, 2007).

8.3.2.1 The EBL Process

EBL has three main steps (Baillargeon & DeJong, 2017), and the first is *triggering*. When infants encounter outcomes they cannot explain based on their current knowledge, the EBL process is triggered. In situations where no existing rule applies, infants may notice unexplained variation in events' outcomes (e.g., infants who have not yet acquired the first support rule, type-of-contact, may notice that objects released in contact with a base sometimes remain stable and sometimes fall). In situations where an existing rule does apply, infants may notice that, while some outcomes support the rule, others contradict it (e.g., infants who have acquired the type-of-contact rule may notice that objects released on top of a base sometimes remain stable, as predicted, but sometimes fall). Either way, exposure to the unexplained outcomes triggers EBL.

The second step in the EBL process is *explanation construction and generalization*. Infants first search for a potential feature whose values consistently map onto the different outcomes they have observed (e.g., when the feature has value x , one outcome is observed; when the feature has value y , a different outcome is observed). If they discover such a feature (infants' statistical-learning or regularity-detection processes must often play a key role in this discovery; Kirkham et al., 2002; Saffran & Kirkham, 2018; Saffran et al., 1996; Wang, 2019), they bring to bear their physical knowledge (i.e., core knowledge and acquired rules) to generate a plausible explanation for how the feature contributed to the observed outcomes. If they can construct such an explanation, they then generalize it, resulting in a candidate rule that incorporates only the relevant feature specified in the explanation.

The final step in the EBL process is *empirical confirmation*. Once a rule has been

hypothesized, it must be evaluated against further empirical evidence, which will serve to either confirm or reject it. If the candidate rule proves accurate in predicting outcomes for a few additional exemplars, it is adopted, becomes part of infants' physical knowledge, and, from then on, helps guide prediction and action.

The EBL process makes clear why infants generally do not acquire rules based on spurious or accidental regularities in their environments: For a regularity to be adopted as a rule for an event category, it must be plausibly (even if shallowly) explained by infants' physical knowledge. Finally, the EBL process also makes clear why infants may require only a few exemplars to acquire a new rule. Because EBL combines both *analytic evidence* (i.e., the explanation that is constructed and generalized into a candidate rule) and *empirical evidence*, it makes possible highly efficient learning.

8.3.2.2 Teaching Experiments

The EBL process does not only make clear how infants acquire and revise their physical rules: It also suggests how infants might be "taught" a rule they have not yet acquired via exposure to EBL-designed observations (Baillargeon & DeJong, 2017; Wang, 2019; Wang & Baillargeon, 2008a; Wang & Kohne, 2007). One teaching experiment, for example, sought to teach eleven-month-olds the support rule of proportional distribution, which is typically not acquired until about thirteen months (Baillargeon & DeJong, 2017).

Infants first received three pairs of teaching trials. In each pair, an experimenter's gloved hand placed the right half of an asymmetrical box's bottom surface on a base and then released the box. Consistent with physical laws, the box fell when released with its smaller end on the base (*small-on* event), but it

remained stable when released in the reverse orientation, with its larger end on the base (*large-on* event). Each teaching pair involved a different asymmetrical box (e.g., a box shaped like a letter B on its back, a right-triangle box, and a staircase-shaped box). Following the teaching trials, infants saw two static test displays in which half of an L-shaped box's bottom surface lay on a base. In the *unexpected* display, the box's smaller end was supported; in the *expected* display, the box's larger end was supported.

Infants detected the violation in the unexpected test display, suggesting that they had acquired the proportional-distribution rule during the three pairs of teaching trials. How did these trials facilitate EBL? First, each small-on event contradicted infants' proportion-of-contact rule (i.e., the box fell even though half of its bottom surface rested on the base), and these unexplained outcomes triggered the EBL process. Second, because in each teaching pair the small-on and large-on events differed only in the box's orientation, infants could rapidly zero in on this information in their quest for an explanation. By bringing to bear their physical knowledge, infants could reason that (1) since an inert object falls when unsupported (in accordance with the gravity principle) but remains stable when released on a base because the base passively blocks its fall (in accordance with the solidity principle), then (2) it was plausible that in each teaching trial the base could block the fall of the asymmetrical box when half or more of the entire box was on the base, but not when half or more of the entire box was off the base – the larger unsupported portion of the box then caused it to tip off the base and topple to the apparatus floor. Armed with this explanation, infants could then hypothesize a proportional-distribution rule: An object released with one end resting on a base will remain stable as long as the proportion of the entire object that is on

the base is greater than that off the base. Third, infants could confirm this hypothesized rule because across the teaching trials they saw three different asymmetrical boxes all behave in accordance with the rule (e.g., infants might have used the first two boxes to generate the rule, and the third box to confirm it).

Additional experiments indicated that infants no longer learned the proportional-distribution rule (i.e., failed to detect the violation in the unexpected test display) if the teaching trials were modified to disrupt one or more of the EBL steps. Thus, infants did not acquire the rule (1) when shown only teaching events consistent with their proportion-of-contact rule, so that the EBL process was not triggered (e.g., infants saw only large-on events); (2) when shown *reverse* teaching events for which they could construct no plausible explanation (e.g., in each teaching pair, the box now remained stable in the small-on event and fell in the large-on event); and (3) when shown too few distinct exemplars to generate and empirically confirm the rule (e.g., the three teaching pairs involved only two asymmetrical boxes, with one box appearing in both the first and third pairs). Finally, infants also failed to acquire the rule when shown teaching events that could in principle support EBL but made the search for an explanation harder (e.g., salient irrelevant differences were added to the teaching events, making it difficult for infants to rapidly zero in on the box's orientation as a critical feature).

8.3.3 Décalages

When introducing the second wave of research on early physical reasoning, we noted that when young infants are tested with subtle core violations that can be discerned only by attending to the specific properties of objects and their arrangements, they often fail to detect these violations. The research reviewed

in Section 8.3.2 helped explain these failures by showing that (1) infants typically succeed at reasoning about a feature in an event only if they have identified the feature as causally relevant for the event category involved; (2) after a feature is identified as relevant (with EBL playing an important role in this identification process), infants routinely encode information about the feature when representing events from the category; and (3) once information about a feature is included in an event's representation, it is interpreted by infants' physical knowledge, allowing them to detect subtle core violations involving the feature. These include *interaction* violations (i.e., objects interact in ways that are not physically possible given their properties) and *change* violations (i.e., objects spontaneously undergo changes that are not physically possible). To illustrate, consider the occlusion feature height, which is identified at about three and a half months. Infants who have acquired this feature detect an interaction violation if a tall object that is passing behind a screen of the same height fails to appear in a high opening in the screen (Baillargeon & DeVos, 1991; Baillargeon & Graber, 1987), or if a tall object becomes almost fully hidden when lowered behind a short occluder (Hespos & Baillargeon, 2001a; Mou & Luo, 2017). Infants also detect a change violation if an object is either much taller or much shorter after being briefly occluded (Goldman & Wang, 2019; Wang & Baillargeon, 2006). Thus, for any feature identified as causally relevant to an event category, there is *broad generalization within the category*: The feature is encoded for any event from the category (e.g., for any occlusion event), and it allows the detection of many types of violations involving the feature.

In contrast to infants' pervasive and flexible use of identified features *within* each event category, there is no evidence that

infants transfer identified features *between* event categories: Features from one category are not passed on to other categories, even when equally relevant. This means that when infants happen to identify a feature at different ages in different event categories, striking *décalages* (to use a Piagetian term) or lags can be observed in their responses to similar events from the different categories. For example, five-to-six-month-olds detect a change violation if an object surreptitiously changes shape when behind an occluder (Káldy & Leslie, 2005; Wilcox, 1999) or when inside a container (Wang & Onishi, 2017), but not if it changes shape when buried in sand (Newcombe et al., 1999). Similarly, infants as young as three and a half months can detect interaction and change violations involving height in occlusion events, as we just saw, but such violations are not detected until much later in other event categories: at about seven and a half months in containment events, twelve months in covering events, and fourteen months in tube events (Hespos & Baillargeon, 2001a; Wang & Baillargeon, 2006; Wang et al., 2005).

Décalages between event categories have also been observed in action tasks. For example, six-month-olds correctly searched for a tall frog behind a tall as opposed to a short occluder – but they searched randomly when the occluders were replaced with a tall and a short container (the occluders were identical to the fronts of the containers) (Hespos & Baillargeon, 2006). Similarly, after being “taught” to attend to the feature height in covering events, nine-month-olds correctly searched for a tall toy under a tall as opposed to a short cover, thereby showing immediate generalization of the feature to novel events from the category – but they searched randomly when the covers were replaced with a tall and a short tube (the tubes were identical

to the covers without their tops; Wang & Kohne, 2007).

8.3.4 Object-File and Physical-Reasoning Systems

8.3.4.1 Multiple Representations

The growing evidence of marked *décalages* between event categories indicated that (1) identified features are not transferred across categories and (2) weeks or months can separate the identification of the same feature in different categories. Additional evidence indicated that when infants failed to include information about an unidentified feature in an event representation, it did not necessarily mean that they had not registered the feature at all (i.e., that their brains had not encoded it in any way). Strikingly, infants who failed to detect a violation involving a feature could sometimes be shown, using other tasks, to have registered the feature (Wang & Goldman, 2016; Wang & Mitroff, 2009).

In one experiment (Wang & Goldman, 2016), for example, twelve-month-olds saw an experimenter’s hand lower a tall cover (cover condition) or a tall tube (tube condition) over a short block. Next, the hand lifted the cover or tube to reveal either the same block as before (no-change event) or a much taller block (change event). Consistent with prior findings that the feature height is identified at about twelve months in covering events but only at about fourteen months in tube events (Wang & Baillargeon, 2006; Wang et al., 2005), infants in the cover condition detected the change to the block’s height, whereas those in the tube condition did not. However, infants did detect this change in a modified-tube condition in which they were briefly *turned away* from the apparatus while the tube was lowered over the block and lifted back again.

These results suggested that two distinct cognitive systems were involved in infants' responses. One system formed detailed representations of the test objects, including their heights. When infants witnessed no causal interaction in the test trial, as in the modified-tube condition, this first system guided infants' responses, leading to enhanced attention to the novel block in the change event (i.e., infants produced a novelty response). However, when infants did witness a causal interaction, as in the cover and tube conditions, a second system took over, built a specialized representation of the event, and used it to predict how the event would unfold. Because this second system had already identified height as a causally relevant feature for covering events, but not for tube events, it behaved differently in the two conditions. In the cover condition, the second system retrieved the height information from the first system and included it in the event's representation; when interpreted by infants' physical knowledge, this information allowed them to detect the persistence violation in the change event. In the tube condition, in contrast, the second system did not retrieve the height information from the first system, causing infants to fail to detect the persistence violation in the change event.

The notion that infants might form multiple representations of objects and hold information in one representation that they fail to use in another might be puzzling at first. However, this notion echoes extensive findings from the adult literature on change blindness. In particular, these findings show that (1) adults often fail to detect salient changes to attended objects that go briefly out of view, in both laboratory and real-world settings (Rensink et al., 1997; Simons & Levin, 1998; Simons et al., 2002) and (2) adults may overlook a featural change to an object, even though the information necessary to detect this change has been encoded, is maintained, and can be

accessed experimentally via photographic line-ups, probing questions, or more implicit measures (Angelone et al., 2003; Hollingworth et al., 2001; Mitroff et al., 2004).

8.3.4.2 The Two-System Model

In recent years, several multi-system models have been proposed (Baillargeon et al., 2011, 2012; Wang & Baillargeon, 2008b), building on prior research in the adult and infant literature (Huttenlocher et al., 1991, 2002; Kahneman et al., 1992; Leslie et al., 1998; Pylyshyn, 1989, 2007; Rips et al., 2006).

In the most recent of these models (Lin et al., 2021; Stavans et al., 2019), the two cognitive systems that contribute to early physical reasoning are the *object-file* (OF) system (Gordon & Irwin, 1996; Kahneman et al., 1992) and the *physical-reasoning* (PR) system (Baillargeon et al., 2011; Wang & Baillargeon, 2008b). The two systems serve different functions and have at least partly distinct neural substrates (Fischer et al., 2016; Grill-Spector et al., 2001). When infants see objects (e.g., in a picture book, a static scene, or an event), the OF system builds a temporary representation of the “where” and “what” information about each object, drawing on incoming perception as well as on stored knowledge, and it updates this information as needed. If the objects are involved in a causal interaction, the PR system also becomes engaged. It builds a specialized representation of the interaction that contains a subset of the information in the objects' files, and it uses this representation, together with its physical knowledge (i.e., core knowledge and acquired rules), to predict how the interaction will unfold.

To illustrate how the two systems operate, imagine that two objects, A and B, are resting on an apparatus floor. As infants view this static scene, the OF system builds a temporary representation of each object, which includes

spatiotemporal (“where”) information as well as *identity* (“what”) information. Each type of information comprises broad categorical descriptors as well as more fine-grained featural descriptors. Now, imagine that A and B become involved in a causal interaction. This engages the PR system, which then builds a specialized representation of this event, in two steps. In the first, the PR system uses the OF system’s spatiotemporal and identity categorical descriptors to categorize the event. For example, if B is identified as a container, the event is categorized as an *occlusion* event if A moves behind B, as a *collision* event if A hits B, and as a *containment* event if A is placed inside B. Once the PR system has categorized the event, it assigns event-specific roles to the objects (e.g., if A moves behind B, then A is assigned the role of occludee and B that of occluder). In the second step, the PR system accesses the list of features it has identified as causally relevant for predicting outcomes in the event category selected, and it then taps the OF system for information about these – and only these – features. The retrieved information (e.g., about the relative heights and widths of the occludee and occluder and about the shape and pattern of the occludee) is then added to the event’s representation. Finally, the PR system brings to bear its physical knowledge to interpret the categorical and featural information in the event’s representation and guide infants’ responses.

The two-system model helps explain all of the findings summarized to this point in the chapter. First, it explains why very young infants can already detect some core violations (the spatiotemporal and identity categorical information included in the PR system’s event representations is sufficient, when interpreted by the core knowledge, to allow the detection of these violations). Second, it explains why infants become able to detect more subtle core violations with development (once the PR

system has identified a feature as causally relevant for an event category, it routinely taps the OF system for information about this feature when representing events from the category, making possible the detection of violations involving the feature). Finally, it explains why infants who detect a subtle core violation in an event from one category may fail to do so in a similar event from a different category, even though the featural information necessary to detect this violation is available in the OF system (the PR system only taps the OF system for information about identified features).

8.3.4.3 Carryover Effects

The two-system model also suggested new directions for research. In carryover experiments, researchers asked the following question: When infants see a sequence of two events that involve the same objects but belong to different event categories (e.g., a cover is placed first in front of and then over an object), does the PR system discard the featural information it requested from the OF system for the first event’s representation, or does it *carry over* this information to the second event’s representation, to save time and effort? If there is such a carryover, it should have positive consequences whenever (1) the second event depicts a subtle core violation that involves a particular feature and (2) this feature has been identified in the first event’s category but not the second event’s category. This is because information about the feature will be included in the first event’s representation and, when carried over to the second event’s representation, will allow infants to detect the violation in the event.

Several experiments have now demonstrated carryover effects in infants’ detection of interaction violations (Baillargeon et al., 2009a; Wang, 2011; Wang & Baillargeon, 2005) and

change violations (Wang & Onishi, 2017). For example, eight-month-olds detected an interaction violation when a short cover was first slid in front of a tall object, returned to its initial position, and then lifted and lowered over the object until it became fully hidden (Wang & Baillargeon, 2005). The height information that was carried over from the occlusion to the covering event enabled infants to detect a persistence violation that is typically not detected until about twelve months (Wang et al., 2005). Similarly, four-and-a-half-month-olds detected an interaction violation when a tall object was first slid in front of a short container, returned to its initial position, and then lifted and lowered into the container until it became almost fully hidden (Wang, 2011). Here again, the carryover of height information from the occlusion to the containment event enabled infants to detect a persistence violation that is typically not detected until about seven and a half months (Hespos & Baillargeon, 2001a). Additional results indicated that this effect was eliminated when a twenty-second delay was inserted between the two events (infants either were turned away from the apparatus or saw the tall object being moved back and forth next to the container) (Wang, 2011). There are thus temporal limits to carryover effects: As the interval between the two events increases, the PR system becomes more likely to discard the featural information from the first event and tap the OF system for the featural information it has identified as relevant to the second event.

8.3.4.4 Priming Effects

In carryover experiments, by definition, infants reason about a feature in a first event and then carry over the feature to a second, target event; in priming experiments, researchers asked whether infants would still succeed if instead of a first event they saw a *static array*

that highlighted the feature but gave them no opportunity to reason about its effect on an event's outcome (Baillargeon et al., 2009a; Lin et al., 2021; Wang, 2019). The rationale was that if (1) the static array rendered information about the feature more salient in the OF system and (2) this caused the OF system to spontaneously pass on this information when the PR system represented the target event, then (3) the PR system should be able to use the information to guide infants' responses to the event.

One priming experiment with twelve-month-olds, for example, focused on the feature color in containment events (Lin et al., 2021). Infants were assigned to a baseline or a priming condition. In the baseline condition, infants saw test events in which a brightly colored (e.g., orange) doll was lowered into a container too small to hold more than one doll. When lifted again, the doll was either the same as before (no-change event) or a different color (e.g., purple; change event). Infants failed to detect the change to the doll's color, indicating that their PR system had not yet identified color as a containment feature. The priming condition was identical to the baseline condition except that prior to the test events, infants saw a static array of four dolls differing only in color (orange, purple, yellow, pink). Infants now detected the change to the doll's color, suggesting that (1) the static array highlighted the information about the dolls' colors in the OF system; (2) this color information was spontaneously passed on to the PR system; and (3) when interpreted by the PR system's physical knowledge, this information allowed infants to detect the persistence violation in the change event. Additional priming experiments (Lin et al., 2021) indicated that following exposure to static arrays of objects differing only in height, eight-to-ten-month-olds succeeded at reasoning about height in tube events (recall that this feature is typically identified at about

fourteen months; Wang et al., 2005): They detected a persistence violation if a tall object was much shorter after being briefly lowered into a tall tube, and they searched for a tall object in a tall as opposed to a short tube. Similarly, following exposure to static arrays of asymmetrical objects (Lin et al., 2021), seven-month-olds succeeded at reasoning about proportional distribution in support events (recall that this feature is typically identified at about thirteen months; Baillargeon & DeJong, 2017): They detected a gravity violation if an L-shaped box remained stable with its larger end unsupported.

Like the carryover experiments discussed in Section 8.3.4.3, these priming experiments demonstrate that when information about a feature is fortuitously included in an event representation, however it comes to be so, infants then bring to bear their physical knowledge to interpret this information, allowing them to succeed at VOE and action tasks involving the feature – even six months before they typically identify the feature!

8.3.5 Information-Processing Capabilities

In describing the two-system model, we focused mainly on simple situations involving a single event. What happens when two or more events occur side by side in quick succession, with each event involving different objects? As might be expected, infants' limited information-processing capacities begin to restrict how well they represent and reason about each event.

To illustrate, consider experiments in which six-month-olds saw occlusion events involving two identical screens, screen-1 and screen-2, and two objects that differed in shape, A and B (e.g., a red disk and a red triangle; Applin & Kibbe, 2019; Káldy & Leslie, 2005; Kibbe & Leslie, 2011, 2019). In the test events, A was

hidden behind screen-1, and then B was hidden behind screen-2. Infants detected a change violation if screen-2 was lifted to reveal A, but not if screen-1 was lifted to reveal B, suggesting that they had difficulty keeping track of the identity of the first-hidden object. Infants detected other violations involving this object, however, indicating that they did represent some information about it. For example, infants detected a violation if screen-1 was lifted to reveal no object at all (Kibbe & Leslie, 2011). Furthermore, when A and B differed in their ontological categories in that one was human-like and one was not (e.g., a doll's head and a ball), infants detected a violation if screen-1 was lifted to reveal B instead of A (Kibbe & Leslie, 2019).

Together, these results suggest that when A was hidden behind screen-1, the PR system began building an event representation, using the spatiotemporal and identity categorical descriptors provided by the OF system. However, the PR system was unable to adequately deal with the featural information for this event representation: When B was hidden behind screen-2, the PR system had to begin building an event representation for this second event, and it did not have sufficient information-processing resources to simultaneously (1) build this second event representation and (2) complete and/or maintain the first one. Thus, while the second event representation included both categorical and featural information about B, the first event representation included only categorical information about A (i.e., the PR system could not retrieve the featural information about A from the OF system and/or could not bind or maintain this information). As a result, infants detected a violation if screen-1 was lifted to reveal no object or an object from a different ontological category than A, but not an object that simply differed in shape from A.

These results support the two-system model, and also make clear how limitations in infants'

information-processing capacities can hamper the operation of the PR system. Further results with older infants make the same point (Káldy & Leslie, 2003; Kibbe & Leslie, 2013): At nine months, infants detected a violation if either screen-1 or screen-2 was lifted to reveal the wrong object, but they failed with a more challenging situation involving three screens and three objects (for related findings with adults, see Strickland & Scholl, 2015).

8.4 Third Wave: Developments in the Object-File System

Although much of infants' physical reasoning could be explained by considering developments in their PR system and their information-processing capacities, as we saw in Sections 8.3.4 and 8.3.5, there remained a critical difficulty. Experiments with occlusion or containment events indicated that although infants detected interaction and change violations involving features they had identified for these categories, they nevertheless failed to detect *individuation* violations involving these same features. An individuation violation is a type of persistence violation in which fewer objects are revealed at the end of an event than were presented during the event, as though one or more objects had magically disappeared. Xu, Carey, and their colleagues were the first to report this baffling failure (Van de Walle et al., 2000; Xu & Carey, 1996; Xu et al., 2004), and similar results were subsequently obtained in a wide range of individuation tasks (for reviews, see Baillargeon et al., 2012; Stavans et al., 2019). In one task, for example, an experimenter's hand brought out two objects in alternation on either side of a large screen; for present purposes, let us assume that the objects belonged to the same basic-level category and differed only in their featural properties (e.g., a large red ball with blue dots and a small yellow ball with white stripes).

After several repetitions of this occlusion event, the screen was removed to reveal only one of the objects. Infants aged twelve months and younger failed to detect this violation, suggesting that they did not clearly expect to see two objects when the screen was removed, and hence that they were unable to individuate the objects in the occlusion event (i.e., to determine how many individual objects were present; Lin & Baillargeon, 2018; Lin et al., 2019; Stavans et al., 2019; Xu et al., 2004). To researchers familiar with the literature on early physical reasoning, these results were puzzling. By twelve months, most infants have identified size, pattern, and color as occlusion features, and the persistence principle dictates that these features cannot undergo spontaneous changes (e.g., a ball cannot spontaneously change size, pattern, or color). Why, then, did infants not infer that two objects were present behind the screen?

Controversy over the causes of infants' individuation failures persisted for many years, because most accounts could explain only a subset of available findings. However, it eventually became clear that by extending the two-system model described in the last section to consider not only developments in the PR system but also *developments in the OF system*, one could reconcile the findings of individuation tasks with those of other physical-reasoning tasks.

8.4.1 Individuation in the OF and PR Systems

Stavans et al. (2019) proposed that infants' individuation failures stem from catastrophic (as opposed to reconcilable) disagreements between the OF and PR systems. Here, we describe four key assumptions of their account.

8.4.1.1 Different Bases for Individuation

When a physical event comes to an end, infants successfully track the objects involved

past the endpoint of the event as long as the OF and PR systems agree on how many objects are present. To individuate objects, each system uses somewhat different information: The OF system uses *categorical* information (i.e., the categorical descriptors in the objects' files), whereas the PR system uses *both categorical and featural* information (i.e., the categorical and featural information in the event's representation).

Why does the OF system not use the featural as well as the categorical information in its files to individuate objects in physical events? After all, in object-recognition tasks, the OF system does use the featural information at its disposal to detect changes (Oakes et al., 2006; Wang & Goldman, 2016; Wang & Mitroff, 2009). Why are things different in physical-reasoning tasks? According to the two-system model, the main reason has to do with infants' limited information-processing resources. During an event, the OF and PR systems are both engaged, but the PR system has priority: It must operate rapidly, online, to make sense of the unfolding event and predict its outcome. While this is happening (and taking up a sizeable portion of infants' information-processing resources), the OF system can do little more than track the objects in the event by checking their categorical descriptors. Thus, if two objects that come into view in alternation have different descriptors, the OF system infers that two objects are present; if they have the same descriptors, however, it infers that a single object is present and updates its featural properties.

In challenging situations that tax their information-processing resources, adults, too, tend to focus on objects' categorical descriptors; provided these are maintained across views, they fail to notice changes to objects' features (unless, of course, these changes are perceptually highly salient). For example, in an experiment inspired by the work of Xu

and Carey (1996), Simons and Levin (1998) embedded an occlusion event in a novel social interaction on a college campus. An actor who carried a map and was dressed as a construction worker (e.g., a young White man wearing a plain hard hat, black shirt, and white pants) approached individual students and asked for directions. In each case, the interaction between the actor and the student was interrupted by two confederates who passed between them, carrying a door. While occluded, the actor surreptitiously switched positions with one of the confederates, another young White man who also carried a map and was dressed as a construction worker, though in different clothing (e.g., a hard hat with a logo, a tool belt, a light blue shirt, and tan pants). Most students failed to notice the change to the actor, suggesting that they selectively compared the pre- and post-change actors' categorical descriptors (e.g., young, White, male construction worker requesting directions) and mistakenly inferred that a single actor was present because these descriptors remained constant across views.

8.4.1.2 Developments

In each system, significant developments occur with age in how objects are represented. In the OF system, more fine-grained spatiotemporal and identity categorical descriptors come to be used in objects' files. With respect to identity descriptors, for example, infants under twelve months typically do not spontaneously encode an isolated object's basic-level category, such as ball, toy duck, or cup (Pauen, 2002; Xu & Carey, 1996). However, they do encode more abstract or ontological descriptors, such as whether the object is human-like or non-human (Bonatti et al., 2002; Kibbe & Leslie, 2019), whether it is animate or inanimate (Setoh et al., 2013; Surian & Caldi, 2010), and whether it is a container (open at the

top), a cover (open at the bottom), a tube (open at both ends), or a closed object (Mou & Luo, 2017; Wang et al., 2005). By their first birthday, infants begin to spontaneously encode objects' basic-level categories (Cacchione et al., 2013; Xu & Carey, 1996).

Turning to the PR system, two types of changes occur with development. First, because categorical descriptors are passed on to the PR system for its event representations, the OF system's more fine-grained categorical descriptors will also be available to the PR system. Second, as we saw in the last section of the chapter, the PR system includes more and more detailed information about objects' properties and arrangements as it learns, event category by event category, what features are causally relevant for predicting outcomes.

8.4.1.3 Catastrophic Disagreements

It follows from the preceding discussions that under some conditions, the OF and PR systems will disagree on how many objects are present. In particular, consider an occlusion event in which two objects (e.g., two different balls, as before) emerge in alternation on either side of a screen. Because the OF system can establish a continuous spatiotemporal trace between successive emergences, it assigns similar spatiotemporal categorical descriptors to each object. Disagreements between the two systems occur when (1) the OF system also assigns the same identity categorical descriptors to each object (e.g., non-human, inanimate, ball) and hence infers that a single object is present behind the screen, while (2) the PR system encodes distinct featural information about each object (e.g., large, red, blue dots; small, yellow, white stripes) and hence infers that two objects are present behind the screen. Stavans et al. (2019) refer to such disagreements as *catastrophic*. Before the screen is lowered, the OF signals

that a single object is present behind the screen, whereas the PR system signals that two objects are present. At this point, the OF realizes that its object file is corrupted: It does not cleanly refer to a single object in the world but instead contains a tangled mix of information that pertains to two separate objects. The OF system then discards its corrupted file, leading infants to have no expectation at all about how many objects will be revealed when the screen is lowered.

8.4.1.4 Reconcilable Disagreements

In catastrophic disagreements, the OF system represents a single object in a hiding location, whereas the PR system represents two objects in the same location. The systems cannot recover from such disagreements, leading to individuation failures. However, they *can* recover from other types of disagreements. In particular, consider an occlusion event in which two objects (e.g., two different balls) emerge in alternation on either side of a screen. Finally, one of the objects stops in plain view next to the screen, which is then lowered. The OF system will assign the same categorical descriptors to each object, will infer that a single object is present, and will conclude that this object is now resting in view, leaving no object behind the screen. In contrast, the PR system will encode distinct featural information about each object, will infer that two objects are present, and will conclude that while one is resting in view, the other remains hidden behind the screen. In this situation, the OF and PR systems have no disagreement about the object in view; their disagreement is only about whether objects remain behind the screen. The OF system assumes that there are none, whereas the PR system signals that one object still remains. Because the OF system currently has no object file linked to the area behind the screen, it can respond to the PR

system's signal by adding one object file for that area. The OF and PR systems are then in agreement, leading infants to expect one object when the screen is lowered.

8.4.2 Predictions

The two-system model can explain a wide range of individuation findings, and here we focus on three predictions in particular (see Stavans et al., 2019, for additional predictions).

8.4.2.1 Prediction 1: Categorical Descriptors

According to the two-system model, young infants should succeed at an individuation task whenever the OF system assigns different identity categorical descriptors to the objects. In line with this analysis, twelve-month-olds (who use basic-level as well as ontological descriptors) succeed when tested with objects from different basic-level categories (e.g., a toy duck and a ball; Van de Walle et al., 2000; Xu & Carey, 1996), and nine-to-ten-month-olds (who use ontological but not basic-level descriptors) succeed when tested with objects from different ontological categories (e.g., a doll and a ball; Bonatti et al., 2002; Decarli et al., 2020; Surian & Caldi, 2010).

Although infants under twelve months typically fail to individuate objects that differ only in their basic-level categories, they succeed if induced to encode these categories via experimental manipulations (Futó et al., 2010; Stavans & Baillargeon, 2018; Xu, 2002). For example, in a *language-based* manipulation (Xu, 2002), nine-month-olds heard a distinct label (e.g., "Look, a duck!" or "Look, a ball!") as each object came into view during the occlusion event. Following this manipulation, infants detected a violation when the screen was lowered to reveal only one of the objects. Similarly, in a *function-based* manipulation

(Stavans & Baillargeon, 2018), four-month-olds first watched functional demonstrations for two different tools, one at a time (e.g., in one trial, a masher was used to compress sponges, and in another trial, tongs were used to pick them up). The two tools were then brought out in alternation from behind a screen, and infants detected a violation when the screen was lowered to reveal only one of the tools.

Finally, although twelve-month-olds typically fail to individuate objects they encode as merely featurally distinct, they succeed if they first see the objects play different roles in other events (recall that the OF system includes both incoming and stored information in objects' files). Thus, in a *role-based* manipulation (Lin et al., 2019), thirteen-month-olds first saw two blocks that differed only in pattern and color play different event roles in relation to a toy (e.g., in one trial, one block supported the toy, and in another trial, the other block was supported by the toy). The two blocks were then brought out in alternation from behind a screen, and infants detected a violation when the screen was lowered to reveal only one of the blocks.

8.4.2.2 Prediction 2: Catastrophic Disagreements

According to the two-system model, when the OF system signals that one object is present in a hiding location but the PR system signals that two objects are present, the OF system discards its corrupted object file, leading infants to hold no expectation at all about how many objects are present. Consistent with this analysis, after seeing two objects they encoded as merely featurally distinct emerge in alternation from behind a screen, eleven-month-olds detected no violation when the screen was lowered to reveal *no object at all* (Stavans et al., 2019). Similarly, after seeing

two objects they encoded as merely featurally distinct being lifted, one at a time, from a large container, nine-month-olds detected no violation if the container *remained silent when shaken*, as though empty (Stavans et al., 2019).

8.4.2.3 Prediction 3: Reconcilable Disagreements

According to the two-system model, when the OF system signals that no object remains in a hiding location but the PR system signals that one object remains, the OF system then adds an object file for that location. In line with this analysis, positive findings have been obtained with five-to-eleven-month-olds in a variety of *remainder* tasks (Lin & Baillargeon, 2019; McCurry et al., 2009; Stavans et al., 2019; Wilcox & Baillargeon, 1998; Wilcox & Schweinle, 2002; Xu & Baker, 2005). In one experiment, for example, eleven-month-olds first saw two objects they encoded as merely featurally distinct emerge in alternation from behind a screen (Lin & Baillargeon, 2019). Next, one of the objects paused in plain view, and the screen was lowered to reveal an empty area – only the paused object was visible next to the screen. Infants detected the violation in this event, suggesting that the OF system successfully added an object file when the PR system signaled that an object remained behind the screen. Results were positive even if *three* featurally distinct objects emerged in alternation in the occlusion event. As long as the OF system assumed that no object remained behind the screen at the end of the event, (1) the PR system could signal that two objects still remained behind the screen and (2) the OF system could add two object files pointing to that area, leading infants to detect a violation when the screen was lowered to reveal only one object.

8.5 Conclusions

In this chapter, we sought to offer a historical overview that made clear how three successive waves of research on early physical reasoning over the past four decades not only led to a deeper understanding of the development of this ability but also helped shed light on the cognitive architecture needed to support it. In the first wave, demonstrations of early sensitivity to persistence, inertia, and gravity led to the suggestion that a *skeletal framework of core principles* guides infants' reasoning about physical events. In the second wave came the realization that infants often fail to adhere to these principles in their predictions and actions due to *limitations in their PR system*: Early in development, event representations in the PR system tend to be very sparse and often lack the featural information necessary for infants to respond appropriately. Over time, however, event representations become richer and more detailed as the PR system forms event categories and identifies causally relevant features for each category. Finally, the third wave made clear that even when featural information is included in an event representation, infants may still fail to adhere to the core principle of persistence in their predictions and actions due to *limitations in their OF system*: While the PR system uses both the categorical and featural information in its event representations to individuate objects in an event, the OF system uses only the categorical information in its object files for this purpose. When the two systems disagree as an event comes to an end (e.g., the PR system signals that two objects are present behind a screen, whereas the OF system signals that a single object is present), infants fail to track the objects past the endpoint of the event.

The two-system model provides a coherent, integrative framework for the findings from these three waves of research. As we saw, this

model not only reconciles findings that initially appeared puzzling or even contradictory but also makes novel predictions that are being tested in different laboratories. Nevertheless, the two-system model still leaves many questions unanswered. For example, at what age does infants' OF system begin to use featural information to individuate objects in physical events, and what are the mechanisms responsible for this development? Could the OF system be induced to use featural information for this purpose at an earlier age, via experimental manipulations? The next wave of research should bring answers to these questions.

References

- Aguiar, A., & Baillargeon, R. (1999). 2.5-month-old infants' reasoning about when objects should and should not be occluded. *Cognitive Psychology*, *39*, 116–157.
- Aguiar, A., & Baillargeon, R. (2002). Developments in young infants' reasoning about occluded objects. *Cognitive Psychology*, *45*, 267–336.
- Ahmed, A., & Ruffman, T. (1998). Why do infants make A not B errors in a search task, yet show memory for the location of hidden objects in a nonsearch task? *Developmental Psychology*, *34*, 441–453.
- Anderson, E. M., Hespos, S. J., & Rips, L. J. (2018). Five-month-old infants have expectations for the accumulation of nonsolid substances. *Cognition*, *175*, 1–10.
- Angelone, B. L., Levin, D. T., & Simons, D. J. (2003). The relationship between change detection and recognition of centrally attended objects in motion pictures. *Perception*, *32*, 947–962.
- Applin, J. B., & Kibbe, M. M. (2019). Six-month-old infants predict agents' goal-directed actions on occluded objects. *Infancy*, *24*, 392–410.
- Baillargeon, R. (1987). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology*, *23*, 655–664.
- Baillargeon, R. (1991). Reasoning about the height and location of a hidden object in 4.5- and 6.5-month-old infants. *Cognition*, *38*, 13–42.
- Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. E. Granrud (ed.), *Visual Perception and Cognition in Infancy* (pp. 265–315). Hillsdale, NJ: Erlbaum.
- Baillargeon, R. (1995). A model of physical reasoning in infancy. In C. Rovee-Collier, & L. P. Lipsitt (eds.), *Advances in Infancy Research* (Vol. 9, pp. 305–371). Norwood, NJ: Ablex.
- Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science*, *3*, 2–13.
- Baillargeon, R., & Carey, S. (2012). Core cognition and beyond: The acquisition of physical and numerical knowledge. In S. Pauen (ed.), *Early Childhood Development and Later Outcome* (pp. 33–65). Cambridge: Cambridge University Press.
- Baillargeon, R., & DeJong, G. F. (2017). Explanation-based learning in infancy. *Psychonomic Bulletin & Review*, *24*, 1511–1526.
- Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: Further evidence. *Child Development*, *62*, 1227–1246.
- Baillargeon, R., & Graber, M. (1987). Where's the rabbit? 5.5-month-old infants' representation of the height of a hidden object. *Cognitive Development*, *2*, 375–392.
- Baillargeon, R., Graber, M., DeVos, J., & Black, J. (1990). Why do young infants fail to search for hidden objects? *Cognition*, *36*, 225–284.
- Baillargeon, R., Li, J., Gertner, Y., & Wu, D. (2011). How do infants reason about physical events? In U. Goswami (ed.), *The Wiley-Blackwell Handbook of Childhood Cognitive Development*, *11* (2nd ed., pp. 11–48). Oxford: Blackwell.
- Baillargeon, R., Li, J., Ng, W., & Yuan, S. (2009a). An account of infants' physical reasoning. In A. Woodward, & A. Needham (eds.), *Learning and the Infant Mind* (pp. 66–116). New York: Oxford University Press.
- Baillargeon, R., Needham, A., & DeVos, J. (1992). The development of young infants' intuitions about support. *Early Development and Parenting*, *1*, 69–78.

- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, *20*, 191–208.
- Baillargeon, R., Stavans, M., Wu, D., Gertner, Y., Setoh, P., Kittredge, A. K., & Bernard, A. (2012). Object individuation and physical reasoning in infancy: An integrative account. *Language Learning and Development*, *8*, 4–46.
- Baillargeon, R., Wu, D., Yuan, S., Li, J., & Luo, Y. (2009b). Young infants' expectations about self-propelled objects. In B. Hood, & L. Santos (eds.), *The Origins of Object Knowledge* (pp. 285–352). Oxford: Oxford University Press.
- Bogartz, R. S., Shinsky, J. L., & Speaker, C. J. (1997). Interpreting infant looking: The event set \times event set design. *Developmental Psychology*, *33*, 408–422.
- Bonatti, L., Frot, E., Zangl, R., & Mehler, J. (2002). The human first hypothesis: Identification of conspecifics and individuation of objects in the young infant. *Cognitive Psychology*, *44*, 388–426.
- Boudreau, J. P., & Bushnell, E. W. (2000). Spilling thoughts: Configuring attentional resources in infants' goal-directed actions. *Infant Behavior and Development*, *23*, 543–566.
- Cacchione, T., Schaub, S., & Rakoczy, H. (2013). Fourteen-month-old infants infer the continuous identity of objects on the basis of nonvisible causal properties. *Developmental Psychology*, *49*, 1325–1329.
- Carey, S. (2011). *The Origin of Concepts*. New York: Oxford University Press.
- Casasola, M. (2008). The development of infants' spatial categories. *Current Directions in Psychological Science*, *17*, 21–25.
- Cashon, C. H., & Cohen, L. B. (2000). Eight-month-old infants' perception of possible and impossible events. *Infancy*, *1*, 429–446.
- Dan, N., Omori, T., & Tomiyasu, Y. (2000). Development of infants' intuitions about support relations: Sensitivity to stability. *Developmental Science*, *3*, 171–180.
- Daum, M. M., Prinz, W., & Aschersleben, G. (2009). Means-end behavior in young infants: The interplay of action perception and action production. *Infancy*, *14*, 613–640.
- Decarli, G., Franchin, L., Piazza, M., & Surian, L. (2020). Infants' use of motion cues in object individuation processes. *Journal of Experimental Child Psychology*, *197*, 104868.
- Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, *64*, 135–168.
- Fischer, J., Mikhael, J. G., Tenenbaum, J. B., & Kanwisher, N. (2016). Functional neuroanatomy of intuitive physical inference. *Proceedings of the National Academy of Sciences (USA)*, *113*, E5072–E5081.
- Futó, J., Téglás, E., Csibra, G., & Gergely, G. (2010). Communicative function demonstration induces kind-based artifact representation in preverbal infants. *Cognition*, *117*, 1–8.
- Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cognitive Science*, *14*, 79–106.
- Goldman, E. J., & Wang, S. H. (2019). Comparison facilitates the use of height information by 5-month-olds in containment events. *Developmental Psychology*, *55*, 2475–2482.
- Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception & Psychophysics*, *58*, 1260–1277.
- Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*, 1409–1422.
- Haith, M. M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior and Development*, *21*, 167–179.
- Hauf, P., Paulus, M., & Baillargeon, R. (2012). Infants use compression information to infer objects' weights: Examining cognition, exploration, and prospective action in a preferential-reaching task. *Child Development*, *83*, 1978–1995.
- Hespos, S. J., & Baillargeon, R. (2001a). Infants' knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science*, *12*, 141–147.

- Hespos, S. J., & Baillargeon, R. (2001b). Reasoning about containment events in very young infants. *Cognition*, *78*, 207–245.
- Hespos, S. J., & Baillargeon, R. (2006). Décalage in infants' knowledge about occlusion and containment events: Converging evidence from action tasks. *Cognition*, *99*, B31–B41.
- Hespos, S. J., & Baillargeon, R. (2008). Young infants' actions reveal their developing knowledge of support variables: Converging evidence for violation-of-expectation findings. *Cognition*, *107*, 304–316.
- Hespos, S. J., Ferry, A. L., Anderson, E. M., Hollenbeck, E. N., & Rips, L. J. (2016). Five-month-old infants have general knowledge of how nonsolid substances behave and interact. *Psychological Science*, *27*, 244–256.
- Hollingworth, A., Williams, C. C., & Henderson, J. M. (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin & Review*, *8*, 761–768.
- Huettel, S. A., & Needham, A. (2000). Effects of balance relations between objects on infant's object segregation. *Developmental Science*, *3*, 415–427.
- Huttenlocher, J., Duffy, S., & Levine, S. (2002). Infants and toddlers discriminate amount: Are they measuring? *Psychological Science*, *13*, 244–249.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376.
- Hyde, D. C., Aparicio Betancourt, M., & Simon, C. E. (2015). Human temporal-parietal junction spontaneously tracks others' beliefs: A functional near-infrared spectroscopy study. *Human Brain Mapping*, *36*, 4831–4846.
- Hyde, D. C., Simon, C. E., Ting, F., & Nikolaeva, J. I. (2018). Functional organization of the temporal-parietal junction for theory of mind in preverbal infants: A near-infrared spectroscopy study. *Journal of Neuroscience*, *38*, 4264–4274.
- Jin, K. S., Houston, J. L., Baillargeon, R., Groh, A. M., & Roisman, G. I. (2018). Young infants expect an unfamiliar adult to comfort a crying baby: Evidence from a standard violation-of-expectation task and a novel infant-triggered-video task. *Cognitive Psychology*, *102*, 1–20.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, *24*, 175–219.
- Káldy, Z., & Leslie, A. M. (2003). Identification of objects in 9-month-old infants: Integrating “what” and “where” information. *Developmental Science*, *6*, 360–373.
- Káldy, Z., & Leslie, A. M. (2005). A memory span of one? Object identification in 6.5-month-old infants. *Cognition*, *97*, 153–177.
- Kampis, D., Parise, E., Csibra, G., & Kovács, Á. M. (2015). Neural signatures for sustaining object representations attributed to others in preverbal human infants. *Proceedings of the Royal Society B: Biological Sciences*, *282*, 20151683.
- Keen, R. E., & Berthier, N. E. (2004). Continuities and discontinuities in infants' representation of objects and events. In R. V. Kail (ed.), *Advances in Child Development and Behavior* (Vol. 32, pp. 243–279). San Diego, CA: Elsevier Academic Press.
- Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 234–262). Oxford: Clarendon Press.
- Kibbe, M. M., & Leslie, A. M. (2011). What do infants remember when they forget? Location and identity in 6-month-olds' memory for objects. *Psychological Science*, *22*, 1500–1505.
- Kibbe, M. M., & Leslie, A. M. (2013). What's the object of object working memory in infancy? Unraveling “what” and “how many.” *Cognitive Psychology*, *66*, 380–404.
- Kibbe, M. M., & Leslie, A. M. (2019). Conceptually rich, perceptually sparse: Object representations in 6-month-old infants' working memory. *Psychological Science*, *30*, 362–375.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy:

- Evidence for a domain general learning mechanism. *Cognition*, 83, B35–B42.
- Kominsky, J. F., Strickland, B., Wertz, A. E., Elsner, C., Wynn, K., & Keil, F. C. (2017). Categories and constraints in causal perception. *Psychological Science*, 28, 1649–1662.
- Kosugi, D., & Fujita, K. (2002). How do 8-month-old infants recognize causality in object motion and that in human action? *Japanese Psychological Research*, 44, 66–78.
- Kotovskiy, L., & Baillargeon, R. (1994). Calibration-based reasoning about collision events in 11-month-old infants. *Cognition*, 51, 107–129.
- Kotovskiy, L., & Baillargeon, R. (1998). The development of calibration-based reasoning about collision events in young infants. *Cognition*, 67, 311–351.
- Kotovskiy, L., & Baillargeon, R. (2000). Reasoning about collision events involving inert objects in 7.5-month-old infants. *Developmental Science*, 3, 344–359.
- Kovács, Á. M., Téglás, E., & Endress, A. D. (2010). The social sense: Susceptibility to others' beliefs in human infants and adults. *Science*, 330, 1830–1834.
- Leslie, A. M. (1994). ToMM, ToBy, and Agency: Core architecture and domain specificity. In L. A. Hirschfeld, & S. A. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 119–148). New York: Cambridge University Press.
- Leslie, A. M. (1995). A theory of agency. In D. Sperber, D. Premack, & A. J. Premack (eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 121–149). Oxford: Clarendon Press.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25, 265–288.
- Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: developing “what” and “where” systems. *Trends in Cognitive Sciences*, 2, 10–18.
- Lin, Y., & Baillargeon, R. (2018). Infants individuate objects with distinct prior event roles. *Paper presented at the Biennial International Congress of Infant Studies*, June 2018, Philadelphia, PA.
- Lin, Y., & Baillargeon, R. (2019). Testing a new two-system model of early individuation. *Paper presented at the Biennial Meeting of the Cognitive Development Society*, September 2019, Louisville, KY.
- Lin, Y., Li, J., Gertner, Y., Ng, W., Fisher, C. L., & Baillargeon, R. (2021). How do the object-file and physical-reasoning systems interact? Evidence from priming effects with object arrays or novel labels. *Cognitive Psychology*, 125, 101368.
- Lin, Y., Stavans, M., & Baillargeon, R. (2019). Infants can use many types of categories to individuate objects. *Paper presented at the Biennial Meeting of the Society for Research in Child Development*, March 2019, Baltimore, MD.
- Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proceedings of the National Academy of Sciences (USA)*, 116, 17747–17752.
- Luo, Y., & Baillargeon, R. (2005). When the ordinary seems unexpected: Evidence for incremental physical knowledge in young infants. *Cognition*, 95, 297–328.
- Luo, Y., Kaufman, L., & Baillargeon, R. (2009). Young infants' reasoning about physical events involving inert and self-propelled objects. *Cognitive Psychology*, 58, 441–486.
- Mascalzoni, E., Regolin, L., Vallortigara, G., & Simion, F. (2013). The cradle of causal reasoning: Newborns' preference for physical causality. *Developmental Science*, 16, 327–335.
- McCurry, S., Wilcox, T., & Woods, R. (2009). Beyond the search barrier: A new task for assessing object individuation in young infants. *Infant Behavior and Development*, 32, 429–436.
- Merced-Nieves, F. M., Aguiar, A., Dzwilewski, K. L. C., Musaad, S., Korrick, S. A., & Schantz, S. L. (2020). Association of prenatal maternal perceived stress with a sexually dimorphic measure of cognition in 4.5-month-old infants. *Neurotoxicology and Teratology*, 77, 106850.
- Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, 66, 1268–1281.

- Mou, Y., & Luo, Y. (2017). Is it a container? Young infants' understanding of containment events. *Infancy, 22*, 256–270.
- Needham, A., & Baillargeon, R. (1993). Intuitions about support in 4.5-month-old infants. *Cognition, 47*, 121–148.
- Newcombe, N., Huttenlocher, J., & Learmonth, A. (1999). Infants' coding of location in continuous space. *Infant Behavior and Development, 22*, 483–510.
- Oakes, L. M., Ross-Sheehy, S., & Luck, S. J. (2006). Rapid development of feature binding in visual short-term memory. *Psychological Science, 17*, 781–787.
- Pauen, S. (2002). The global-to-basic level shift in infants' categorical thinking: First evidence from a longitudinal study. *International Journal of Behavioral Development, 26*, 492–499.
- Piaget, J. (1952). *The Origins of Intelligence in Children*. New York: International Universities Press.
- Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic Books.
- Pylyshyn, Z. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition, 32*, 65–97.
- Pylyshyn, Z. W. (2007). *Things and Places: How the Mind Connects with the World*. Cambridge, MA: MIT Press.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8*, 368–373.
- Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review, 113*, 1–30.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928.
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology, 69*, 181–203.
- Saxe, R., Tenenbaum, J., & Carey, S. (2005). Secret agents: 10- and 12-month-old infants' inferences about hidden causes. *Psychological Science, 16*, 995–1001.
- Saxe, R., Tzelnic, T., & Carey, S. (2006). Five-month-old infants know humans are solid, like inanimate objects. *Cognition, 101*, B1–B8.
- Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental Psychology, 43*, 149–158.
- Setoh, P., Wu, D., Baillargeon, R., & Gelman, R. (2013). Young infants have biological expectations about animals. *Proceedings of the National Academy of Sciences (USA), 110*, 15937–15942.
- Shinsky, J. L. (2002). Infants' object search: Effects of variable object visibility under constant means-end demands. *Journal of Cognition and Development, 3*, 119–142.
- Simons, D. J., Chabris, C. F., Schnur, T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition, 11*, 78–97.
- Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review, 5*, 644–649.
- Southgate, V., & Vernetti, A. (2014). Belief-based action prediction in preverbal infants. *Cognition, 130*, 1–10.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review, 99*, 605–632.
- Spelke, E. S., Kestenbaum, R., Simons, D. J., & Wein, D. (1995a). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology, 13*, 113–142.
- Spelke, E. S., Phillips, A., & Woodward, A. L. (1995b). Infants' knowledge of object motion and human action. In D. Sperber, D. Premack, & A. J. Premack (eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 44–78). Oxford: Clarendon Press.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science, 348*, 91–94.
- Stavans, M., & Baillargeon, R. (2018). Four-month-old infants individuate and track simple tools

- following functional demonstrations. *Developmental Science*, 21, e12500.
- Stavans, M., Lin, Y., Wu, D., & Baillargeon, R. (2019). Catastrophic individuation failures in infancy: A new model and predictions. *Psychological Review*, 126, 196–225.
- Strickland, B., & Scholl, B. J. (2015). Visual perception involves event-type representations: The case of containment versus occlusion. *Journal of Experimental Psychology: General*, 144, 570–580.
- Surian, L., & Caldi, S. (2010). Infants' individuation of agents and inert objects. *Developmental Science*, 13, 143–150.
- Thelen, E., & Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Perception and Action*. Cambridge, MA: MIT Press.
- Ullman, T. D., Spelke, E., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21, 649–665.
- Van de Walle, G. A., Carey, S., & Prevor, M. (2000). Bases for object individuation in infancy: Evidence from manual search. *Journal of Cognition and Development*, 1, 249–280.
- Wang, S. (2011). Priming 4.5-month-old infants to use height information by enhancing retrieval. *Developmental Psychology*, 47, 26–38.
- Wang, S. (2019). Regularity detection and explanation-based learning jointly support learning about physical events in early infancy. *Cognitive Psychology*, 113, 101219.
- Wang, S., & Baillargeon, R. (2005). Inducing infants to detect a physical violation in a single trial. *Psychological Science*, 16, 542–549.
- Wang, S., & Baillargeon, R. (2006). Infants' physical knowledge affects their change detection. *Developmental Science*, 9, 173–181.
- Wang, S., & Baillargeon, R. (2008a). Can infants be “taught” to attend to a new physical variable in an event category? The case of height in covering events. *Cognitive Psychology*, 56, 284–326.
- Wang, S., & Baillargeon, R. (2008b). Detecting impossible changes in infancy: A three-system account. *Trends in Cognitive Sciences*, 12, 17–23.
- Wang, S., Baillargeon, R., & Brueckner, L. (2004). Young infants' reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition*, 93, 167–198.
- Wang, S., Baillargeon, R., & Paterson, S. (2005). Detecting continuity violations in infancy: A new account and new evidence from covering and tube events. *Cognition*, 95, 129–173.
- Wang, S., & Goldman, E. J. (2016). Infants actively construct and update their representations of physical events: Evidence from change detection by 12-month-olds. *Child Development Research*, article 3102481.
- Wang, S., Kaufman, L., & Baillargeon, R. (2003). Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events. *Infant Behavior and Development*, 26, 529–568.
- Wang, S., & Kohne, L. (2007). Visual experience enhances infants' use of task-relevant information in an action task. *Developmental Psychology*, 43, 1513–1522.
- Wang, S., & Mitroff, S. R. (2009). Preserved visual representations despite change blindness in infants. *Developmental Science*, 12, 681–687.
- Wang, S., & Onishi, K. H. (2017). Enhancing young infants' representations of physical events through improved retrieval (not encoding) of information. *Journal of Cognition and Development*, 18, 289–308.
- Wang, S., Zhang, Y., & Baillargeon, R. (2016). Young infants view physically possible support events as unexpected: New evidence for rule learning. *Cognition*, 157, 100–105.
- Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, 43, 337–375.
- Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition*, 72, 125–166.
- Wilcox, T., & Baillargeon, R. (1998). Object individuation in infancy: The use of featural information in reasoning about occlusion events. *Cognitive Psychology*, 37, 97–155.

- Wilcox, T., & Chapa, C. (2004). Priming infants to attend to color and pattern information in an individuation task. *Cognition*, *90*, 265–302.
- Wilcox, T., Nadel, L., & Rosser, R. (1996). Location memory in healthy preterm and full-term infants. *Infant Behavior and Development*, *19*, 309–323.
- Wilcox, T., & Schweinle, A. (2002). Object individuation and event mapping: Developmental changes in infants' use of featural information. *Developmental Science*, *5*, 132–150.
- Wilcox, T., Smith, T., & Woods, R. (2011). Priming infants to use pattern information in an object individuation task: The role of comparison. *Developmental Psychology*, *47*, 886.
- Wilson, R. A., & Keil, F. C. (2000). The shadows and shallows of explanation. In F. C. Keil, & R. A. Wilson (eds.), *Explanation and Cognition* (pp. 87–114). Cambridge, MA: MIT Press.
- Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, *358*, 749–750.
- Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, *85*, 223–250.
- Xu, F., & Baker, A. (2005). Object individuation in 10-month-old infants using a simplified manual search method. *Journal of Cognition and Development*, *6*, 307–323.
- Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, *30*, 111–153.
- Xu, F., Carey, S., & Quint, N. (2004). The emergence of kind-based object individuation in infancy. *Cognitive Psychology*, *49*, 155–190.
- Zhang, Y., & Wang, S. (2019). Violation to infant faulty knowledge induces object exploration by 7.5-month-olds in support events. *Paper presented at the Biennial Meeting of the Cognitive Development Society*, September 2019, Louisville, KY.