Core Cognition and Beyond:

The Acquisition of Physical and Numerical Knowledge

Renée Baillargeon

University of Illinois

and

Susan Carey

Harvard University

Only human beings create the conceptual repertoire that underlies the lexicons of natural languages and that subserves the institutions of science, mathematics, government, religion, and art. Accounting for this astonishing feat is one of the foremost challenges in the cognitive and psychological sciences. As a matter of logic, any account must have the following structure: a specification of the initial representational repertoire that serves as the input to subsequent learning, a characterization of the differences between the initial and adult repertoires, and a characterization of the learning mechanisms that make possible conceptual development.

In Part I (written by Renée Baillargeon), we first consider the question of the initial repertoire by presenting a case study of the system of representations that underlies reasoning about the physical world. We illustrate the properties of what has been called "core cognition" in the domain of physics, and also characterize how learning underlies changes in this system of representation within infancy. In Part II (written by Susan Carey), we turn to a case study of the acquisition of concepts of number. We characterize the core cognition systems with numerical content, outline the discontinuities between the initial representations and later ones, and describe the bootstrapping process that underlies the change.

Part I: The Acquisition of Physical Knowledge

Over the past two decades, substantial progress has been made in understanding how infants reason and learn about physical events. This research has led to the development of a new account of the development of infants' physical reasoning (for recent reviews, see Baillargeon, Li, Gertner, & Wu, 2010; Baillargeon, Li, Ng, & Yuan, 2009a; Baillargeon, Wu, Gertner, Setoh, Kittredge, & Stavans, in press). In what follows, we briefly describe this account.

Like several other researchers, we assume that infants are born equipped with a *physical-reasoning (PR) system*—an abstract computational system that provides them with a skeletal causal framework for making sense of the displacements and interactions of objects and other physical entities (e.g., Gelman, 1990; Leslie, 1995; Premack & Premack, 2003; Spelke, Breinlinger, Macomber, & Jacobson, 1992). Before we describe how the PR system operates, three general comments are in order.

First, it should be understood that the PR system operates without conscious awareness. Infants are not aware of the causal framework they use when reasoning about physical events, any more than young children are aware of the grammar of their language as they begin to understand and produce sentences.

Second, the PR system allows infants to reason and learn about the simple everyday physical events that were familiar to our distant evolutionary ancestors: for example, occlusion events (i.e. events in which an object moves or is placed behind another object, or occluder), containment events (i.e. events in which an object is placed inside a container), and collision events (i.e. events in which an object hits another object). In infancy, the PR system has relatively little to say about events that involve complex artifacts whose causal mechanisms are opaque to most adults—artifacts such as phones, computers, radios, and televisions. Although infants may learn to operate some of these artifacts, pedagogical processes may be necessary to support these acquisitions (e.g., Csibra & Gergerly, 2009; Futó, Téglás, Csibra, & Gergely, 2010).

Finally, although our focus here is on the PR system, it should be kept in mind that this system is only a part of the complex cognitive architecture that underlies infants' responses to objects and events. Following other researchers, we have argued that, in addition to the PR system, at least two other systems are involved in infants' responses: an *object-tracking* (OT) and an *object-representation* (OR) system (e.g., Baillargeon et al., 2010; Wang & Baillargeon, 2008b; Wang & Mitroff, 2009). To illustrate the distinction between the three systems, consider a simple static display involving a block and a can standing apart on an apparatus floor. As infants begin to attend to the objects, the OT system assigns an *index* to each object; each index functions as an attentional pointer that "sticks" to its object, enabling infants to keep track of it when it moves—to know where it is without having to search for it (e.g., Leslie, Xu, Tremoulet, & Scholl, 1998; Pylyshyn, 1989, 1994; Scholl & Leslie, 1999). As soon as the OT system assigns indexes to the block and can, the OR system begins to build a temporary *file* for each object, listing both individual (e.g., color) and relational (e.g. relative height) features (e.g., Huttenlocher, Duffy, & Levine, 2002; Kahneman, Treisman, & Gibbs, 1992; Needham, 2001; Rose, Gottfried, Melloy-Carminar, & Bridger, 1982). If an experimenter then places the block inside the can, the PR system also becomes involved: the objects are now engaged in a physical interaction, and the PR system's main purpose is that of interpreting and predicting the outcomes of such interactions (e.g., will the block fit through the can's opening? will the block protrude above the can, or be visible through its sidewalls? if a block is later removed from the can, is it the same block as was seen before or a different one?). In the first year of life, object representations in the PR system often contain only a small subset of the information included in the object files of the OR system; as we explain more fully below, event representations in the PR system are initially very sparse (no doubt to facilitate learning) and become richer as infants gradually learn what information is useful for predicting and interpreting outcomes (for a fuller discussion of the links between the OR and PR systems, see Baillargeon et al., 2010).

Core Concepts and Core Principles

When infants watch a physical event, the PR system builds a specialized *physical representation* of the event. Any information included in this physical representation becomes subject to the PR system's causal framework, which encompasses both core concepts and core principles (e.g., Baillargeon et al., 2009a; Carey & Spelke, 1984; Gelman, 1990; Leslie, 1995; Premack & Premack, 2003; Spelke et al., 1992).

Core concepts invoke unobservable elements that help explain events' outcomes. Examples of core concepts include 'force' and 'internal energy' (these are listed in quote marks to emphasize that they are only primitive versions of the concepts used by scientists). When infants see an object hit another object, the PR system represents a *force*—like a directional arrow—being exerted by the first object onto the second one (e.g., Leslie, 1995; Leslie & Keeble, 1987). Furthermore, when infants see a novel object begin to move or change direction on its own, they categorize it as a self-propelled object and endow it with an *internal source of energy*; infants recognize that a self-propelled object can use its internal energy directly to control its own motion and indirectly—through the application of force—to control the motion of other objects (e.g., Baillargeon, Wu, Yuan, Li, & Luo, 2009b; Gelman, Durgin, & Kaufman, 1995; Leslie, 1995; Luo, Kaufman, & Baillargeon, 2009; Pauen & Träuble, 2009; Premack, 1990; Saxe, Tenenbaum, & Carey, 2005; Woodward, Phillips, & Spelke, 1993).

Core principles constrain infants' expectations about objects' displacements and interactions. Examples of core principles include 'persistence', 'inertia', and 'gravity' (again, these are listed in quote marks to emphasize that they are only primitive versions of the principles used by scientists). The principle of *persistence* states that, all other things being equal, objects persist, as they are, in time and space (e.g., Baillargeon, 2008; Baillargeon et al., 2009a). The persistence principle has many corollaries (e.g., Baillargeon, 2008; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995b), which specify that an object cannot occupy the same space as another object (solidity) and cannot spontaneously appear or disappear (continuity), break apart (cohesion), fuse with another object (boundedness), and change size, shape, pattern, or color (unchangeableness).

The principle of *inertia* states that, all other things being equal, an object in motion will follow a smooth path, without abrupt changes in direction or speed (e.g., Kochukhova & Gredebäck, 2007; Luo et al., 2009). Finally, the principle of *gravity* states that, all other things being equal, an object will fall when released in midair (e.g., Needham & Baillargeon, 1993; Premack & Premack, 2003).

Structural Information

When infants watch a physical event, the PR system begins by representing the *structural information* about the event. This structural information includes both *spatiotemporal* and *categorical* information. The spatiotemporal information describes how the objects are arranged and how this arrangement changes over time as the event unfolds (e.g., Kestenbaum, Termine, & Spelke, 1987; Needham & Ormsbee, 2003; Slater, 1995; Yonas & Granrud, 1984). The categorical information specifies what kinds of objects are involved in the event, by providing categorical descriptors for each object. Early descriptors include (1) abstract *ontological* descriptors, such as whether the objects in the event are human or non-human, agentive or non-agentive, and inert or self-propelled (e.g., Bonatti, Frot, & Mehler, 2005; Bonatti, Frot, Zangl, & Mehler, 2002; Csibra, 2008; Johnson, Shimizu, & Ok, 2007; Luo et al., 2009; Saxe, Tzlenic, & Carey, 2007; Surian & Caldi, 2010; Träuble & Pauen, 2011b), and (2) primitive *functional* descriptors, such as whether the objects in the event are closed, open at the top to form containers, open at the bottom to form covers, or open at both ends to form tubes (e.g., Hespos & Baillargeon, 2001b; Wang & Baillargeon, 2006; Wang, Baillargeon, & Paterson, 2005; Wang & Kohne, 2007; see also Träuble & Pauen, 2007, 2011a).

Both the spatiotemporal and the categorical information about an event help specify how many objects are involved in the event. For example, if two identical objects stand apart on a table and a screen is lifted to hide them, the spatiotemporal information will specify that two objects are present behind the screen (e.g., Aguiar & Baillargeon, 1999, 2002). Similarly, if a human disappears behind a large screen and what reappears is a non-human object, the categorical information will specify that two distinct objects are involved in the event, one human and one non-human (e.g., Bonatti et al., 2002, 2005).

With experience, the PR system begins to identify distinct event categories—kinds of causal interactions in which objects play specific roles. In addition to occlusion, containment, and collision events (which were described earlier), early event categories include covering events (i.e. events which a cover is placed over an object), tube events (i.e. events in which an object is placed inside a tube), and arrested-motion events (i.e. events in which an object is brought short against an extended surface such as a wall or floor). As a rule, events that do not involve causal interactions and have no physical consequences (e.g., events in which an object is simply placed next to another object) are not identified as event categories. When infants watch an event from a known category, the PR system uses the structural information available in the physical representation to *categorize* the event (e.g., Casasola, Cohen, & Chiarello, 2003; Hespos & Baillargeon, 2006; McDonough, Choi, & Mandler, 2003; Quinn, 2007; Wang & Baillargeon, 2006; Wilcox & Chapa, 2002) and to assign specific *roles* to the objects in the event (e.g., Leslie & Keeble, 1987; Onishi, 2011). As an example, consider a simple event involving two identical blocks, block-A and block-B. If block-A is used to hit block-B, the event is categorized as a collision event, with block-A as the 'hitter' and block-B as the 'hittee'. If block-B is lowered behind block-A, the event is categorized as an occlusion event, with block-A as the 'occluder' and block-B as the 'occludee'. After watching one of these events repeatedly, infants look reliably longer if the two objects change roles (e.g., if block-A becomes the hittee in the collision event).

The structural information about an event thus captures its essence: it specifies how many objects are involved (e.g., two objects), what kinds of objects they are (e.g., non-human, non-agentive, inert, closed objects), what kind of causal interaction the objects are engaged in (e.g., a collision event), and what role each object plays in this interaction (e.g., object-A is the hitter, object-B is the hittee).

Over time, at least two changes take place in the structural information infants encode about events. One change concerns the *categorical* information: as infants begin to routinely encode objects in terms of their taxonomic categories (e.g., a spoon, a cookie), these specific categorical descriptors also come to be included at the structural level of the PR system, along with the more general (ontological and functional) categorical descriptors discussed above (e.g., Xu & Carey, 1996; Xu, 2002). The second change concerns the *spatiotemporal* information: not only is this information represented more accurately over time (e.g., as depth perception improves; Yonas & Granrud, 1984), but descriptions of object interactions also become more precise. For example, infants initially fail to distinguish between events in which an object is released *on top of* or *against* another object; in each case, infants simply encode “object released in contact with another object”, and they hold no particular expectation about whether the released object should remain stable or fall. By 4.5 to 5.5 months of age, however, infants identify *support* as an event category: they recognize that when an object, or ‘supportee’, is placed on top of another object, or ‘supporter’, the supportee’s fall is blocked by the supporter. As a result, infants now expect objects to remain stable when released on top of, but not against, other objects (Li, Baillargeon, & Needham, 2011).

Detecting structural violations. In the first months of life, infants’ physical representations typically include only structural information and are therefore rather sparse. Nevertheless, this information is still sufficient, when interpreted by the PR system’s causal framework, to allow infants to hold at least some expectations about physical events. To uncover these expectations, researchers often use *violation-of-expectation* tasks. In a typical task, infants see two test events: an expected event, which is consistent with the expectation examined in the experiment, and an unexpected event, which violates this expectation. With appropriate controls, evidence that infants look reliably longer at the unexpected than at the expected event is taken to indicate that they possess the expectation under investigation and detect the violation in the unexpected event.

For example, young infants detect a *persistence violation* (see **Figure 1**) when a cover is lowered over an object, slid to the side, and then lifted to reveal no object (Leslie, 1995; Wang et al., 2005); when an object is lowered inside an open container which is then slid forward and to the side to reveal the object standing in the container's initial position (Hespos & Baillargeon, 2001b); when an object is placed behind a screen, which then rotates through the space occupied by the object (Baillargeon, 1987; Baillargeon, Spelke, & Wasserman, 1985); and when a cover is lowered over a closed object and then lifted to reveal an open object (Wu & Baillargeon, 2011).

In addition, young infants detect an *inertia violation* when an inert object, after being set into motion, abruptly reverses direction to return to its starting position (Luo et al., 2009); and they detect a *gravity violation* if an inert object remains suspended when released in midair (e.g., Luo et al., 2009; Needham & Baillargeon, 1993). Interestingly, young infants do not view these last two events as violations if the object is self-propelled rather than inert (Kochukhova & Gredebäck, 2007; Luo et al., 2009). This is not to say that young infants believe self-propelled objects are not subject to the same principles as inert objects. Rather, infants assume that a novel self-propelled object may be able to use its internal energy to exert some control over its horizontal and vertical displacements: in particular, it may be able to change direction at will and to "resist" falling. At the same time, young infants recognize that there are limits to what self-propelled objects can do, despite their internal energy: for example, infants appreciate that self-propelled objects cannot spontaneously disappear or pass through obstacles (e.g., Aguiar & Baillargeon, 1999; Luo et al., 2009; Saxe, Tzelnic, & Carey, 2006; Spelke, Kestenbaum, Simons, & Wein, 1995a).

Kinds of explanations. The discussion in the preceding paragraph illustrates a key feature of physical reasoning in early infancy: young infants' explanations for events tend to be shallow, abstract explanations almost entirely divorced of mechanistic details. Upon seeing that a novel box initiates its own motion, for example, young infants endow it with internal energy and assume it can also use this energy to "resist" falling when released in midair. As adults, we may naturally find this assumption puzzling: we possess sufficient physical knowledge to realize that (outside of the realm of science-fiction) a self-propelled box would be highly unlikely to remain perfectly stationary in midair.

Infants' physical reasoning brings to mind Keil's (1995) suggestion that adults' concepts are "embedded in theory-like structures which owe their origins to a small but diverse set of fundamental modes of construal...one key part of these early modes of construal may be more general expectations...[that] exist before any specific explanation or detailed intuitive theory, and thus indicate *kinds of explanations* rather than any particular explanation" (pp. 260-261, italics added). In line with Keil's suggestion, we believe that the PR system's core concepts and principles provide infants with shallow kinds of explanations, rather than with specific or detailed mechanistic explanations, for events.

Variable Information

We have just seen that, in the first few months of life, the PR system includes only structural information in its physical representations of events. Although this information captures many essential elements, it is still very limited. If a spoon is placed inside a pot, for example, the structural information about the event specifies that a non-human, non-agentive, inert, closed object has been placed inside a non-human, non-agentive, inert container—but it does not specify the size, shape, pattern, and color of either object. How does this more detailed information come to be added to infants' physical representations?

For each event category, infants gradually identify a host of *variables* that help them interpret and predict outcomes within the category (e.g., Baillargeon, Needham, & DeVos, 1992; Hespos & Baillargeon, 2008; Kotovsky & Baillargeon, 1998; Luo & Baillargeon, 2005; Wang, Kaufman, Baillargeon, 2003; Wilcox, 1999). A variable both calls infants' information to a certain type of information in an event (e.g., features of objects or their arrangements) and provides a causal rule for interpreting this information. To illustrate, consider some of the variables infants identify for containment events (see **Figure 2**). By about 4 months of age, most infants have identified *width* as a relevant variable, and they now detect a violation when a wide object becomes fully hidden inside a narrow container (Wang, Baillargeon, & Brueckner, 2004). By about 7.5 months of age, most infants have identified *height* as a containment variable: they now detect a violation if a tall object becomes almost fully hidden inside a short container (Hespos & Baillargeon, 2001a). By about 9.5 months of age, most infants have identified *container-surface* as a relevant variable: they now detect a violation if an object becomes fully hidden when placed inside a container that is made of a transparent or holey material (Luo & Baillargeon, 2011). Finally, by about 14.5 months of age, infants identify color as a containment variable: they now detect a violation if a yellow toy is lowered inside a container (large enough to hide only one toy), and a purple toy is then removed from the container (Setoh & Baillargeon, 2011).

With the gradual identification of variables, infants' physical representations become increasingly richer (see **Figure 3**). After representing the structural information about an event and using this information to categorize the event, the PR system accesses the list of variables that have been identified as relevant for predicting outcomes in the category selected. The PR system then gathers information about each variable and includes this information in the physical representation of the event. This variable information is then interpreted by the variable rules as well as by the PR system's causal framework.

To illustrate this process, consider what variable information 7.5-month-olds would include in their physical representation of a containment event in which a ball was lowered inside a box. By 7.5 months, as we just saw, width and height have been identified as containment variables, but container-surface and color have not. Thus, infants would include information about the relative widths and heights of the ball and box in their physical representation of the event, but *not* information about the container's surface or about the ball's color. As a rule, the PR system does not include information about variables that have not yet been identified in its physical representation of an event (e.g., Hespos & Baillargeon, 2001a; Kotovsky & Baillargeon, 1998; Luo & Baillargeon, 2005; Newcombe, Huttenlocher, & Learmonth, 1999).

Décalages. We have just seen that, for each event category, infants gradually identify variables that enable them to better predict outcomes within the category. One might ask whether variables identified in the context of one event category are generalized to other categories, when equally relevant. For example, after infants learn to attend to height information in containment events, do they also attend to this information in covering or in tube events? Interestingly, the answer to such questions has turned out to be negative: a variable identified in the context of one category is *not* generalized to other categories, even when equally relevant (e.g., Hespos & Baillargeon, 2001a, 2006; Wang et al., 2005; Wang & Baillargeon, 2006; Wang & Kohne, 2007). Infants thus learn separately about each event category.

In some cases, infants may identify a variable at about the same age in different event categories. For example, the variable width is identified at about the same age in occlusion and in containment events (Wang et al., 2004). In other cases, however, several months may separate the acquisition of the same variable in different event categories, resulting in marked lags or *décalages* (to use a Piagetian term) in infants' responses to similar events in different event categories. As a case in point, consider the variable height (see **Figure 4**). Although this variable is identified at about 7.5 months in containment events, as we just saw (Hespos & Baillargeon, 2001a), it is identified earlier, at about 3.5 months, in occlusion events (Baillargeon & DeVos, 1991); it is identified later, at about 12 months, in covering events (Wang et al., 2005); and it is identified later still, at about 14.5 months, in tube events (Wang et al., 2005).

Décalages have been observed not only in violation-of-expectation tasks, as we just saw, but also in action tasks (e.g., Hespos & Baillargeon, 2006; Wang & Kohne, 2007). In one experiment, for example, 6- and 7.5-month-olds first played with a tall stuffed frog (Hespos & Baillargeon, 2006). Next, the frog was placed behind a large screen, which was then removed to reveal a tall and a short occluder (occlusion condition) or a tall and a short container (containment condition). The occluders were identical to the front halves of the containers; two frog feet protruded on either side of each occluder or through small holes at the bottom of each container. At both ages, infants were reliably more likely to search for the frog behind the tall as opposed to the short occluder; however, only the 7.5-months-olds (who had identified height as a containment variable) were reliably more likely to search for the frog inside the tall as opposed to the short container. Control infants who did not see the frog tended to reach about equally for the two occluders or containers.

Explanation-based learning. We suggested earlier that infants learn, with experience, what variables are helpful for interpreting and predicting outcomes in each event category. How does this learning process occur? Building on work in machine learning by DeJong (1993, 1997), we have proposed that the identification of a variable depends on an *explanation-based learning* (EBL) process that involves three main steps (e.g., Baillargeon et al., 2009a; Wang & Baillargeon, 2008).

First, infants must notice *contrastive outcomes* relevant to the variable. This occurs when infants build similar physical representations for two or more events—and notice that the events have contrastive outcomes. For example, consider the variable height in covering events, which is typically identified at about 12 months of age (e.g., Wang et al., 2005). We suppose that at some point prior to 12 months of age, infants begin to notice—as they manipulate covers and objects, or as they observe others doing so—that when a cover is lowered over an object, the object sometimes remains partly visible beneath the cover and sometimes does not. Infants thus notice contrastive outcomes they cannot predict based on their current variable knowledge: similar physical representations ("cover lowered over object") lead to contrastive outcomes ("object remains partly visible beneath cover" versus "object becomes fully hidden"), suggesting that a crucial piece of information is missing from the representations.

At this point, infants begin to search for the *conditions* that map onto these contrastive outcomes. Specifically, infants attempt to determine under what condition one outcome is observed, and under what condition the other outcome is observed. Eventually, infants uncover a regularity linking each outcome with a distinct condition (we assume that infants' statistical learning mechanisms help detect these regularities; e.g., Fiser & Aslin, 2002; Saffran, 2009). In the case of the variable height in covering events, infants detect that objects remain partly visible when placed under covers that are shorter than the objects, and become fully hidden when placed under covers that are as tall as or taller than the objects.

Finally, infants attempt to generate an *explanation* for the condition-outcome regularity detected, based on their prior knowledge. According to the EBL process, *only* condition-outcome regularities for which explanations can be provided are recognized as new variables. These explanations are typically very shallow (e.g., Keil, 1995; Luo et al., 2009; Wilson & Keil, 2000), but they still serve to integrate new variables with infants' existing causal knowledge (by the same token, these explanations also serve to prevent infants from learning incorrect or spurious variables). In the case of the variable height in covering events, infants' principle of persistence can provide a ready explanation for their observations: because an object continues to exist and retains its height when under a cover, it can become fully hidden only if its height is equal to, or shorter than, that of the cover.

After a new variable has been identified (i.e. is added to the list of variables relevant to an event category), infants begin to routinely include information about the variable in their physical representations of events from the category.

The EBL process helps make clear why infants learn separately about each event category. Infants do not compare arbitrary groups of events and look for invariants or critical variables that might explain similarities or differences among the events. The only situation that can trigger the identification of a variable is one where events with similar physical representations yield (as yet unpredicted or unexplained) contrastive outcomes. The learning process is thus highly constrained: it is designed to compare apples with apples, and not apples with rabbits or spoons.

Teaching experiments. The EBL process predicts that infants who have not yet identified a variable in an event category should be able to identify the variable—even several months before they would normally do so—if exposed in the laboratory (or the home) to appropriate observations for the variable. And indeed, a number of "teaching" experiments have now provided evidence for this prediction (e.g., Baillargeon, 2002; Wang & Baillargeon, 2008a; Wang & Kohne, 2007).

For example, in a series of experiments, Wang and her colleagues “taught” 9-month-old infants the variable height in covering events (recall that this variable is typically not identified until about 12 months of age; e.g., Wang et al., 2005). Infants received three pairs of teaching trials. In each pair of trials, a tall and a short cover (that differed only in height) were lowered over a tall object; infants could see that the object remained partly visible beneath the short cover, but became fully hidden under the tall cover. Different covers were used in the three pairs of teaching trials. Next, the infants received either a violation-of-expectation task or an action task involving novel covers and objects. In the violation-of-expectation task, infants looked reliably longer (even after a 24-hour delay) when a tall object became fully hidden under a short as opposed to a tall cover (Wang & Baillargeon, 2008a). In the action task, infants searched correctly for a tall object under a tall as opposed to a short cover (Wang & Kohne, 2007).

From an EBL perspective, these results are readily interpretable. During the teaching trials, (1) the infants noticed that events with similar physical representations led to contrastive outcomes; (2) they uncovered the specific height conditions that mapped onto these outcomes; and (3) they built an explanation for this condition-outcome regularity using their prior knowledge. Height was then added to the list of variables identified as relevant to covering events. When the infants next encountered covering events, they attended to the height information in the events, which enabled them to detect the violation in the violation-of-expectation task and to search correctly in the action task.

Two additional results supported this analysis. First, infants failed at the violation-of-expectation task if they received inappropriate teaching trials for which no explanation was possible (Wang & Baillargeon, 2008a; see also Newcombe, Sluzenski, & Huttenlocher, 2005). In this experiment, false bottoms were inserted into the teaching covers, rendering them all 2.5 cm deep; when the covers were rotated forward to reveal their interiors, the infants could see that they were all shallow. Thus, in each pair of teaching trials, the infants still observed that the tall object became fully hidden under the tall cover and partly hidden under the short cover—but they could no longer build an explanation for this condition-outcome regularity, because the tall and short covers were now equally shallow (i.e. it did not make sense that the tall object became fully hidden under the tall but shallow covers). Second, infants failed at the action task if they received appropriate teaching trials but were tested with tubes instead of covers (Wang & Kohne, 2007). When the tops of the tall and short covers were removed to form tubes, infants searched for the tall object in either the tall or the short tube, suggesting that they had identified height as a variable relevant to covering events and did not generalize this variable to tube events.

Together, the results summarized in this section suggest that infants can be taught a new variable in an event category through brief exposure to appropriate observations for the variable. Furthermore, infants who are taught a new variable immediately attend to information about the variable in situations presenting different stimuli and calling for different responses—but only when these situations involve events from the *same* category. The EBL process ensures broad, yet circumscribed, generalization: a variable identified in an event category is attended to in *any* event from the category—but *only* in events from the category.

Part II: The Acquisition of Numerical Knowledge

A representation of a given object is a representation of a unique individual. Representations of individuals are inputs into a variety of quantificational computations, such as which of two objects is bigger, which of two sets is more numerous, exactly how many objects are in a given container, and so on. We now turn to one type of quantification—by number.

As we will see, the development of number representations provides a case study of conceptual discontinuities: children must build representational resources qualitatively different from the initial state. In this way, number representations differ from object representations. Core cognition supports learning about objects, as Part I demonstrates, but this learning does not require forming object representations with very different properties from those available to young infants. I shall argue that core cognition contains two systems of representation with numerical content: *analog magnitude representations of number* and *parallel individuation of small sets of entities in working-memory models*. Both systems of representation continue to articulate thought throughout life and play crucial roles in later mathematical development. However, as I will show, the representations within them, alone, cannot express mathematical concepts such as integer or fraction. Thus, we need to understand how the latter concepts arise.

This question—how humans create representational resources that are discontinuous with those that are the input to the learning processes that create them—is of great scientific interest. Moreover, this case study is important for social and educational reasons also. Preschoolers vary vastly in their mastery of counting and the simple arithmetical algorithms that depend upon counting (e.g., counting up to add), and this variability predicts academic success in elementary school more than does variability in reading readiness, vocabulary, or many other predictors (Duncan et al., 2007). It is important, then, to understand what is difficult about acquiring the earliest mathematical knowledge and to understand how this learning is achieved. In what follows, I will sketch answers to these questions, detailing the core cognition systems with numerical content, showing how they are discontinuous with the integers, and describe a learning process through which children navigate the conceptual achievement of creating a representational resource that is capable of expressing natural numbers.

Core Cognition System 1: Analog Magnitude Representations of Number

Human adults, human infants, and non-human animals deploy a system of analog magnitude representations of number. Number is represented by a physical magnitude that is roughly proportional to the number of individuals in the set being enumerated. **Figure 5** depicts an external analog magnitude representational system in which length represents number. A psychophysical signature of analog magnitude representations is that the discriminability of any two magnitudes depends on their ratio. That is, discriminability is in accordance with Weber's law. This is a coding scheme widely used to represent dimensions of experience, such a loudness, time, brightness, length, size, intensity of pain, and many others. Animals and humans do not confuse these different dimensions of experience, but they use similar representational systems to encode them.

Dehaene (1997) and Gallistel (1990) review the evidence for the long evolutionary history of analog magnitude number representations. Animals as disparate as pigeons, rats, and non-human primates all represent number, and number discriminability satisfies Weber's law. In the past years, four different laboratories have provided unequivocal evidence that preverbal infants form analog magnitude representations of number as well (Brannon, 2002; Brannon, Abbot & Lutz, 2004; Lipon & Spelke, 2003, 2004; McCrink & Wynn, 2004; Wood & Spelke, 2005; Xu & Spelke, 2000; Xu, Spelke & Goddard, 2005). The first paper in this flurry of studies was by Fei Xu and Elizabeth Spelke, who solved the problem of how to control for other possible bases of judgment (cumulative surface area, element size, density) in a large number habituation paradigm. The authors habituated 6-month-old infants either to displays containing 8 dots, or to displays containing 16 dots. Possible confounds between number and other variables were controlled either by equating the two series of stimuli on those variables, or by making the test displays equidistant from the habituation displays on them. Infants who were habituated to 8-dot displays recovered interest when shown the novel 16-dot displays, while generalizing habituation to the novel 8-dot displays. Those habituated to 16-dot displays showed the reverse pattern. Subsequent studies duplicated this design (and the positive result) with 16-dot vs. 32 dot comparisons and with 4 dot vs. 8 dot comparisons.

That analog magnitude representations support these discriminations is shown by the fact that success is a function of the ratio of the set sizes. In all of the above studies, in which 6-month-old infants succeeded with a 2:1 ratio, they failed in comparisons that involved a 3:2 ratio (i.e., they failed to discriminate 8-dot from 12-dot arrays, 16-dot from 24-dot arrays, and 4-dot from 6-dot arrays). Also, these researchers have found that sensitivity improves by 9 months of age. Infants of this age succeed at 3:2 comparisons across a wide variety of absolute set sizes, but fail at 4:3 comparisons. Subsequent studies showed analog magnitude representations of number of different kinds of individuals (jumps, sounds), with the same profiles of sensitivity (Lipton & Spelke, 2004; Wood & Spelke, 2005).

In all the studies presented so far, we can be confident it is number infants are responding to, because every other variable has been equated either across the habituation stimuli or across the test stimuli. Of course, if the analog magnitude representations underlying performance in these habituation studies are truly numerical representations, number relevant computations other than establishing numerical equivalence should be defined over them, and indeed this is so. Elizabeth Brannon (2002) showed that 11-month-old infants represent numerical order using analog magnitude representations of sets. Koleen McCrink and Karen Wynn showed that 9-month-olds can manipulate sets of objects in the analog magnitude range to support addition, subtraction, and ratio computation (McCrink & Wynn, 2004). In sum, analog magnitude representations of number are available at least by 6 months of age. Preverbal infants represent the approximate cardinal value of sets, and compute numerical equivalence, numerical order, addition, subtraction, and ratios over these representations.

Core Cognition System 2: Parallel Individuation of Small Sets

Science moves rapidly, and the infant studies reviewed above came relatively late in the history of studies designed to show that infants are sensitive to number. The first studies, some 20 years earlier, concerned *small* sets—discriminations among sets of 1, 2 and 3 objects. These include many 2 vs. 3 habituation studies and Wynn's 1 + 1 = 2 or 1 violation-of-expectation study (e.g., Antell & Keating, 1983; Starkey & Cooper, 1980, Wynn, 1992b). In the latter study, infants were shown a single object on a stage, which was then hidden by a screen. Next, another object was introduced behind the screen. Finally, the screen was lowered to reveal either one object or two objects. Infants' attention was drawn to the unexpected outcome of a single object. Although some have suggested that analog magnitude number representations underlie success in these experiments (e.g., Dehaene, 1997), the evidence conclusively implicates a very different representational system (Feigenson & Carey, 2003, 2005; Feigenson, Carey, & Hauser, 2002; Scholl & Leslie, 1999; Simon, 1997; Uller, Carey, Huntley-Fenner, & Klatt, 1999). In this alternative representational system, number is only implicitly encoded; there are no symbols for number at all, not even analog magnitude ones. Instead, the representations include a symbol for each individual in an attended set. Thus, a set containing one apple might be represented: "O" (an iconic object file) or "apple" (a symbol for an individual of the kind apple) and a set containing two apples might be represented "O O" or "apple apple," and so forth. These representations consist of one symbol (file) for each individual, and when the content of a symbol is a spatiotemporally determined object, it is an object file (Kahnemann et al., 1992).

Infants also create working-memory models of small sets of other types of individuals, such as sound bursts or events, and so I shall call the system of representation “parallel individuation” and the explicit symbols within it “individual files.”

There are many reasons to favor individual file representations over analog magnitude representations as underlying performance in most of the infant small number studies (see Carey, 2009, for a more thorough review). First, and most important, success on many spontaneous number representation tasks involving small sets do not show the Weber-fraction signature of analog magnitude representations; rather they show the set-size signature of individual file representations. Individuals in small sets (sets of 1, 2 or 3) can be represented, and sets outside of that limit cannot, even when the sets to be contrasted have the same Weber-fraction as those small sets where the infant succeeds at that age. This is the set size signature of individual file representations.

Here I describe just one paradigm that elicits the set size signature of parallel individuation (see Carey, 2009, for others). An infant watches as each of two opaque containers, previously shown to be empty, is baited with a different number of graham crackers. For example, the experimenter might put two crackers in one container and three in the other. After placement, the parent allows the infant to crawl toward the containers. The dependent measure is which container the baby chooses. Ten- to 12-month-olds infants succeed at 1 vs. 2, 2 vs. 3, and 1 vs. 3, but fail at 3 vs. 4, 2 vs. 4, and even 1 vs. 4 (Feigenson & Carey, 2005; Feigenson et al., 2002). Although 1:4 is a more favorable ratio than 2:3, infants fail at 1 vs. 4 comparisons and succeed at 2 vs. 3. Note also that 5 crackers are involved in each choice, so the total length of time of placements is equated over these two comparisons. This is a striking result. Infants could succeed at 1 vs. 4 comparisons on many different bases: putting 4 crackers into a bucket takes much longer, draws more attention to that bucket, and so on, yet infants are at chance. Although infants could solve this problem in many different ways, apparently they are attending to each cracker and creating a model of what's in the container that contains one object-file for each cracker. As soon as one of the sets exceeds the limits on parallel individuation (apparently three at this age; see also the manual-choice paradigm in Feigenson & Carey, 2003, 2005), performance falls apart. This finding provides very strong evidence that parallel individuation underlies success on this task.

The purpose of parallel individuation is to create working-memory models of small sets of individuals, in order to represent spatial, causal, and intentional relations among them. Unlike analog magnitude number representations, the parallel-individuation system is not a dedicated number representation system. Far from it! The symbols in the parallel-individuation system explicitly represent individuals. This ability is closely tied to knowledge development in the domain of physical reasoning (see first part of chapter): Only because the infant has core knowledge about persistence, he/she can be sure that the number of individuals perceived in a given event does not change magically. Combining this knowledge with the individuation system, it is possible to understand occlusion or containment relations involving more than one object. Imagine, for example, an occlusion event in which two boxes are placed behind a screen. Since both boxes are temporarily out of view, the infant needs to keep track of how many individuals are occluded, and to represent their relation (occlusion) to the screen. **Figure 6** depicts several different possible individual file representations of two boxes. In none of these alternative models is there a symbol that has the content "two;" rather the symbols in working memory represent the boxes. The whole model {box box} represents two boxes, of course, but only implicitly.

If parallel individuation models do not include symbols for number, why am I discussing these models in the present context? The answer is that they are shot through with numerical content, even though that numerical content is merely implicit in the computations that index and represent small sets, that govern the opening of new individual files, that update working-memory models of sets as individuals are added or subtracted, and that compare sets on numerical criteria. The creation of a new individual file requires principles of individuation and numerical identity (see Part I. above; numerical identity means sameness in the sense of "same one"); models must keep track of whether this object (sound, event etc.), seen now, is the same *one* as that perceived before. The decision the system makes dictates whether an additional individual file is established, and this guarantees that a model of a set of three boxes will contain three box symbols. Computations of numerical identity are (as their name says) numerical computations. Also, the opening of a new individual file in the presence of other active files provides an implicit representation of the process of adding one to an array of individuals. Finally, working-memory models of two sets of individuals can be simultaneously maintained, and when individual-file models are compared on the basis of 1-1 correspondence, the computations over these symbols establish numerical equivalence and numerical order (for evidence of such computations, see Feigenson, 2005; Feigenson & Carey, 2003).

Conceptual Discontinuity

In the adult representational system, verbal numerals represent positive integers: they are summary symbols that represent cardinal values of sets, and the system of verbal numerals satisfies Peano's axioms. The central axiom is that each integer has a unique successor; for each integer, the next integer is n + 1. Implicitly or explicitly, the adult meanings of verbal numerals must be formulated over concepts such as *exactly 1*, *plus, exact numerosity, set, successor*.

Within this set of primitives, "four" can be defined in many different ways. Here are two: 1) four is the cardinal value of any set that can be put into 1-1 correspondence with {x, y, z, w}, where "x," "y," "z," and "w" refer to numerically distinct individuals and {} denotes a set; and 2) four is the successor of three. These two ways of defining "four" are formally equivalent, in the sense that each definition determines exactly the same concept, namely, *four*. For "four" to have the same meaning as it does in the adult representational system, its meaning must be equivalent to both of these. Of course, it is an open question just what set of primitives underlies the child's first successful representation of *four*. Whatever it is, it must provide the expressive power of the primitives listed above.

Neither of the core cognition systems discussed in the previous sections has the capacity to represent the integers. Parallel individuation includes no summary symbols for number at all, and has an upper limit of 3 or 4 on the size of sets it represents. Analog magnitude representations include summary symbols for cardinal values that are embedded within a system of arithmetical computations, but they represent only approximate cardinal values; there is no representation of exactly 1, and therefore no representation of + 1. Analog magnitude representations cannot even resolve the distinction between 10 and 11 (or any two successive integers beyond its discrimination capacity), and so cannot express the successor function. Thus, neither system can represent *10*, let alone *342,689,455*. This analysis makes precise the senses in which the concepts expressed by the verbal numeral list are qualitatively different from those representations that precede it. The numeral list has more expressive power—it can represent infinitely more concepts than can either of the core cognition systems with numerical content.

Because the numeral list is qualitatively different from each of the core cognition systems with numerical content, it is indeed difficult to learn. American middle-class children learn to recite the count list and to carry out the count routine in response to the probe “how many” shortly after their second birthday. They do not learn how counting represents number for another 1 ½ or 2 years, however. Young two-year-olds first assign a cardinal meaning to “one,” treating other numerals as equivalent plural markers that contrast in meaning with “one.” Some 7 to 9 months later they assign cardinal meaning to “two,” but still take all other numerals to mean essential “some,” contrasting only with “one” and “two.” They then work out the cardinal meaning of “three” and then of “four.” This protracted period of development is called the “subset”-knower stage, for children have worked out cardinal meanings for only a subset of the numerals in their count list. So far, we still do not know exactly what representations underlie the meanings of “one” through “four” in the subset-knower stage. It seems likely, however, that they draw on the resources of parallel individuation. LeCorre and Carey (2007) proposed a system of “enriched parallel individuation” in which the child creates long-term memory models of sets of 1 through 4 objects, mapping the verbal numerals to them using a rule that specifies that any set that can be put in 1-1 correspondence to the memory representation of a singleton set has “one” object, and similarly for sets of 2 through 4.

Many different tasks, which make totally different information-processing demands on the child, confirm that subset-knowers differ qualitatively from children who have worked out how counting represents number. Subset-knowers cannot create sets of sizes specified by their unknown numerals (Wynn, 1990, 1992b), cannot estimate the cardinal values of sets outside their known numeral range (Le Corre, Brannon, Van de Walle, & Carey, 2006), do not know what set-size is reached if 1 individual is added to a set labeled with a numeral outside their known numeral range (Sarnecka & Carey, 200), and so on. Children who succeed on one of these tasks succeed on all of them. Furthermore, a child diagnosed as a “one”-knower on one task is also a “one”-knower on all of the others, ditto for “two”-knowers, “three”-knowers and “four”-knowers.

Thus, learning how the verbal count list works is extremely difficult, the process unfolding over a two-year period. Adults who lack a count list (because their language does not contain one) demonstrate profound limitations in non-verbal numerical reasoning (Gordon, 2004; Pica, Lemer, Izard, & Dehaene, 2004; Spaepen, Coppola, Spelke, Carey, & Goldin-Meadow, 2011), and preschoolers who do not fully master the numeral list representation of number are at a profound disadvantage in the elementary school curriculum. How is creating a representation of integers achieved?

Acquiring Representations of Integers

Ultimately, learning requires adjusting expectations, representations, and actions to data. Abstractly, all of these learning mechanisms are variants of hypothesis-testing algorithms. The representations most consistent with the available data are strengthened; those hypotheses are accepted. However, in cases of developmental discontinuity, the learner does not initially have the representational resources to state the hypotheses that will be tested, to represent the variables that could be associated or could serve as an input to a Bayesian learning algorithm. Carey (2009) describes one learning mechanism that underlies conceptual discontinuities—*Quinian bootstrapping*. Quinian bootstrapping can create new representational machinery, new concepts that articulate hypotheses previously unstatable.

In Quinian bootstrapping episodes, mental symbols are established that correspond to newly coined or newly learned explicit symbols. These are initially placeholders, getting whatever meaning they have from their interrelations with other explicit symbols. As is true of all word learning, newly learned symbols must necessarily be initially interpreted in terms of concepts already available. But at the onset of a bootstrapping episode, these interpretations are only partial—the learner (child or scientist) does not yet have the capacity to formulate the concepts that the symbols will come to express.

The bootstrapping process uses the set of interrelated symbols in the placeholder structure to model the phenomena in the domain, where the phenomena are represented in terms of whatever concepts the child or scientist has available. Both structures (the placeholder structure and the system of available concepts) provide constraints, some only implicit and instantiated in the computations defined over the representations. These constraints are respected as much as possible in the course of the modeling activities, which include analogy construction and monitoring, limiting case analyses, thought experiments, and inductive inference.

Bootstrapping Representations of Natural Number

In the case of the construction of the numeral list representation of the integers, the memorized count list is the placeholder structure. Its initial meaning is exhausted by the relation among the external symbols: they are stably ordered. "One, two, three, four…" initially has no more meaning for the child than "a, b, c, d…" The details of the subset-knower period suggest that the resources of parallel individuation, enriched by the machinery of linguistic set-based quantification, provide the partial meanings children assign to the placeholder structures that get the bootstrapping process off the ground. The meaning of the word "one" could be subserved by a mental model of a set of a single individual {i}, along with a procedure that determines that the word "one" can be applied to any set that can be put in 1-1 correspondence with this model. Similarly "two" is mapped onto a long-term memory model of a set of two individuals {j k}, along with a procedure that determines that the word "two" can be applied to any set that can be put in 1-1 correspondence with this model. And so on for "three" and "four." This proposal requires no mental machinery not shown to be in the repertoire of infants—parallel individuation, the capacity to compare models on the basis of 1-1 correspondence, and the set-based quantificational machinery that underlies the singular/plural distinction and makes possible the representation of dual and trial markers. The work of the subset-knower period of numeral learning, which extends in English-learners between ages 2:0 and 3:6 or so, is the creation of the long-term memory models and computations for applying them that constitute the meanings of the first numerals the child assigns numerical meaning to.

Once these meanings are in place, and the child has independently memorized the placeholder count list and the counting routine, the bootstrapping proceeds as follows. The child notices the identity between the words "one," "two," "three," and "four" that express the numerical meanings captured by the machinery of enriched parallel individuation and the first four words in the count list. The child must try to align these two independent structures. The critical analogy is between order on the list and order in a series of sets related by *one additional individual.* This analogy supports the induction that any two successive numerals will refer to sets such that the numeral farther in the list picks out a set that is 1 greater than that earlier in the list.

This proposal illustrates all of the components of bootstrapping processes: placeholder structures whose meaning is provided by relations among external symbols, partial interpretations in terms of available conceptual structures, modeling processes (in this case analogy), and an inductive leap. The greater representational power of the numeral list than that of any of the systems of core cognition from which it is built derives from combining distinct representational resources—a serially ordered list, the numerical content of parallel individuation (which is largely embodied in the computations carried out over sets represented in memory models with one symbol for each individual in the set). The child creates symbols that express information that previously existed only as constraints on computations. Numerical content does not come from nowhere, but the process does not consist of defining “seven” in terms of mental symbols for number available to infants.

Carey (2009) illustrates how Quinian bootstrapping underlies conceptual discontinuities of two types: 1) those that involve the construction of representational resources with more expressive power such is the case in mastering the verbal integer list (sketched here) or the construction of rational number, and 2) episodes of theory changes involving successive conceptual systems that are incommensurable with each other. I draw my examples from early childhood through historical theory changes achieved by metaconceptually sophisticated scientists.

Perhaps you have noted a surprising fact about my account of the initial construction of integer representations. The process makes no use of one of the systems of core cognition with numerical content—the analog-magnitude (AM) system. This is so, and Carey (2009) summarizes the evidence for this surprising turn of events. However, another episode of Quinian bootstrapping achieves the integration of AM number representations with verbal numerals within 6 months of the child’s becoming a “cardinal principle-knower” (Le Corre & Carey, 2007). This integration plays a very important role in subsequent arithmetic learning and numerical reasoning. Impairments in the AM system through brain damage lead to breakdowns of mathematical reasoning in adults, and are implicated in dyscalculia, a specific learning disability in mathematics (Butterworth, 1999; Dehaene, 1997). Two recent studies have shown that even within normally developing children, those with high resolution in their non-verbal AM system outperform those with low resolution on school arithmetic learning, beginning in kindergarten (Gilmore, McCarthy & Spelke, 2010; Halberda, Mazzocco, and & Feigenson, 2008). Thus, although Quinian bootstrapping is required to construct representations of integers, as well as fractions and other later developing mathematical concepts, the representations within core cognition play an important role in learning and reasoning throughout all of life.

Conclusions: Core Cognition and Beyond

Systems of core cognition, such as the physical-reasoning system sketched in Part I, and the analog-magnitude and parallel-individuation systems sketched in Part II, are learning devices. They make it possible for infants to represent, and hence learn about, the objects in their world, and their spatial, causal, and quantitative relations to each other. As described in Part I, during infancy the PR system itself is enriched through explanation-based learning (EBL), as the child identifies new variables that are relevant to the predictions and explanations that constitute physical reasoning. The learning achieved in this way enriches the physical-reasoning system, but does not transcend it. It does not require creating new representational primitives; the variables identified are already available as input to the statistical processes that are part of EBL. EBL goes beyond mere statistical learning, for it is constrained by the causal/explanatory knowledge embedded in core cognition. Similarly, the quantitative machinery of the core systems of number representation makes it possible for the child to learn quantitative regularities about their world.

The conceptual representations that constitute systems of core cognition are part of the building blocks of later developing conceptual representations, even those that are qualitatively different from core cognition, expressing concepts not found within it. As sketched in Part II, the resources of parallel individuation provide a large part of the meanings assigned to verbal numerals during the subset-knower stage. However, they cannot provide the full meaning of verbal numerals; parallel individuation cannot represent the meaning of “seven” or “ten,” for example, because working memory has a limit of three or four items during infancy and the preschool years. Quinian bootstrapping is required for the child to work out the trick of using serial order in the count list to determine the cardinal value of any numeral in the list, that is, to come to use counting to implement the successor function. And indeed, Sarnecka and Carey (2008) showed that children who have learned how counting represents number (children called “cardinal principle-knowers,” or “CP-knowers,” in the literature) differ from subset-knowers precisely in this way. Only CP-knowers can derive without counting that if you add one to a set containing 5 objects, you will have 6 objects.

Our goal in this chapter was two-fold. First, we provided worked examples of the type of research that reveals rich representational resources early in infancy, representational resources that provide building blocks for later conceptual understanding of the world, illustrating that these are dynamic learning mechanisms. Second, we illustrated the point that mastery of the adult conceptual repertoire often requires creating entirely new representational resources (in this case the integer-list representation of number). Through such constructions, children become able to think thoughts unavailable to infants or to any non-human animal.

Acknowledgements

The preparation of this chapter was supported by a grant from the National Institute of Child Health and Human Development to Renée Baillargeon (HD-021104), and by grants from NICHHD (HD-028338) and NSF (REC-0633955) to Susan Carey. We would like to thank Amélie Bernard and Peipei Setoh for helpful comments.

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.

Antell, S., & Keating, D. P. (1983). Perception of numerical invariance in neonates. *Child Development*, *54*, 695-701.

Baillargeon, R. (1987). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology*, *23*, 655-664.

Baillargeon, R. (2002). The acquisition of physical knowledge in infancy: A summary in eight lessons. In U. Goswami (Ed.), *Blackwell handbook of childhood cognitive development* (pp. 47-83). Oxford: Blackwell.

Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science, 3,* 2-13.

Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: Further evidence. *Child Development*, *62*, 1227-1246.

Baillargeon, R., Li, J., Gertner, Y., & Wu, D. (2010). How do infants reason about physical events? In U. Goswami (Ed.), *The Wiley-Blackwell handbook of childhood cognitive development, second edition* (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 5-month-old infants. *Cognition*, *20*, 191-208.

Baillargeon, R., Wu, D., Gertner, R., Setoh, P., Kittredge, A. K., Stavans, M. (in press). Object individuation and physical reasoning in infancy: An integrative account. *Language Learning and Development.*

Baillargeon, R., Wu, D., Yuan, S., Li, J., & Luo, Y. (2009b). Young infants' expectations about self-propelled objects. In B. M. Hood & L. R. Santos (Eds.), *The origins of object knowledge* (pp. 285-352). Oxford: Oxford University Press.

Bonatti, L., Frot, E., & Mehler, J. (2005). What face inversion does to infants' numerical abilities. *Psychological Science*, *16*, 506-510.

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.

Brannon, E. M. (2002). The development of ordinal numerical knowledge in infancy. *Cognition*, *83*, 223-240.

Brannon, E. M., Abbott, S., & Lutz, D. J. (2004). Number bias for the discrimination of large visual sets in infancy. *Cognition*, *93*, B59-B68.

Butterworth, B. (1999). *What counts: How every brain is hardwired for math*. New York: Free Press.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Carey, S., & Spelke, E. S. (1994). Domain-specific knowledge and conceptual change. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and* *culture* (pp. 169-200). New York, NY: Cambridge University Press.

Casasola, M., Cohen, L., & Chiarello, E. (2003). Six-month-old infants’ categorization of containment spatial relations. *Child Development*, *74,* 679-693.

Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, *107*, 705-717.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13,* 148-153.

Dehaene, S. (1997). *The number sense*. Oxford: Oxford University Press.

DeJong, G. F. (1993). *Investigating explanation-based learning*. Boston: Kluwer Academic Press.

DeJong, G. F. (1997). Explanation-based learning. In A. Tucker (Ed.), *Encyclopedia of computer science* (pp. 499-520). Boca Raton, FL: CRC Press.

Duncan, G., Dowsett, C., Claessens, A. Magnuson, K., Huston, A., Klebanov, P., Pagani, L. Feinstein, L.,Engel, M., Brooks-Gunn, J., Sexton, H., & Duckworth, K. (2007). School readiness and later achievement. *Developmental Psychology*, *43*, 128-1446.

Feigenson, L. (2005). A double-dissociation in infants' representation of object arrays. *Cognition*, *95,* B37–B48.

Feigenson, L., & Carey, S. (2003). Tracking individuals via object files: evidence from infants' manual search. *Developmental Science, 6*, 568-584.

Feigenson, L., & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition*, *97*, 295-313.

Feigenson, L., Carey, S., & Hauser, M. D. (2002). The representations underlying infants' choice of more: object files versus analog magnitudes. *Psychological Science*, *13*, 150-156.

Fiser, J., & Aslin, R. N. (2002) Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, *99*, 15822-15826.

Futó, J., Téglás, E., Csibra, G., & Gergely, G. (2010). Communicative function demonstration induces kind-based artifact representation in preverbal infants. *Cognition, 117,* 1-8.

Gallistel, C. R. (1990). *The organization of learning.* Cambridge, MA: MIT Press.

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.

Gelman, R., Durgin, F., & Kaufman, L. (1995). Distinguishing between animates and inanimates: Not by motion alone. In D. Sperber, D. Premack & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 150-184). Oxford: Clarendon Press.

Gilmore, C.K., McCarthy, S. E., & Spelke, E. S. (2010). Non-symbolic arithmetic abilities and mathematics achievement in the first year of formal schooling. *Cognition*, *115*, 394-406.

Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science, 306*, 496-499.

Halberda, J., Mazzocco, M. M. M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature, 455,* 665–668.

Hespos, S. J., & Baillargeon, R. (2001a). Infants' knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science*, *12*, 140-147.

Hespos, S. J., & Baillargeon, R. (2001b). Knowledge about containment events in very young infants. *Cognition*, *78*, 204-245.

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.

Huttenlocher, J., Duffy, S., & Levine, S.C. (2002). Infants and toddlers discriminate amount: Are they measuring? *Psychological Science*, *13,* 244-249.

Johnson, S. C., Shimizu, Y. A., & Ok, S-J. (2007). Actors and actions: The role of agent behavior in infants' attribution of goals. *Cognitive Development*, *22*, 310-322.

Kahneman, D., Treisman, A., & Gibbs, B. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology,* 24, 175-219.

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.

Kestenbaum, R., Termine, N., & Spelke, E. S. (1987). Perception of objects and object boundaries by 3-month-old infants. *British Journal of Developmental Psychology*, *5*, 367-383.

Kochukhova, O., & Gredeback, G. (2007). Learning about occlusion: Initial assumptions and rapid adjustments. *Cognition, 105,* 26-46.

Kotovsky, L., & Baillargeon, R. (1998). The development of calibration-based reasoning about collision events in young infants. *Cognition, 67,* 311-351.

Le Corre, M., Brannon, E. M., Van de Walle, G. A., & Carey, S. (2006). Re-visiting the competence/performance debate in the acquisition of the counting principles. *Cognitive Psychology, 52,* 130-169.

Le Corre, M., & Carey, S. (2007). One, two, three, four, nothing more: How numerals are mapped onto core knowledge of number in the acquisition of the counting principles. *Cognition, 105*, 395-438.

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' system. *Trends in Cognitive Sciences, 2,* 10-18.

Li., J., Baillargeon, R., & Needham, A. (2011). Learning about support events in young infants. Manuscript in preparation.

Lipton, J. S., & Spelke, E. S. (2003). Origins of number sense: Large number discrimination in human infants. *Psychological Science, 15,* 396-401.

Lipton, J. S., & Spelke, E. S. (2004). Discrimination of large and small numerosities by human infants. *Infancy, 5,* 271-290.

Luo, Y., & Baillargeon, R. (2005). When the ordinary seems unexpected: Evidence for rule-based physical reasoning in young infants. *Cognition, 95,* 297-328.

Luo, Y., & Baillargeon, R. (2011). Infants' reasoning about transparent occluders and containers. Manuscript in preparation.

Luo, Y., Kaufman, L., & Baillargeon, R. (2009). Young infants' reasoning about events involving inert and self-propelled objects. *Cognitive Psychology, 58,* 441-486.

McDonough, L., Choi, S., & Mandler, J. M. (2003). Understanding spatial relations: Flexible infants, lexical adults. *Cognitive Psychology, 46,* 229-259.

McCrink, K., & Wynn, K. (2004). Large-number addition and subtractrion by 9-month old infants. *Psychological Science, 15*, 776-781.

Needham, A. (2001). Object recognition and object segregation in 4.5-month-old infants. *Journal of Experimental Child Psychology, 78,* 3-24.

Needham, A., & Ormsbee, S. M. (2003). The development of object segregation during the first year of life. In R. Kimchi, M. Behrmann, & C. Olson (Eds.), *Perceptual organization in vision: Behavioral and neural perspectives* (pp. 205-232). Mahwah, NJ: Erlbaum.

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.

Newcombe, N. S., Sluzenski, J., & Huttenlocher, J. (2005). Pre-existing knowledge versus on-line learning: What do infants really know about spatial location? *Psychological Science, 16,* 222-227.

Onishi, K. H. (2011). Perception of event roles in infancy. Manuscript in preparation.

Pauen, S., & Träuble, B. (2009). How 7-month-olds interpret ambiguous motion events: Category-specific reasoning in infancy. *Cognitive Psychology, 59,* 275-295.

Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian Indigene group. *Science, 306,* 499-503.

Premack, D. (1990). The infant's theory of self-propelled objects. *Cognition, 36,* 1-16.

Premack, D., & Premack, A. (2003). *Original intelligence: Unlocking the mystery of who we are*. New York: McGraw-Hill.

Pylyshyn, Z. W. (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial index model. *Cognition, 32,* 65-97.

Pylyshyn, Z. W. (1994). Some primitive mechanisms of spatial attention. *Cognition, 50,* 363-384.

Quinn, P. (2007). On the infant's prelinguistic conception of spatial relations: Three developmental trends and their implications for spatial language learning. In J. M. Plumert & J. P. Spencer (Eds.), *The emerging spatial mind* (pp. 117-141). New York: Oxford University Press.

Rose, S. A., Gottfried, A. W., Melloy-Carminar, P., & Bridger, W. H. (1982). Familiarity and novelty preferences in infant recognition memory: Implications for information processing. *Developmental Psychology*, *18,* 704-713.

Saffran, J. R. (2009). What can statistical learning tell us about infant learning? In A. Woodward & A. Needham (Eds.), *Learning and the infant mind* (pp. 29-46). New York: Oxford University Press.

Sarnecka, B., & Carey, S. (2008). How counting represents number: What children must learn and when they learn it. *Cognition, 108,* 662-67.

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.

Scholl, B. J., & Leslie, A. M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In E. Lepore & Z. Pylyshyn (Eds.), *What is Cognitive Science*? (pp. 26-73). Oxford: Blackwell.

Setoh, P., & Baillargeon, R. (2011, March). *How do infants generalize new physical expectations?* Paper presented at the Biennial meeting of the Society for Research in Child Development, Montreal, Canada.

Simon, T. J. (1997). Reconceptualizing the origins of number knowledge: A "non-numerical" account. *Cognitive Development*, *12*, 349-372.

Slater, A. (1995). Visual perception and memory at birth. In C. Rovee-Collier & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 9, pp. 107-162). Norwood: Ablex.

Spaepen, E., Coppola, M.' Spelke. E, Carey, S., & Goldin-Meadow, S. (2011). Number without a language model. *Proceedings of the National Academy of Sciences*, *108*, 3163-3168.

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*, 1-30.

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.

Starkey, P., & Cooper, R. (1980). Perception of numbers by human infants. *Science*, *210*, 1033-1035.

Surian, L., & Caldi, S. (2010). Infants' individuation of agents and inert objects. *Developmental Science*, *13*, 143-150.

Träuble, B., & Pauen, S. (2007). The role of functional information for infant categorization. *Cognition, 105*, 362-379.

Träuble, B., & Pauen, S. (2011a). Cause of effect—what matters? How 12-month-old infants learn to categorize artifacts. *British Journal of Developmental Psychology*. DOI: 10.1348/026151009X479547.

Träuble, B., & Pauen, S. (2011b). Infants' reasoning about ambiguous motion events: The role of spatiotemporal and dispositional status information. *Cognitive Development, 26*, 1-15.

Uller, C., Carey, S., Huntley-Fenner, G., & Klatt, L. (1999). What representations might underlie infant numerical knowledge. *Cognitive Development*, *14*, 1-36.

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 and solidity violations in infancy: A new account and new evidence from covering events. *Cognition*, *95*, 129-173.

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 11-month-old infants. *Developmental Science*, *12*, 681-687.

Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition*, *72*, 125-166.

Wilcox, T., & Chapa, C. (2002). Infants' reasoning about opaque and transparent occluders in an object individuation task. *Cognition*, *85*, B1-B10.

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.

Wood, J., & Spelke, E. (2005). Infants' enumeration of actions: Numerical discrimination and its signature limits. *Developmental Science*, *8*, 173-181.

Woodward, A. L., Phillips, A., & Spelke, E. S. (1993). Infants' expectations about the motion of animate versus inanimate objects. *Proceedings of the Fifteenth Annual Meeting of the Cognitive Science Society* (pp. 1087-1091). Hillsdale, NJ: Erlbaum.

Wu, D., & Baillargeon, R. (2011, March). *Young infants detect open-close but not variable changes to objects in covering events.* Paper presented at the Biennial meeting of the Society for Research in Child Development, Montreal, Canada.

Wynn, K. (1990). Children's understanding of counting. *Cognition*, *36*, 155-193.

Wynn, K. (1992a). Children's acquisition of the number words and the counting system. *Cognitive Psychology*, *24,* 220-251.

Wynn, K. (1992b). Addition and subtraction by human infants. *Nature*, *358*, 749-750.

Xu, F. (2002). The role of language in the acquiring object kind concepts in infancy. *Cognition*, *85*, 223-250.

Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, *30,* 111-153.

Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month old infants. *Cognition*, *74,* B1-B11.

Xu, F., Spelke, E., & Goddard, S. (2005). Number sense in human infants. *Developmental Science*, *8,* 88-101.

Yonas, A., & Granrud, C. E. (1984). The development of sensitivity to kinetic, binocular, and pictorial depth information in human infants. In D. Engle, D. Lee, & M. Jeannerod (Eds.), *Brain mechanisms and spatial vision* (pp. 113-145). Dordrecht: Martinus Nijhoff.

Figure Captions

Figure 1: Examples of basic persistence violations young infants can detect. Infants are surprised if a cover is lowered over an object, slid to the side, and then lifted to reveal no object (Wang et al., 2005); if an object is placed inside a container which is then slid forward and to the side to reveal the object in the container's original position (Hespos & Baillargeon, 2001b); if a screen rotates through the space occupied by a box in its path (Baillargeon, 1987); and if a cover is lowered over a closed object (e.g., a toy duck) and then lifted to reveal an open object (e.g., a toy duck with a large hole in its midsection; Wu & Baillargeon, 2011).

Figure 2: Some of the variables infants identify as they learn about containment events. By about 4 months of age, infants are surprised if a wide object becomes fully hidden inside a narrow container (Wang et al., 2004). By about 7.5 months, infants are surprised if a tall object becomes almost fully hidden inside a short container (Hespos & Baillargeon, 2001a). By about 9.5 months, infants are surprised if an object becomes hidden inside a transparent container (Luo & Baillargeon, 2011). Finally, by about 14.5 months, infants are surprised if an object changes color (e.g., from purple to yellow) when briefly lowered inside a container (too small to hide more than one object) (Setoh & Baillargeon, 2011).

Figure 3: Schematic model of the physical-reasoning system: How infants represent and interpret the structural and the variable information about a physical event.

Figure 4: Décalages in infants' identification of the variable height across different event categories. By about 3.5 months of age, infants are surprised if a tall object remains fully hidden when passing behind an occluder with a short midsection (Baillargeon & DeVos, 1991). By about 7.5 months, infants are surprised if a tall object becomes almost fully hidden inside a short container (Hespos & Baillargeon, 2001a). By about 12 months, infants are surprised if a tall object becomes fully hidden under a short cover (Wang et al., 2005). Finally, by about 14.5 months, infants are surprised if a tall object becomes fully hidden inside a short tube (Wang et al., 2005). In the latter two cases, infants are first allowed to inspect the cover or tube in a brief orientation procedure.

Figure 5: External analog magnitude representation of number in which number is represented by line length.

Figure 6: Two versions of the memory structures that might subserve parallel individuation of small sets of objects. In one, each object is represented by an object-file that abstracts away from specific features (OBJ). In the other, each object is represented by an object-file on which shape, color, texture and spatial extent features have been bound.

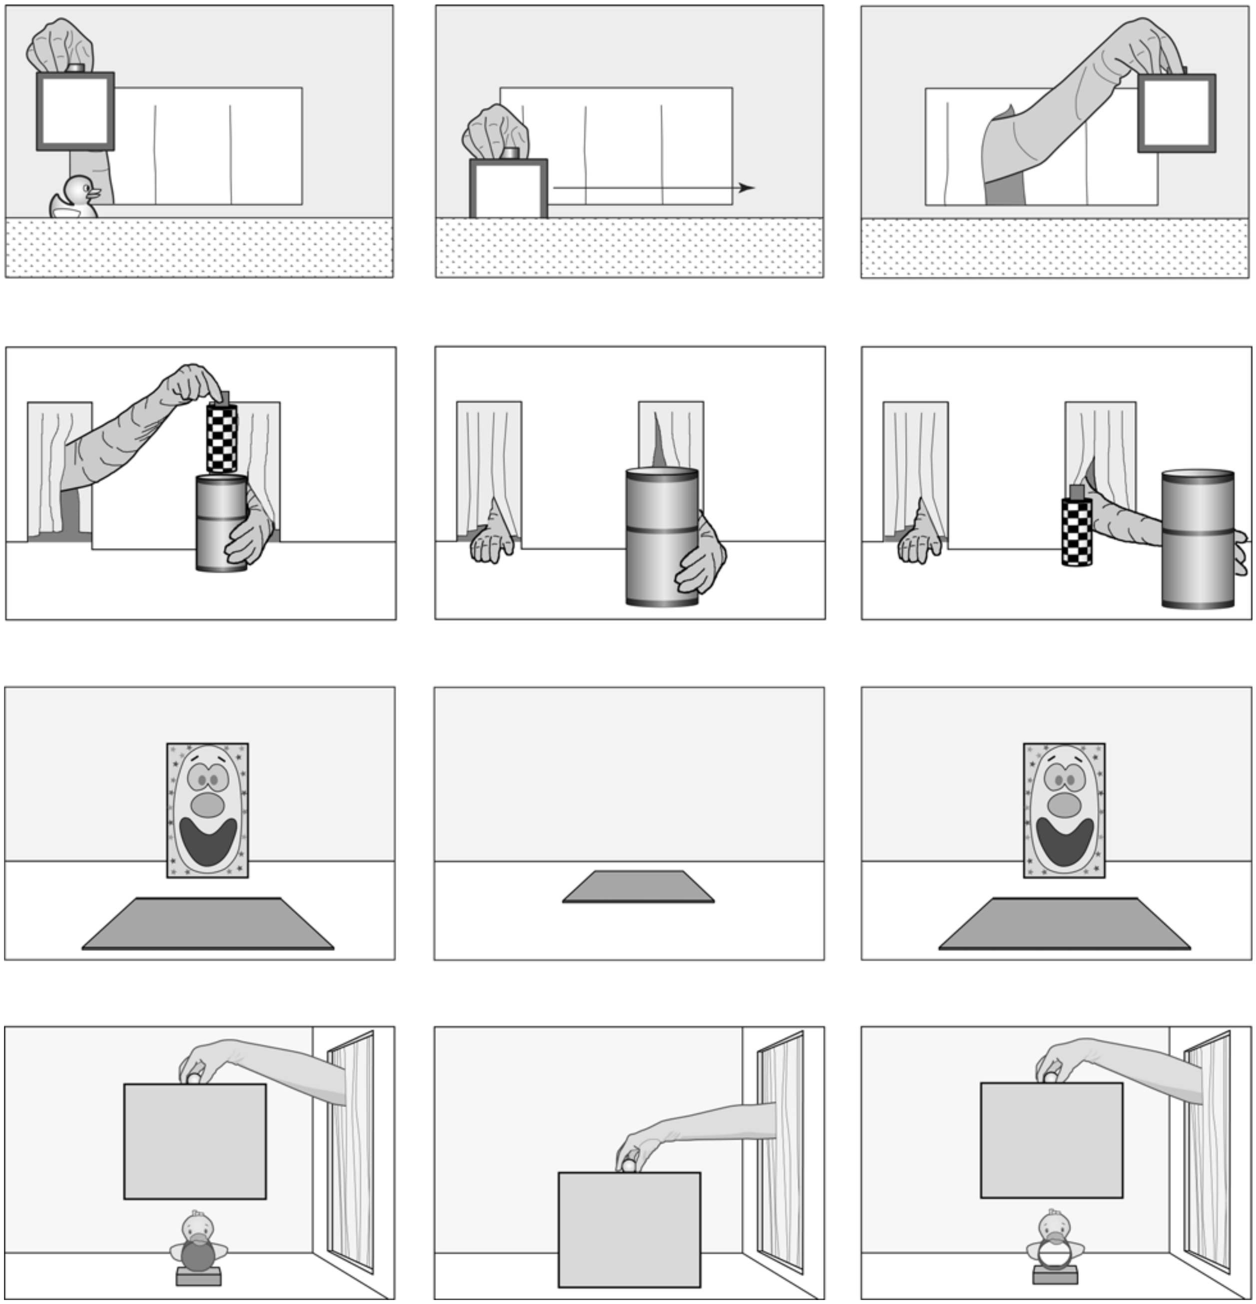

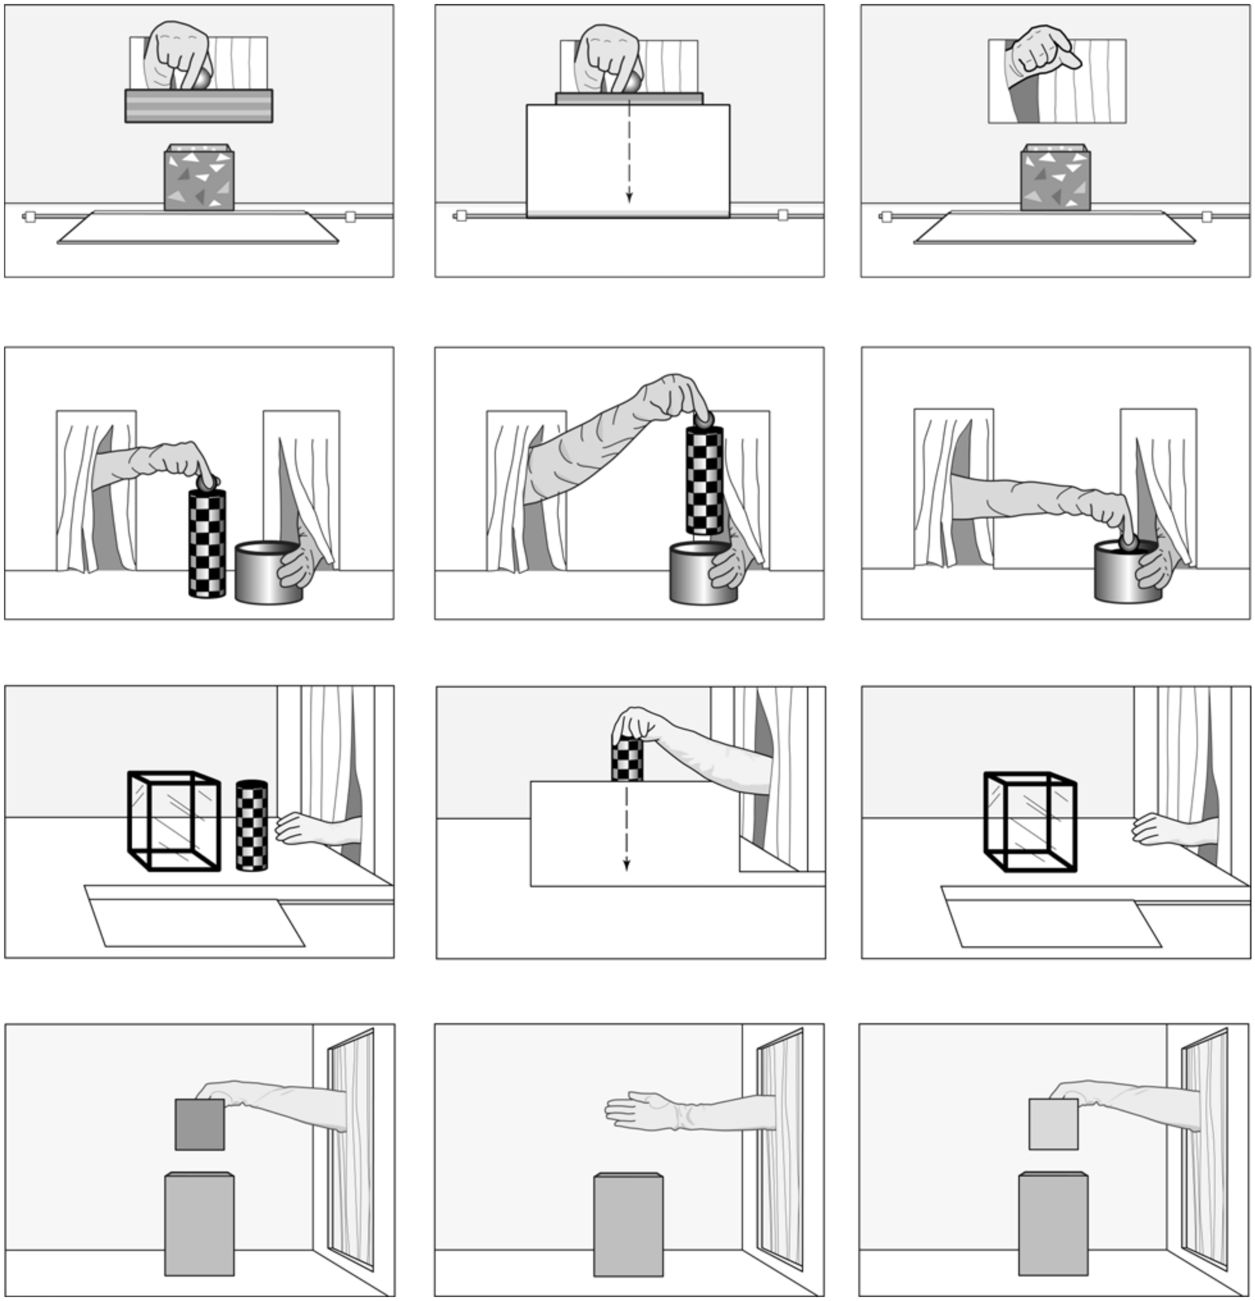

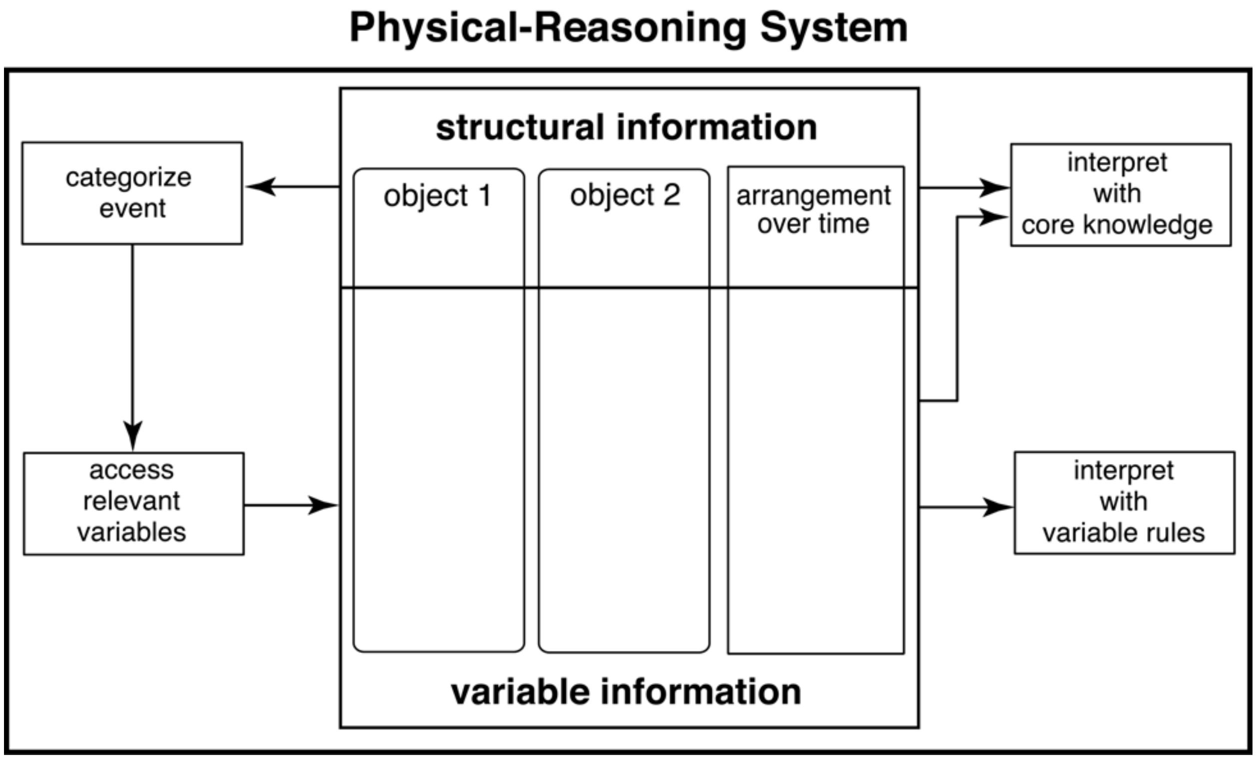
Physical-Reasoning System
structural information
categorize event
object 1
object 2
arrangement over time
interpret with core knowledge
access relevant variables
interpret with variable rules
variable information

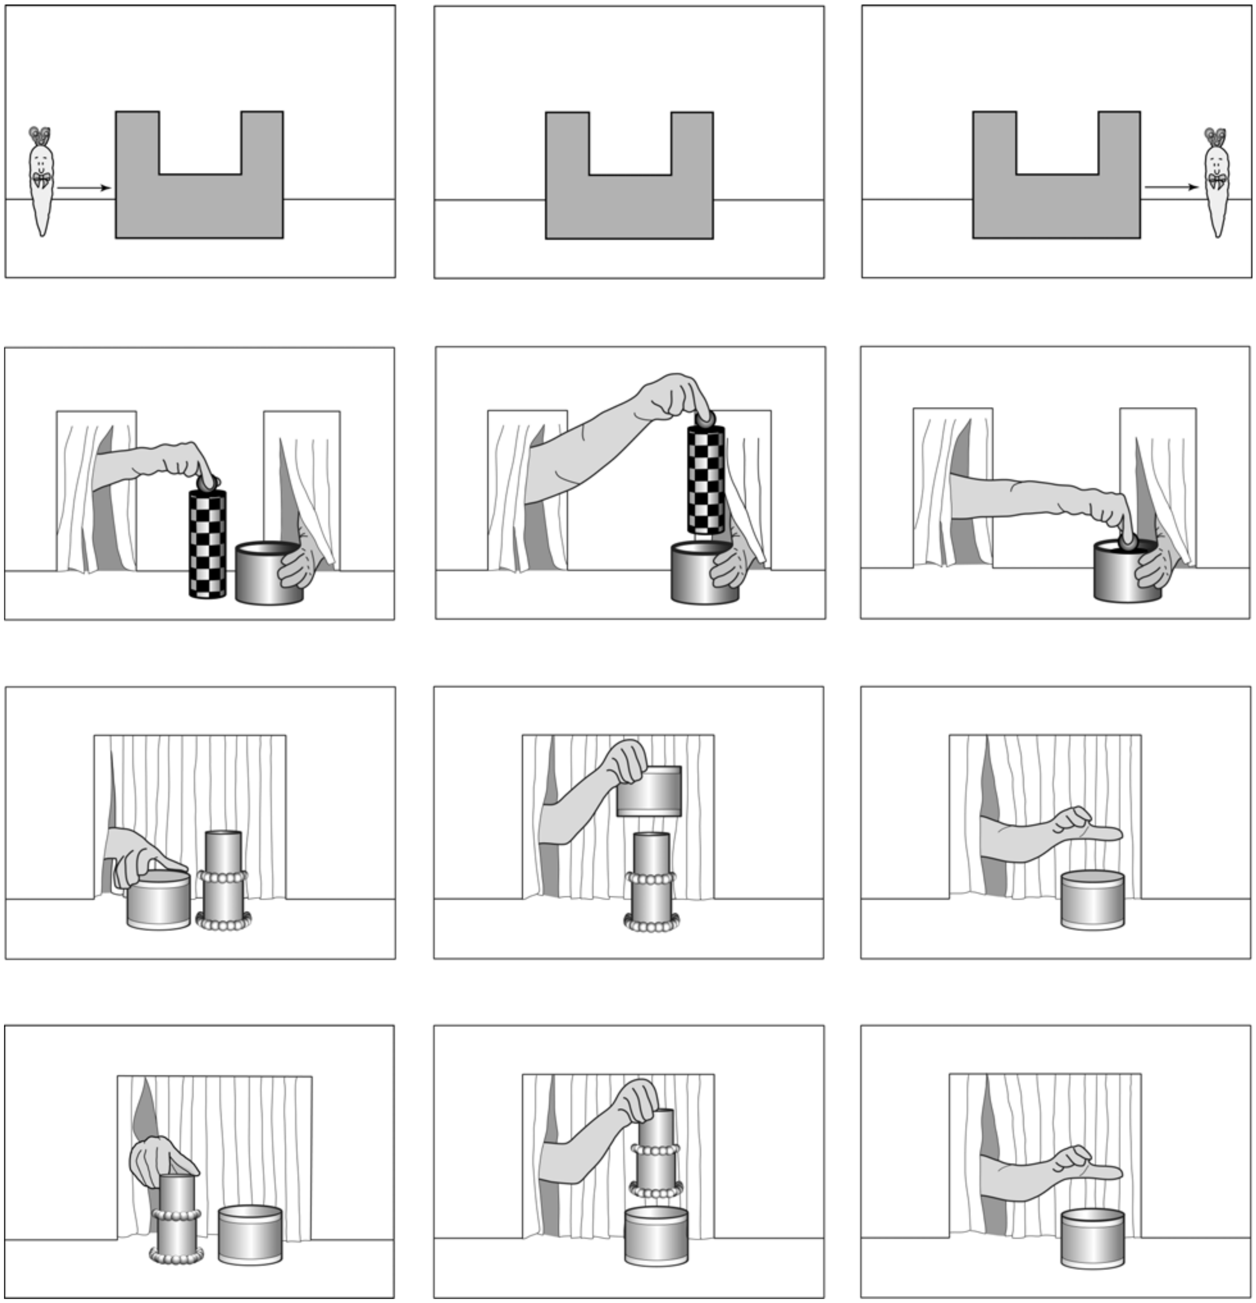

Figure 5

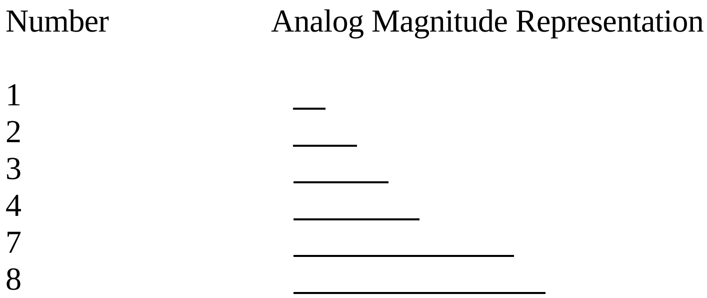

Figure 6

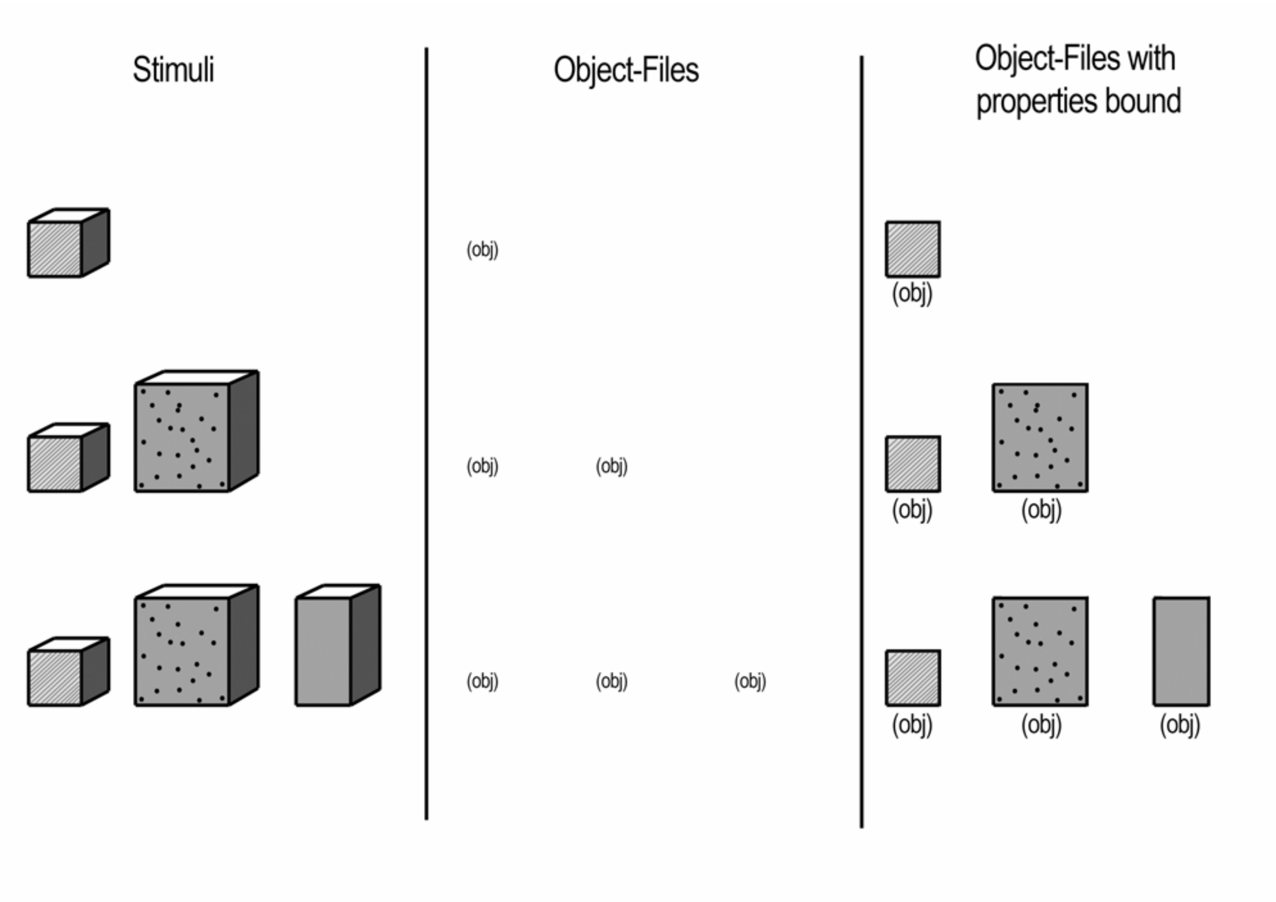
Stimuli
Object-Files
Object-Files with properties bound
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)
(obj)