

ARTICLE WITH PEER COMMENTARIES AND RESPONSE

Infants' reasoning about hidden objects: evidence for event-general and event-specific expectations

Renée Baillargeon

Department of Psychology, University of Illinois, USA

For commentaries on this article see Hood (2004), Leslie (2004) and Bremner and Mareschal (2004).

Abstract

Research over the past 20 years has revealed that even very young infants possess expectations about physical events, and that these expectations undergo significant developments during the first year of life. In this article, I first review some of this research, focusing on infants' expectations about occlusion, containment, and covering events, all of which involve hidden objects. Next, I present an account of infants' physical reasoning that integrates these various findings, and describe new experiments that test predictions from this account. Finally, because all of the research I discuss uses the violation-of-expectation method, I address recent concerns about this method and summarize new findings that help alleviate these concerns.

Introduction

As adults, we possess a great deal of knowledge about the physical world, which we use for many different purposes: for example, to predict and interpret the outcomes of physical events, to guide our actions on objects, to interpret others' actions, and even to entertain or deceive others. Over the past 20 years, my collaborators and I have been studying how infants use their developing physical knowledge to predict and interpret the outcomes of the physical events they observe.

As we all know, Piaget (1952, 1954) was the first researcher to examine the development of infants' physical knowledge. Through his observations and writings, Piaget raised many fascinating questions about infants' understanding of objects, space, time and causality. Unfortunately, Piaget did not have access to the sophisticated new methods available to us today, and so his conclusions tended to underestimate infants' physical knowledge and reasoning abilities. These new methods have yielded two general findings: (1) even very young infants possess expectations about various physical events, and (2) these expectations undergo significant developments during the first year of life (for recent reviews, see Baillargeon,

2001, 2002). In this article, I illustrate these general findings by focusing on one small portion of infants' physical knowledge, namely, infants' ability to predict and interpret the outcomes of physical events involving *hidden objects*.

Recent research suggests that infants form distinct event categories, such as containment, support and collision events. The evidence for these event categories comes from several subfields of infant cognition, including category discrimination, physical reasoning, perseveration and object individuation (e.g. Aguiar & Baillargeon, 2003; Casasola, Cohen & Chiarello, 2003; Hespos & Baillargeon, 2001a; McDonough, Choi & Mandler, 2003; Munakata, 1997; Needham & Ormsbee, 2003; Wilcox & Baillargeon, 1998a; Wilcox & Chapa, 2002; for a partial review, see Baillargeon & Wang, 2002). In this article, I focus on three event categories that involve hidden objects: *occlusion* events (which are events in which one object moves or is placed behind a nearer object, or occluder); *containment* events (which are events in which an object is placed inside a container); and *covering* events (which are events in which a rigid cover is lowered over an object).

Most of the research I will review used the violation-of-expectation (VOE) method (e.g. Baillargeon, 1998;

Address for correspondence: Renée Baillargeon, Department of Psychology, University of Illinois, 603 E. Daniel, Champaign, IL 61820, USA; e-mail: rbaillar@s.psych.uiuc.edu

Wang, Baillargeon & Brueckner, 2004). In a typical experiment, 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 infants (1) possess the expectation under investigation; (2) detect the violation in the unexpected event; and (3) are 'surprised' by this violation. The term 'surprised' is intended here simply as a short-hand descriptor, to denote a state of heightened interest or attention induced by an expectation violation. Throughout the article, I will use interchangeably the phrases 'detect a violation', 'are surprised by a violation' and 'respond with increased attention to a violation'.

The article is organized into five main sections. First, I discuss very young infants' expectations about hidden objects. Second, I explore several different ways in which these expectations develop during the first year of life. Third, I point out some apparent discrepancies between the findings discussed in the first and second sections, and outline a new account of infants' physical reasoning that attempts to make sense of these discrepancies. Fourth, I describe two lines of research that test predictions from this account. Finally, I consider recent concerns about the VOE method, and evaluate these concerns in light of the findings reviewed in the previous sections as well as additional findings.

1. In the beginning

The youngest infants tested successfully to date with the VOE method are 2.5-month-old infants. To my knowledge, there are now six reports indicating that these young infants can detect violations in occlusion, containment and covering events. Rather than discussing these experiments in detail, I simply describe the violations that the infants in these experiments succeeded in detecting.

Occlusion events (see Figure 1)

Spelke, Breinlinger, Macomber and Jacobson (1992) showed 2.5-month-old infants two barriers standing a short distance apart on the right end of a platform. A screen was lowered to hide the barriers, and then an experimenter's hand placed a ball on the left end of the platform and hit it gently so that it rolled behind the screen. Finally, the screen was raised to reveal the ball resting against the second barrier. The infants looked reliably longer at this event than at a similar, expected event, suggesting that they believed that the ball contin-

ued to exist after it became hidden, and realized that it could not roll to the second barrier when the first barrier blocked its path.

Wilcox, Nadel and Rosser (1996) showed 2.5-month-old infants a toy lion resting on one of two placemats. Next, screens hid the placemats, and an experimenter's hand entered the apparatus and retrieved the lion from behind the incorrect screen. The infants detected the violation in this event, suggesting that they believed that the lion continued to exist after it became hidden, and realized that it could not be retrieved from behind one screen when it was hidden behind the other screen.

In a series of experiments, Andrea Aguiar, Yuyan Luo and I showed 2.5-month-old infants events in which an object moved behind one of two screens separated by a gap; after a few seconds, the object reappeared from behind the other screen (Aguiar & Baillargeon, 1999; Luo & Baillargeon, in press). The same positive results were obtained whether the screens were symmetrical or asymmetrical, and whether the object was a short toy mouse or a tall cylinder. In all cases, the infants responded with increased attention, suggesting that they believed that the object continued to exist after it became hidden, and realized that it could not disappear behind one screen and reappear from behind the other screen without appearing in the gap between them.

Containment events (see Figure 2)

Sue Hespos and I found that 2.5-month-old infants could detect two different containment violations (Hespos & Baillargeon, 2001b). In one violation, an experimenter rotated a tall container forward to show the infants its closed top. Next, the experimenter placed the container upright on the apparatus floor and then lowered an object into the container through its closed top. In the other violation, an experimenter lowered an object inside a container with an open top. Next, the experimenter slid the container forward and to the side to reveal the object standing in the container's initial position. The infants looked reliably longer at these events than at similar, expected events, suggesting that they believed that the object continued to exist after it became hidden, and realized that it could not pass through the closed top or the closed walls of the container.

Covering events (see Figure 2)

Finally, Su-hua Wang, Sarah Paterson and I recently found that infants aged 2.5 to 3 months could detect two different covering violations (Wang, Baillargeon & Paterson, in press). In one violation, the infants first saw a toy duck resting on the left end of a platform.

Occlusion events




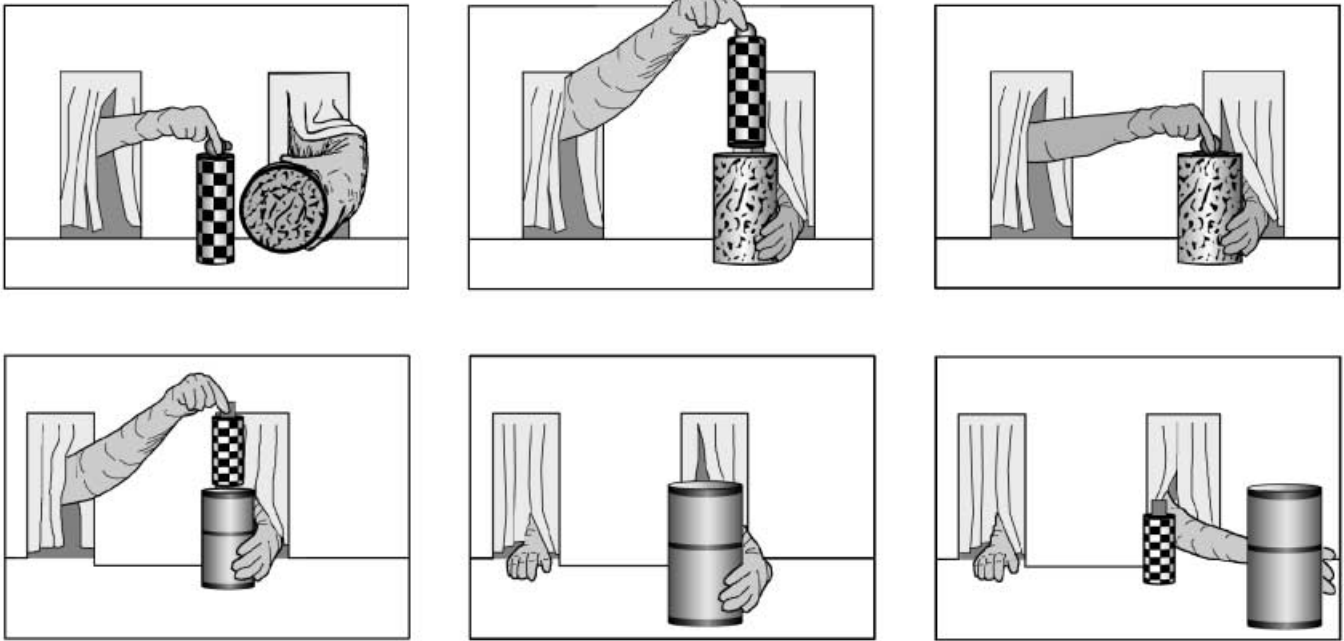
Figure 1 Occlusion violations detected by 2.5-month-old infants: row 1, Spelke et al. (1992); row 2, Wilcox et al. (1996); row 3: Aguiar and Baillargeon (1999).

Next, an experimenter's hand lowered a cover over the duck. The hand slid the cover to the right end of the platform and then lifted the cover to reveal no duck. In the other violation, the middle of the platform was hidden by a screen slightly taller than the duck. The hand lowered the cover over the duck, slid the cover behind the left half of the screen, lifted it above the screen, moved it to the right, lowered it behind the right half of the screen, slid it past the screen, and finally lifted it to reveal the duck. The infants were surprised by these violations, suggesting that they believed that the duck continued to exist after it became hidden, and expected it to move with the cover when the cover was slid but not lifted to a new location.

Conclusions

It certainly is impressive that infants as young as 2.5 months of age can detect the various occlusion, containment and covering violations I have just described. But how do they come to do so? It does not seem likely that very young infants would have repeated opportunities to observe *all* of these events, and to learn to associate each event with its outcome. A more likely possibility, I believe, is that suggested by Spelke and her colleagues (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke *et al.*, 1992; Spelke, Phillips & Woodward, 1995b): that from an early age infants interpret physical events in accord with general principles of *continuity* (objects exist continuously in time

Containment events



Covering events

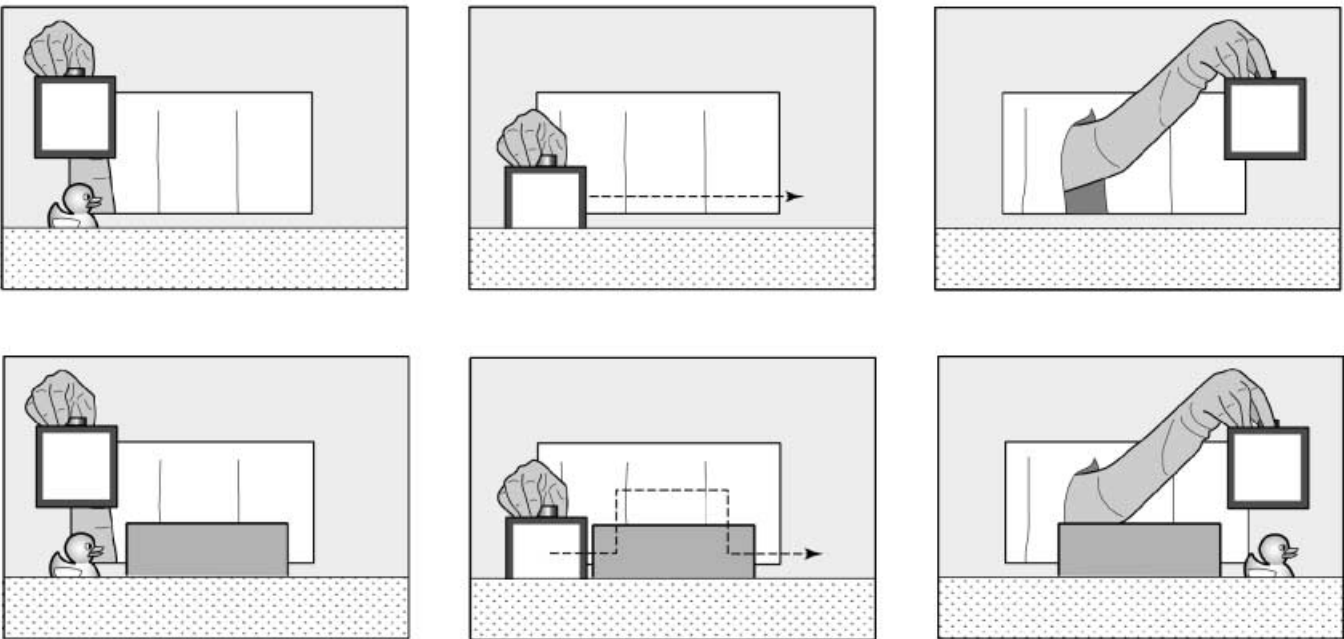


Figure 2 Top two rows: containment violations detected by 2.5-month-old infants, Hespous and Baillargeon (2001b); bottom two rows: covering violations detected by 2.5- to 3-month-old infants, Wang et al. (in press).

and space) and *solidity* (for two objects to each exist continuously, the two cannot exist at the same time in the same space). We return in Section 3 to the question of whether these principles are likely to be innate or learned.

2. Developments

The evidence that 2.5-month-old infants already possess expectations about occlusion, containment and covering events does not mean that little or no development remains to take place. In fact, research over the past 10 years has identified many different ways in which infants' expectations develop during the first year. In this section, I discuss three such developments: (a) generating explanations for occlusion violations; (b) identifying variables to better predict the outcomes of occlusion events; and (c) identifying similar variables in containment and covering events (e.g. Baillargeon & Luo, 2002).

2A. Generating explanations

We have known for many years that infants are sometimes able to generate explanations for violations involving hidden objects (e.g. Baillargeon, 1994b; Spelke & Kestenbaum, 1986; Spelke, Kestenbaum, Simons & Wein, 1995a; Xu & Carey, 1996). In a recent series of experiments, Andrea Aguiar and I explored the early development of this ability (Aguiar & Baillargeon, 2002).

In one experiment, 3- and 3.5-month-old infants were first habituated to a toy mouse moving back and forth behind a large screen; the mouse disappeared at one edge of the screen and reappeared, after an appropriate interval, at the other edge. Next, a window was created in the upper or lower half of the screen, and the mouse again moved back and forth behind the screen. In the high-window event, the mouse was shorter than the bottom of the window and did not become visible when passing behind the screen. In the low-window event, the mouse should have become visible, but it again did not appear in the window.

The 3-month-old infants looked reliably longer at the low- than at the high-window event, suggesting that they (1) believed that the mouse continued to exist after it became hidden behind the screen; (2) realized that the mouse could not disappear at one edge of the screen and reappear at the other edge without traveling the distance behind the screen; and (3) expected the mouse to become visible in the low window and were surprised that it did not. In contrast to the 3-month-olds, the 3.5-month-olds tended to look equally at the two test events. Our interpretation of this negative result was that these older infants were able to generate an explanation for the

low-window event. Upon seeing that the mouse did not appear in the low window, the infants inferred that *two* mice were involved in the event, one traveling to the left and one to the right of the screen. By positing the presence of a second mouse, the infants were able to make sense of the low-window event, which then no longer seemed surprising to them. Unlike the 3.5-month-olds, the 3-month-olds were not able to spontaneously infer that two mice were present in the apparatus; because they could not make sense of the low-window event, this event remained surprising to them throughout the test trials.

To confirm these interpretations, we conducted several additional experiments (see Figure 3). For example, in one condition 3.5-month-old infants saw the same habituation and test events as before with one exception: at the start of each trial, the screen was briefly lowered to show that only one mouse was present in the apparatus. We reasoned that the 3.5-month-olds in this condition would no longer be able to generate a two-mouse explanation for the low-window event, and they should therefore look reliably longer at this event than at the high-window event. In another condition, 3-month-old infants were shown similar events, except that two mice were revealed when the screen was lowered. We reasoned that if the 3-month-olds in this condition were able to take advantage of this two-mouse 'hint' to make sense of the low-window event, they should tend to look equally at the low- and high-window events. We thus expected the 3- and 3.5-month-old infants in this experiment to show the reverse pattern from that in our initial experiment, and that is exactly what we found: the 3.5-month-old infants, who could no longer generate a two-mouse explanation, now looked reliably longer at the low- than at the high-window event; and the 3-month-old infants, who were shown that two mice were present in the apparatus, now looked about equally at the two events.

In another experiment, 3- and 3.5-month-old infants saw events similar to those in the last experiment, with one exception: when the screen was lowered at the start of each trial, the infants could see one mouse and one small screen that was sufficiently large to hide a second mouse (see Figure 4). We reasoned that, upon seeing that the mouse did not appear in the screen's low window, the 3.5-month-olds might infer that a second mouse had been hidden behind the small screen, and hence might look about equally at the low- and high-window events. As for the 3-month-olds, since these younger infants did not seem to be able to spontaneously generate a two-mouse explanation for the low-window event, we expected that they would look reliably longer at the low- than at the high-window event. In other words, we predicted that the results of this experiment would mirror those of our initial experiment, and that is indeed

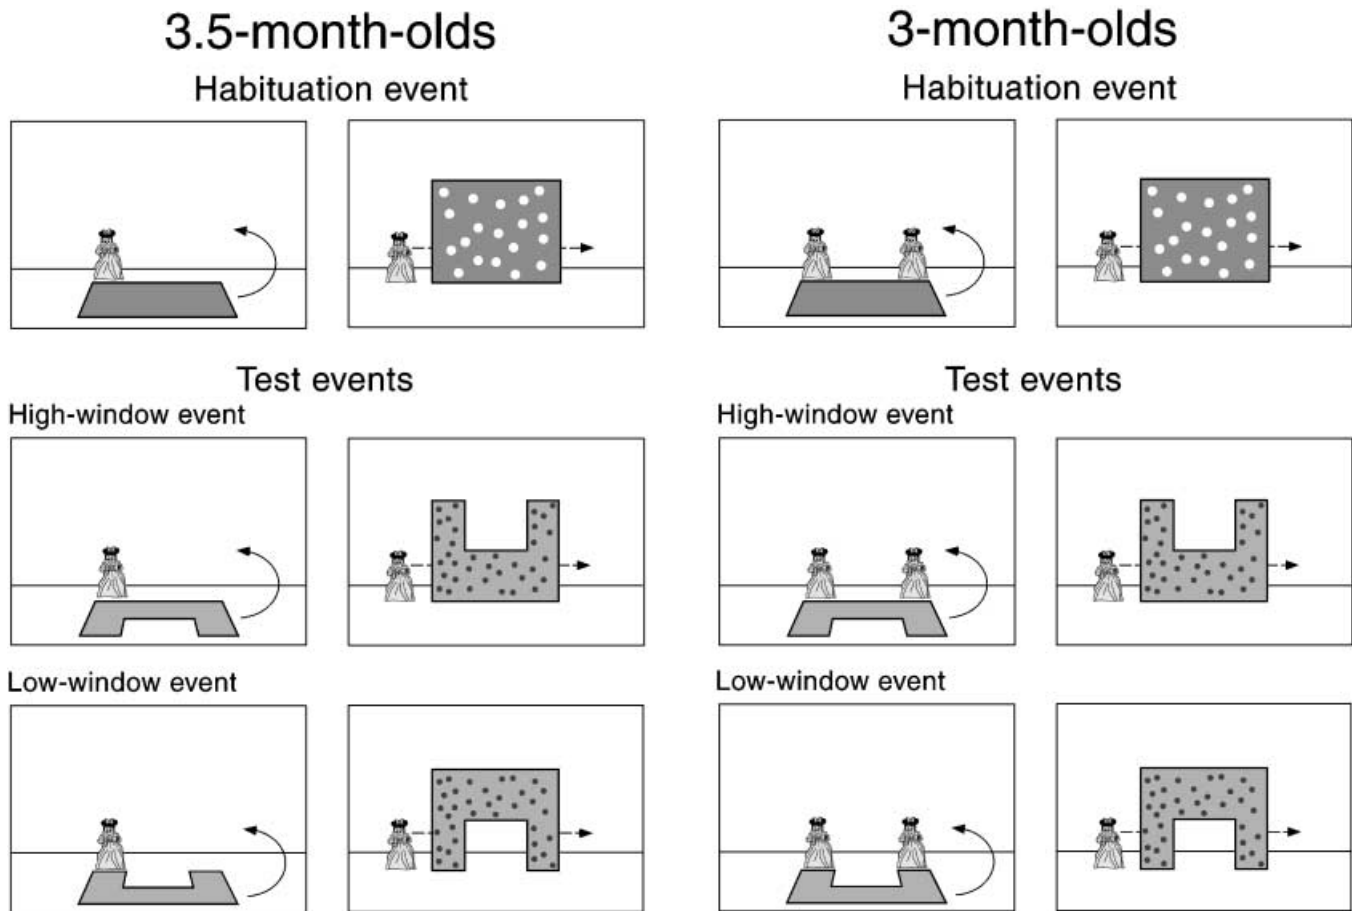


Figure 3 Habituation and test events used by Aguiar and Baillargeon (2002); the screen was lowered at the start of each trial to reveal either one mouse (3.5-month-old infants) or two mice (3-month-old infants).

what we found. The older infants, who could generate an explanation for the low-window event, tended to look equally at the events, whereas the younger infants, who could not generate such an explanation, looked reliably longer at the low- than at the high-window event.


Conclusions

The results we have just discussed support two general conclusions. First, by 3.5 months of age, infants are able to posit additional objects to make sense of at least some occlusion violations. As we will see later on, there are other, more subtle occlusion violations that 3.5- and even 5.5-month-old infants cannot explain in this way (e.g. violations in which the upper portion of an object fails to appear in a high window; see Section 5A). The range of occlusion violations infants solve by inferring the presence of an additional object behind an occluder thus increases steadily with age.

Second, infants younger than 3.5 months of age do not seem to be able to posit additional objects in occlusion events. We saw earlier that 2.5-month-old infants are surprised when an object fails to appear between two screens (Aguiar & Baillargeon, 1999); and we just saw that 3-month-old infants are surprised when an object fails to appear in a screen's low window (Aguiar & Baillargeon, 2002). Why younger infants do not spontaneously posit the presence of additional objects is an interesting question for future research. One possibility is that younger infants are less aware that many objects (such as toy mice) have duplicates, and hence are less likely to invoke such explanations. Alternatively, it may be that, when watching an event, young infants are initially limited to representing objects they directly see or have seen (e.g. when shown a toy mouse that moves across an apparatus and then disappears behind a screen, infants can represent only the mouse and screen). Inferring the presence of additional objects – going beyond


3- and 3.5-month-olds

Habituation event



Test events

High-window event



Low-window event




Figure 4 Habituation and test events used by Aguiar and Baillargeon (2002) with 3- and 3.5-month-old infants; the screen was lowered at the start of each trial to reveal one mouse and one small screen large enough to hide a second mouse.

the information given, to borrow Bruner's (1973) words – may not be possible in the first 3 months of life, and may occur only after appropriate developments have taken place. For example, it may be that in order for infants to posit objects beyond those immediately given, connections must be forged between their physical-reasoning system and a separate, problem-solving system.¹

¹ Could young infants' inability to posit the presence of an additional object behind a screen be due to a more general inability to keep track of multiple objects at the same time? We think not. Recall that the 3-month-old infants in our original mouse experiment were surprised when the mouse failed to appear in the screen's low window (Aguiar & Baillargeon, 2002). This response was eliminated when the screen was first lowered to reveal two mice, but not a mouse and a small screen. If the infants in these last experiments could keep track of three objects – two mice and a large screen, or one mouse, one small screen and one large screen – why did the infants in the original experiment fail to posit the existence of an additional mouse behind the large screen?

2B. Identifying variables in occlusion events

Research over the past 10 years has shown that, when learning about an event category such as support or collision events, infants identify a series of *variables* or rules that enable them to predict outcomes within the category more and more accurately over time (e.g. Baillargeon, Needham & DeVos, 1992; Dan, Omori & Tomiyasu, 2000; Huettel & Needham, 2000; Kotovsky & Baillargeon, 1994, 1998; Sitskoorn & Smitsman, 1995; Wang, Kaufman & Baillargeon, 2003; for reviews, see Baillargeon, 1995, 1998, 2002). Recent evidence suggests that this developmental pattern holds for occlusion events as well. Although infants realize at an early age that an object continues to exist *after* it becomes hidden behind an occluder, as we saw in Section 1, they are rather poor initially at predicting *when* an object behind an occluder should be hidden, *how soon* an object should reappear from behind an occluder, *how long* an object should take to cross a window in an occluder, and so on (e.g. Arterberry, 1997; Aguiar & Baillargeon, 1999; Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a; Lécuyer & Durand, 1998; Luo & Baillargeon, 2004a, in press; Spelke *et al.*, 1995a; Wang *et al.*, 2004; Wilcox, 1999; Wilcox & Schweinle, 2003). With experience, infants identify variables that enable them to predict all of these outcomes more accurately. Due to space limitations, I focus here on the first of these developments.

What are some of the variables infants consider to predict when an object behind an occluder should and should not be hidden (see Figure 5)? At 2.5 months of age, infants appear to use only a simple *behind/not-behind* variable: they expect an object to be hidden when behind an occluder and to be visible when not (Aguiar & Baillargeon, 1999; Lécuyer & Durand, 1998; Luo & Baillargeon, in press). Thus, when a toy mouse moves back and forth behind two screens, infants expect the mouse to be hidden when behind each screen and to be visible when between them, because at that point the mouse does not lie behind any occluder (see Section 1). However, if the screens are connected at the top or bottom, infants now view them as forming a single occluder, and they expect the mouse to remain hidden when behind this occluder. At this age, *any* object is expected to be hidden when behind *any* occluder. Infants thus detect the violation shown in the top row of Figure 5, but not those in the following rows.

At about 3 months of age, infants identify a new occlusion variable, *lower-edge-discontinuity*: they now expect an object to be hidden when behind an occluder with a continuous lower edge, but to be visible when behind an occluder with a discontinuous lower edge (Aguiar & Baillargeon, 2002; Luo & Baillargeon, in press).




Figure 5 Sequence of variables infants identify as they learn when an object behind an occluder should and should not be hidden.

Infants thus detect the violation shown in the second row of Figure 5 (see Section 2A), but not those in the following rows. It is not until infants are about 3.5 to 4 months of age that they identify *height* and *width* as occlusion variables and expect tall objects to remain partly visible when behind short occluders (Baillargeon & DeVos, 1991), and wide objects to remain partly visible when behind narrow occluders (Wang *et al.*, 2004; Wilcox, 1999; Wilcox & Baillargeon, 1998b; see Section 5A and Figure 13 for a fuller description of the width violation in Figure 5).

Finally, at about 7.5 months of age, infants identify *transparency* as an occlusion variable: when an object is

placed behind a transparent occluder, infants now expect the object to be visible through the front of the occluder and are surprised if it is not (Luo & Baillargeon, 2004a, 2004b; see Section 2C and Figure 8 for a fuller description of the transparency violation in Figure 5).²

Errors of omission and commission

The findings I have just summarized indicate that infants' knowledge of when objects behind occluders should and should not be hidden is initially very limited, and improves steadily as they identify relevant variables. This description predicts that young infants who have *not* yet identified a variable should err in two distinct ways in VOE tasks, when shown violation and non-violation events involving the variable. First, infants should respond to violation events consistent with their faulty knowledge as though they were expected. We discussed several instances of such errors above: recall, for example, that infants who have not yet identified height as an occlusion variable are not surprised when (or view as expected a violation event in which) a tall object remains fully hidden when passing behind a short occluder (Aguiar & Baillargeon, 2002; Baillargeon & DeVos, 1991; Luo & Baillargeon, in press). I will refer to this first kind of error – viewing a violation event as expected – as an error of *omission*.

Second, infants should also respond to non-violation events inconsistent with their faulty knowledge as though they were unexpected. In other words, infants should respond to perfectly ordinary and commonplace occlusion events with increased attention, when these events happen to contradict their limited knowledge. I will refer to this second kind of error – viewing a non-violation event as unexpected – as an error of *commission*.

Do young infants with a limited knowledge of occlusion events produce errors of commission as well as errors of omission in their responses to these events? Yuyan Luo

² Readers may wonder why the variable transparency is such a late acquisition. Recent work by Johnson and Aslin (2000) suggests that infants only begin to detect clear, transparent surfaces at about 7 months of age, as a result of developments in their contrast sensitivity, which may in turn be tied to the maturation of the magnocellular system. At this stage, infants do not realize that an object should be visible when behind a transparent occluder (Luo & Baillargeon, 2004a). They have not yet identified transparency as an occlusion variable, and only take into account the variables lower-edge-discontinuity, height and width when reasoning about occlusion events. Thus, when an object is placed behind a transparent occluder that has no openings and is taller and wider than the object, infants expect the object to be hidden and are surprised if it is not (Luo & Baillargeon, 2004a). By 7.5 months of age, infants have identified transparency as an additional occlusion variable, and they now expect an object behind a transparent occluder to be visible through the front of the occluder (Luo & Baillargeon, 2004a, 2004b).




Figure 6 Familiarization and test events used by Luo and Baillargeon (in press).

and I recently conducted a series of experiments that addressed this question (Luo & Baillargeon, in press).

In one experiment, 3-month-old infants were first familiarized with a cylinder that moved back and forth behind a screen; the cylinder was as tall as the screen (see Figure 6). Next, a large portion of the screen's mid-section was removed to create a large opening; a short strip remained above the opening in the discontinuous-lower-edge test event, and below the opening in the continuous-lower-edge test event. For half of the infants, the cylinder did not appear in the opening in either event (CDNA condition); for the other infants, the cylinder appeared (CA condition).

The infants in the CDNA condition were shown two violation test events. However, because at 3 months infants have identified lower-edge-discontinuity but not height as an occlusion variable, we predicted that the infants would view only one of these violation events as unexpected. Specifically, the infants should view the event in which the cylinder failed to appear behind the screen with a discontinuous lower edge as unexpected (a

correct response), but they should view the event in which the cylinder failed to appear behind the screen with a continuous lower edge as expected (an error of omission). The infants should therefore look reliably longer at the discontinuous- than at the continuous-lower-edge event.

Unlike the infants in the CDNA condition, those in the CA condition were shown two non-violation test events. Again, because 3-month-old infants have identified lower-edge-discontinuity but not height as an occlusion variable, we predicted that the infants would view only one of those events as expected. Specifically, the infants should view the event in which the cylinder appeared behind the screen with a discontinuous lower edge as expected (a correct response), but they should view the event in which the cylinder appeared behind the screen with a continuous lower edge as unexpected (an error of commission). The infants should therefore look reliably longer at the continuous- than at the discontinuous-lower-edge event.

The results supported our predictions: the infants in the CDNA condition looked reliably longer at the discontinuous- than at the continuous-lower-edge event, and those in the CA condition showed the opposite looking pattern. Their limited knowledge of occlusion thus (1) led the infants in the CDNA condition to view one of the violation events they were shown as expected (an error of omission), and (2) led the infants in the CA condition to view one of the non-violation events they were shown as unexpected (an error of commission). To put it differently, the infants both failed to detect a violation where there was one, and perceived a violation where there was none.

Conclusions

The evidence reviewed in this section suggests two broad conclusions. First, infants identify a series of variables that enables them to predict the outcomes of occlusion events more and more accurately over time. Second, when infants' knowledge of occlusion is still limited, they err in two distinct ways in their responses to occlusion events, by viewing violation events consistent with their faulty knowledge as non-violations, and by viewing non-violation events inconsistent with their faulty knowledge as violations. Surprise, like beauty, clearly lies in the eye of the beholder.

2C. Identifying similar variables in containment and covering events

We saw in the last section that infants identify a series of variables that enables them to predict the outcomes

of occlusion events more and more accurately over time. Exactly the same developmental pattern has been observed for infants' reasoning about containment and covering events (e.g. Aguiar & Baillargeon, 1998; Hespos & Baillargeon, 2001a, 2001b; Leslie, 1995; Luo & Baillargeon, 2004b; McCall, 2001; Sitskoorn & Smitsman, 1995; Spelke & Hespos, 2002; Wang *et al.*, 2004, in press).

Given that in many cases the same variables affect the outcomes of occlusion, containment and covering events, one might ask whether infants generalize variables identified in one event category to the other categories. For example, the variables height and transparency are equally relevant to occlusion, containment and covering events. When infants have acquired these variables in one category, do they immediately generalize them to the other categories? Recent research from our laboratory suggests that they do not: variables identified in one event category appear to remain tied to that category – they are not generalized to other relevant categories (e.g. Hespos & Baillargeon, 2001a; Luo & Baillargeon, 2004a, 2004b; Onishi, 2000; Wang *et al.*, in press).

To illustrate this point, I will first describe experiments Sue Hespos and I conducted to compare 4.5-month-old infants' ability to reason about height information in containment and in occlusion events (Hespos & Baillargeon, 2001a). The infants were assigned to a containment or an occlusion condition (see Figure 7). The infants in the *containment* condition saw two test events. At the start of each event, an experimenter's gloved hand grasped a knob at the top of a tall cylindrical object; next to the object was a container. The hand lifted the object and lowered it inside the container until only the knob remained visible above the rim. In the tall-container event, the container was as tall as the cylindrical portion of the object; in the short-container event, the container was only half as tall, so that it should have been impossible for the cylindrical portion of the object to become fully hidden inside the container. Prior to the test trials, the infants received familiarization trials in which the containers were rotated forward so that the infants could inspect them. The infants in the *occlusion* condition saw similar familiarization and test events, except that the bottom and back half of each container were removed to create a rounded occluder.

Because height is identified at about 3.5 months of age as an occlusion variable (Baillargeon & DeVos, 1991; see Section 2B), we expected that the infants in the occlusion condition would look reliably longer at the short- than at the tall-occluder test event, and this is precisely what we found. In marked contrast, the infants




Figure 7 Test events used by Hespos and Baillargeon (2001a) in the containment, occlusion and container-as-occluder conditions.

in the containment condition tended to look equally at the short- and tall-container test events. Our interpretation of this negative result was that at 4.5 months of age infants have not yet identified the variable height in containment events: they do not yet realize that a tall object cannot become fully hidden inside a short

container.³ This interpretation led to a striking prediction: infants shown the same test events as in the containment condition but with the object lowered *behind* rather than *inside* each container should be able to detect the violation in the short-container event. In this condition, the containers served simply as occluders, so the infants' performance should mirror that of the infants in the occlusion condition. The results confirmed this prediction: when the object was lowered behind rather than inside the containers, the infants looked reliably longer at the short- than at the tall-container event.

In a subsequent experiment, 5.5-, 6.5- and 7.5-month-old infants were tested with the container condition test events. Only the 7.5-month-old infants detected the violation in the short-container event, suggesting that it is not until infants are about 7.5 months of age that they identify the variable height in containment events.

These results (and control results obtained with a shorter object) suggested two conclusions. First, infants do not generalize variables from occlusion to containment events: they learn separately about each event category. Second, because several months separate the acquisition of the variable height in these two categories, a striking lag or *décalage* (to use a Piagetian term) can be observed in infants' responses to similar events from the categories. We return in Section 5A to the question of why infants might identify the variable height later in containment than in occlusion events.

Additional *décalages*

Due to space limitations, I will describe only briefly two other *décalages* we have recently uncovered (see Figure 8). The first comes from experiments in which Su-hua Wang, Sarah Paterson and I compared 9- to 12-month-old infants' reasoning about the variable height in containment and in covering events (Wang *et al.*, in press). Consistent with the results just described, we found that the 9-month-old infants responded with increased attention to a violation event in which a tall object was lowered inside a short container until it

³ There were of course other possible interpretations for the negative result of the containment condition. For example, it might be suggested that 4.5-month-old infants generally have more difficulty reasoning about containment than occlusion events, or must devote more computational resources to representing containment than occlusion events, and so are less likely to detect containment than occlusion violations. Two sets of findings argued against such interpretations. First, as we saw in Section 1, even 2.5-month-old infants are able to detect violations in containment events (Hespos & Baillargeon, 2001b). Second, as we will see in Section 5A, 4-month-old infants detect width (as opposed to height) violations in both occlusion and containment events (Wang *et al.*, 2004).

became fully hidden. However, it was not until infants were 12 months of age that they responded with increased attention to a similar violation event in which a short cover (the short container turned upside down) was lowered over the tall object until it became fully hidden.

The other *décalage* comes from experiments in which Yuyan Luo and I examined 7.5- to 9.5-month-old infants' reasoning about the variable transparency in occlusion and in containment events (Luo & Baillargeon, 2004a, 2004b). We found that the 7.5-month-old infants responded with increased attention when shown the following occlusion violation. To start, a checkered object stood next to a transparent occluder; the edges of the occluder were outlined with red tape so that they were easily detectable. Next, a screen hid the occluder, and an experimenter's gloved hand grasped the object and lowered it behind the transparent occluder. Finally, the screen was lowered to reveal the transparent occluder with no object visible behind it. Although the 7.5-month-old infants detected this violation, only the 9.5-month-old infants detected a similar violation in which the transparent occluder was replaced with a transparent container.

Conclusions

The research discussed in this section suggests that infants do not generalize variables from occlusion to containment or covering events: they learn separately about each category. When several weeks or months separate the identification of the same variable in these different categories, striking *décalages* arise in infants' responses to similar events from the categories.

3. A new account of infants' physical reasoning


In Section 1, I reviewed evidence that infants as young as 2.5 months detect some violations in occlusion, containment and covering events; and I suggested, following Spelke and her colleagues (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke *et al.*, 1992, 1995b), that from an early age infants interpret physical events in accord with general principles of continuity and solidity. Some of the evidence reviewed in Section 2 may at first seem inconsistent with the notion that infants possess general continuity and solidity principles, for two reasons; these two discrepancies are discussed in turn.

Event-general principles and event-specific expectations


We saw in Section 2 that the expectations infants acquire about physical events are not event-general principles

Height in containment and covering events

9 months




12 months



Transparency in occlusion and containment events

7.5 months



9.5 months




Figure 8 Top two rows: *décalage* in infants' reasoning about height in containment and covering events (Wang et al., in press); bottom two rows: *décalage* in infants' reasoning about transparency in occlusion and containment events (Luo & Baillargeon, 2004a, 2004b).

that are applied broadly to all relevant events, but rather *event-specific* expectations. Infants do not acquire general principles of height or transparency: they identify these variables separately in each event category. For example, infants identify the variable height at about 3.5 months in occlusion events, at about 7.5 months in containment events, and at about 12 months in covering events (Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a; Wang *et al.*, in press). But if infants are capable of acquiring only event-specific expectations,

how could they possess event-general principles of continuity and solidity, and as early as 2.5 months of age?

One possibility is that infants' learning mechanism is initially geared toward acquiring event-general expectations, but soon evolves into a different mechanism capable of acquiring only event-specific expectations. Another possibility, which I think more likely, is that infants' general principles of continuity and solidity are innate (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke *et al.*, 1992, 1995b).

Successes and failures in detecting continuity and solidity violations

Whether one chooses the first or second possibility above, difficulties remain. If infants interpret physical events in accord with general principles of continuity and solidity (whether learned or innate), one might expect them to detect *all* salient violations of these principles. However, we have seen that although some continuity and solidity violations are detected as early as 2.5 months, others are not detected until much later: recall, for example, that infants younger than 7.5 months are not surprised when a tall object becomes hidden inside a short container (Hespos & Baillargeon, 2001a); that infants younger than 9.5 months are not surprised when an object placed inside a transparent container is not visible through the front of the container (Luo & Baillargeon, 2004a); and that infants younger than 12 months are not surprised when a short cover is lowered over a tall object until it becomes fully hidden (Wang *et al.*, in press).

How can we make sense of the fact that infants detect some continuity and solidity violations at a very young age, and others only much later? To address this question, Su-hua Wang and I have been developing a new account of infants' physical reasoning (e.g. Baillargeon, 2002; Wang *et al.*, in press). This account rests on four assumptions (see Figure 9). *First*, when watching a physical event, infants build a specialized physical representation of the event that is used to predict and interpret its outcome. *Second*, all of the information, but only the information, infants include in their physical representation of an event becomes subject to their general principles of continuity and solidity (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke *et al.*, 1992, 1995b).⁴

Third, in the first weeks of life, infants' physical representation of an event tends to be rather impoverished and includes only *basic* spatial and temporal information about the event (e.g. Kestenbaum, Termine & Spelke, 1987; Leslie, 1994; Needham, 2000; Slater, 1995;




Figure 9 Schematic representation of the reasoning account (Baillargeon, 2002; Wang *et al.*, in press).

Spelke, 1982; Yonas & Granrud, 1984). For example, when watching a containment event, infants represent that an object is being lowered inside a container. This basic information captures the essence of the event, but leaves out most of its details: whether the container is wider or taller than the object, whether it is transparent or opaque, and so on.

Fourth, as infants form event categories and learn what variables to consider in each category, they include more and more of this detailed information, or *variable* information, in their physical representations (e.g. Baillargeon, 1991; Dan *et al.*, 2000; Kotovsky & Baillargeon, 1998; Sitskoorn & Smitsman, 1995; Wang *et al.*, 2003; Wilcox, 1999). When watching an event, infants first represent the basic information about the event and use this information to categorize it. Infants then access their knowledge of the event category selected; this knowledge specifies the variables that have been identified as relevant to the category and hence that should be included in the physical representation. Variables not yet identified are typically not included in the representation. Going back to our example, infants who have identified height as a containment variable would include information about the relative heights of the object and container in their representation of the event; this information would then become subject to their general principles of continuity and solidity, making it possible for them to detect violations involving tall objects and short containers. In contrast, infants who have not yet identified height as a containment variable would include *no* height information in their physical representation of the event; as a result, this information would not be available and hence could not be interpreted in accord with their continuity and solidity principles.

How does the reasoning account explain the fact that some continuity and solidity violations are detected at an early age and others only much later? According to the account, very young infants should *succeed* in detecting any continuity and solidity violation, in any event category, as long as this violation involves only the basic spatial and temporal information they can represent.

⁴ Leslie (1994, 1995) has suggested that, from birth, infants interpret physical events in accord with a primitive notion of *force*. When watching an object push another object, for example, infants represent a force – like a directional arrow – being exerted by the first object onto the second one. In Leslie's (1994) own words, infants' physical-reasoning system 'takes, as input, descriptions that make explicit the geometry of the objects contained in a scene, their arrangements and their motions, and onto such descriptions paints the mechanical properties of the scenario' (p. 128). In a similar vein, one might suggest that the principles of continuity and solidity bestow continued existence upon objects: all other things being equal, objects in physical representations are expected to persist through time and space. Although many of the events described in the present article involved forces (e.g. hands lifting, lowering or sliding objects), we focus here only on those aspects of the events that concerned the continuity and solidity principles.

Furthermore, much older infants should *fail* to detect a continuity and solidity violation in an event category, when this violation involves a variable they have not yet identified as relevant to the category and hence do not typically include in their physical representations of events from the category.

Conclusions

I began Section 3 with two discrepancies. First, how can infants possess event-general principles of continuity and solidity, and as early as 2.5 months of age, if they acquire only event-specific expectations? Second, if infants possess such principles, why do they detect some but not other violations of the principles? To make sense of these discrepancies, I suggested that infants' principles of continuity and solidity are innate (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke *et al.*, 1992, 1995b); that these principles can only be applied to the information infants include in their physical representations of events; and that this information is initially limited but becomes richer as infants learn what variables to consider in each event category and begin to include information about these variables in their physical representations.

4. Two tests of the account

The reasoning account described in the last section makes a number of interesting predictions. For example, it predicts that infants younger than 2.5 months should succeed at detecting continuity and solidity violations that involve only the basic information they can represent. To examine this prediction, we are currently setting up new experiments to test 6- to 8-week-old infants. The reasoning account also suggests that infants should succeed at detecting violations involving variables they have not yet identified if primed through contextual manipulations to include information about these variables in their physical representations. There already exists evidence consistent with this prediction (e.g. Kotovsky, Mangione & Baillargeon, cited in Baillargeon, 1995; Wang & Baillargeon, 2004b; Wilcox & Chapa, 2004). Finally, the reasoning account also predicts change-blindness and teaching effects, which are described next.

Change-blindness effects

According to the reasoning account, infants who have not yet identified a variable as relevant to an event category typically do not include information about this variable when representing events from the category. Suhua Wang and I reasoned that if infants do not include

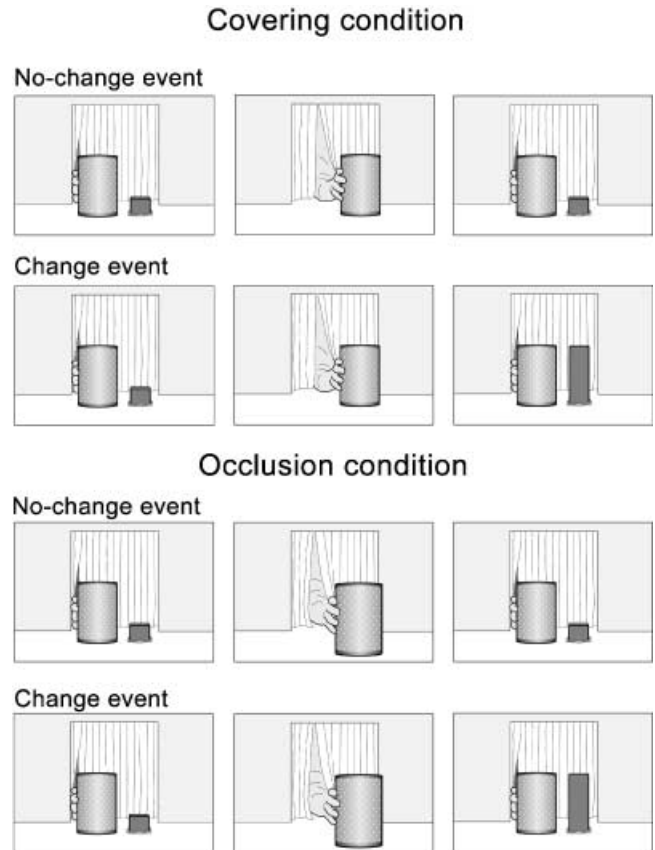


Figure 10 Test events used by Wang and Baillargeon (2004a) in the covering and occlusion conditions.

information about a variable in their physical representation of an event, then they should be unable to detect surreptitious changes involving the variable: in other words, they should be *blind* to such changes (Wang & Baillargeon, 2004a; see also related research on change detection and change-blindness in the adult literature, e.g. Simons, 1996; Simons, Franconeri & Reimer, 2000).

In our first experiment, 11-month-old infants were assigned to a covering or an occlusion condition (see Figure 10). The infants in the *covering* condition saw two test events: a no-change and a change event. At the start of each event, a tall cover stood next to a short object on an apparatus floor. An experimenter's gloved hand lifted the cover and lowered it over the object. After a pause, the hand returned the cover to the apparatus floor. In the no-change event, the object was the same as before when the cover was removed. In the change event, the object was now as tall as the cover. (Note that both the short and the tall object could fit under the cover; the experiment tested not whether infants could judge what object fit under what cover, but whether they could detect a surreptitious change in the height of an object

under a cover.) The infants in the *occlusion* condition saw similar test events except that the cover was lowered in front of, rather than over, the object.

We saw earlier that infants identify the variable height at about 3.5 months in occlusion events (Baillargeon & DeVos, 1991), but only at about 12 months in covering events (Wang *et al.*, in press). Based on these findings, we predicted that the infants in the occlusion condition would include information about the relative heights of the cover and object in their physical representations of the events, and hence would detect the violation in the change event. Conversely, we expected that the infants in the covering condition would include no height information in their physical representations, and hence would fail to detect the violation in the change event.

The results confirmed these predictions: the infants in the occlusion condition looked reliably longer at the change than at the no-change event, whereas those in the covering condition tended to look equally at the two events. In a subsequent experiment, older, 12-month-old infants were tested in the covering condition; as expected, these infants detected the violation in the change event. Thus, as predicted by the reasoning account, only the 11-month-old infants in the covering condition were blind to the surreptitious change in the height of the object.

Teaching effects

According to the reasoning account, infants who have not identified a variable as relevant to an event category typically do not include information about this variable when representing events from the category. My collaborators and I reasoned that if infants could be *taught* a new variable, then they would include information about this variable in their physical representations, which would allow them to detect continuity and solidity violations involving the variable earlier than they would otherwise (e.g. Baillargeon, Fisher & DeJong, 2000; Wang & Baillargeon, 2004c). To illustrate, Su-hua Wang and I recently attempted to teach 9.5-month-old infants the variable height in covering events (Wang & Baillargeon, 2004c); recall that this variable is typically not identified until about 12 months of age (Wang *et al.*, in press).

What might be the key ingredients in a successful teaching experiment? According to a recent account, the process by which infants typically identify a new variable in an event category is one of explanation-based learning (EBL) and involves three main steps (Baillargeon, 2002; for a computational description of EBL in the machine learning literature, see DeJong, 1993, 1997). First, infants notice *contrastive outcomes* for the variable (e.g. in the case of the variable height in covering events, infants notice that when a cover is placed over an object, the

object is sometimes fully and sometimes only partly hidden). Second, infants search for the *conditions* that map onto these outcomes (e.g. infants notice that the object becomes fully hidden when the cover is as tall as or taller than the object, and partly hidden when the cover is shorter than the object). Third, infants build an *explanation* for these condition-outcome data using their prior knowledge, including their event-general knowledge (e.g. the continuity and solidity principles specify that a tall object can extend to its full height inside a tall but not a short cover). According to the EBL account, only condition-outcome observations for which infants can build causal explanations are identified as new variables. These explanations are no doubt shallow (e.g. Keil, 1995; Wilson & Keil, 2000), and they may even be incorrect (e.g. Baillargeon, 2002), but they still serve to integrate new variables with infants' prior causal knowledge.

The EBL account not only describes how infants identify a new variable in an event category: it also suggests how one might go about teaching infants such a variable. That is, it specifies what ingredients might be essential for a successful teaching recipe.

The infants in our experiment first received three pairs of teaching trials (see Figure 11). Each pair consisted of a tall- and a short-cover event. In both events, an experimenter's gloved hand first rotated the cover forward to show its hollow interior. The hand then placed the cover upright next to a tall object, so that the infants could compare their heights. Finally, the hand lifted the cover and lowered it over the object. In the tall-cover event, the cover was taller than the object, so the object became fully hidden; in the short-cover event, the cover was shorter than the object, so only the top portion of the object became hidden. The tall and short covers in each teaching pair differed only in height; the cover not in use in a given pair was placed against the apparatus's back wall. The second and third teaching pairs were identical to the first except that the covers differed in pattern and color.⁵

According to the EBL account, the teaching trials provided the infants with the necessary information to identify the variable height in covering events: (1) the infants saw *contrastive outcomes*, in that the object was sometimes fully and sometimes only partly hidden; (2)

⁵ In previous experiments (e.g. Baillargeon, 1998; Baillargeon *et al.*, 2000), we attempted to teach 11-month-old infants the variable proportional distribution in support events (an asymmetrical object is stable when released on a platform as long as half or more of the whole object is supported). We found that infants learned this variable as long as a different asymmetrical box was used in each of the three pairs of teaching trials. In order to acquire a variable, infants apparently need to see multiple objects behave in a manner consistent with the variable.




Figure 11 Teaching and test events used by Wang and Baillargeon (2004c).

the infants could gather appropriate *condition* data to map onto these contrastive outcomes: because each cover was placed next to the object, the infants could easily compare their heights and note that the object became fully hidden when the cover was taller but not shorter than the object; and finally (3) the infants could bring to bear their continuity and solidity principles to build an *explanation* for these condition-outcome data – to make sense of the fact that the object could extend to its full height inside the tall but not the short covers.

Following the three pairs of teaching trials, the infants saw tall- and short-cover events involving a novel object and novel tall and short covers. In both events, the object became fully hidden. The infants looked reliably longer at the short- than at the tall-cover event. The same positive result was also obtained in a subsequent experiment in which a 24-hour delay separated the teaching and test trials. Thus, after being exposed to the teaching events, 9.5-month-old infants succeeded in

detecting the violation in the short-cover event, 2.5 months before they would normally have done so.

In subsequent experiments, we began examining some of the assumptions behind our teaching trials. In particular, was it important that the infants be exposed to contrastive outcomes? That they be able to gather condition data about the relative heights of the object and covers? And that they be able to build an explanation for these condition-outcome data? How essential were these three ingredients? To answer these questions, we conducted three experiments identical to our original experiment, except that the teaching trials were modified: in each experiment, one key ingredient was removed. In each case, our prediction was that infants would now fail to identify the variable height during the teaching trials, and hence would fail to detect the violation in the short-cover event during the test trials. Negative results were thus expected in all three experiments (see Figure 12).

To address the issue of *contrastive outcomes*, infants were shown teaching events in which the tall object was replaced with a very short object that became fully hidden under the tall and short covers; the infants thus no longer saw contrastive outcomes that could trigger learning. To address the issue of *condition data*, infants watched teaching events in which the cover was never placed next to the tall object on the apparatus floor; instead, the cover was held next to and above the object, making it difficult for the infants to compare their heights. Finally, to address the issue of *explanation*, infants were presented with teaching events identical to those in our original experiment with one exception: false bottoms inside the covers – revealed when the covers were rotated forward – rendered them all equally shallow; the infants thus could no longer make sense of the fact that the tall object became fully hidden under the tall covers.


The results confirmed our predictions: unlike the infants in our original teaching experiment, those in the short-object, no-height-comparison and shallow-cover experiments tended to look equally at the short- and tall-cover test events, suggesting that they had not been able to identify the variable height during the teaching trials and hence could not detect the violation in the short-cover event during the test trials.

Conclusions

According to the reasoning account presented in Section 3, infants who have not identified a variable in an event category typically do not include information about this variable when representing events from the category; as a result, this information is not available and hence cannot be subject to infants' continuity and solidity principles. Consistent with this account, we saw in Section 4 that

Short-object condition

Tall-cover event




Short-cover event



No-height-comparison condition

Tall-cover event



Short-cover event



Shallow-cover condition

Tall-cover event



Short-cover event



Figure 12 Teaching events used by Wang and Baillargeon (2004c) in the short-object, no-height-comparison and shallow-cover experiments.

infants who have not yet identified a variable are blind to surreptitious changes involving the variable; and that infants who are taught the variable can then detect violations involving it, both immediately and after a 24-hour delay.

5. About the VOE method

All of the research I have described to this point made use of the VOE method. This research not only shed light on the development of infants' expectations about occlusion, containment and covering events, but also illustrated the remarkable flexibility of the VOE method. In particular, we saw that the method can be used with infants from a wide age range, that it can be used with sundry physical events, and that it can be used to address many different questions: whether infants can detect a violation, whether they can generate an explanation for a violation, whether they can be taught to detect a violation, and so on.

Over the past few years, a number of researchers have expressed concerns about the VOE method (e.g. Bogartz, Shinsky & Schilling, 2000; Bogartz, Shinsky & Speaker, 1997; Cashon & Cohen, 2000; Haith, 1998, 1999; Haith & Benson, 1998; Munakata, 2001; Munakata, McClelland, Johnson & Siegler, 1997; Schilling, 2000; Thelen & Smith, 1994). In this final section of the article, I would like to address two such concerns; both focused on findings from VOE tasks indicating that young infants can represent hidden objects (see also Aslin, 2000; Baillargeon, 1999, 2000; Lécuyer, 2001; Munakata, 2000; Wang *et al.*, 2004).

5A. Transient-preference accounts

Several researchers have suggested that young infants may look reliably longer at the unexpected than at the expected events in VOE tasks involving hidden objects, not because they possess expectations about such objects, but because the habituation or familiarization trials that are typically included in these tasks induce in them transient and superficial preferences for the unexpected events (e.g. Bogartz *et al.*, 1997, 2000; Cashon & Cohen, 2000; Schilling, 2000; Thelen & Smith, 1994). There are two main ways of addressing these *transient-preference* accounts, and I describe them in turn (for fuller discussion, see Wang *et al.*, 2004).

Using test-only VOE tasks

Transient-preference accounts predict that young infants should fail to give evidence that they can represent hidden objects in VOE tasks with *no* habituation or familiarization trials. Without such trials, infants should have no opportunity to form transient novelty or familiarity preferences that could contribute to their responses in the test trials, and they should therefore tend to look equally at the unexpected and expected events.

With the single exception of the teaching experiments discussed in the last section (Wang & Baillargeon, 2004c), *all* of the experiments described in this article that involved infants aged 7.5 months and older were conducted with




Figure 13 Test events used by Wang *et al.* (2004) in the experimental and control conditions.

test trials only: the infants received no habituation or familiarization trials, so that any positive result obtained in these experiments could not be attributed to transient novelty or familiarity preferences formed during such trials.

To find out whether younger infants would also succeed in a test-only VOE task, Su-hua Wang, Laura Brueckner and I recently conducted an experiment with 4-month-old infants (Wang *et al.*, 2004). This experiment focused on the variable width in occlusion events: would infants realize that a wide object can become fully hidden behind a wide but not a narrow occluder? The infants were assigned to an experimental or a control condition (see Figure 13). The infants in the *experimental* condition saw two test events. At the start of each event, an experimenter's gloved hand held a wide object slightly above and behind a wooden occluder (to facilitate width comparisons). Next, a screen was raised to hide the occluder, and the hand lowered the object to the apparatus floor behind the occluder. Finally, the screen was lowered to reveal the occluder standing alone on the apparatus floor. In the wide-occluder event, the occluder was wider than the object and so could fully hide it. In the narrow-occluder event, the occluder was much narrower

than the object and so should not have been able to fully hide it. The infants in the *control* condition saw similar events except that the object was much narrower and could be hidden behind either the wide or the narrow occluder.

The infants in the experimental condition looked reliably longer at the narrow- than at the wide-occluder event, whereas those in the control condition tended to look equally at the two events. These results suggested that the infants (1) believed that the wide or narrow object continued to exist after it became hidden; (2) recognized that the narrow object could be fully hidden behind either occluder, and that the wide object could be fully hidden behind the wide but not the narrow occluder; and (3) were surprised when this last expectation was violated. Similar results were obtained in a second experiment in which the wide and narrow occluders were replaced with wide and narrow containers.⁶

Thus, infants as young as 4 months of age give evidence that they can represent a hidden object even when tested in a VOE task with no habituation or familiarization trials, only test trials.

⁶ These results suggest that, although there is a *décalage* in infants' reasoning about the variable *height* in occlusion and containment events (see Section 2C.; Hespos & Baillargeon, 2001a), there is little or no *décalage* in their reasoning about the variable *width* in these events. Why is that? According to the EBL account presented in Section 4, in order to identify a variable in an event category, infants must (1) notice contrastive outcomes for the variable; (2) find the conditions that map onto these outcomes; and (3) build an explanation for these condition-outcome data (e.g. Baillargeon, 2002). Infants' difficulty in identifying the variable height in containment as opposed to occlusion events might involve the second step. Prior research (e.g. Baillargeon, 1991, 1994a, 1995) indicates that when infants begin to reason about a continuous variable in an event category, they can reason about the variable qualitatively but not quantitatively: they are not able at first to encode and remember absolute amounts. In order to encode the heights of objects and occluders or containers qualitatively, infants must compare them as they stand *side by side*. Infants may have more opportunities to perform such comparisons with occlusion than with containment events. In the case of occlusion events, infants will often see objects move behind the side edges of occluders, making it easy to compare their heights as they stand next to each other (e.g. when a cereal box is pushed in front of a bowl). In the case of containment events, however, there may be relatively few instances in which objects are placed first next to and then inside containers; caretakers will more often lower objects directly into containers, giving infants no opportunity to compare their heights (e.g. Baillargeon, 2002; Hespos & Baillargeon, 2001a; Wang *et al.*, 2004).

The preceding reasoning predicts that, in containment events, infants should identify the variable width before the variable height, because each time an object is lowered inside a container (e.g. when a spoon is lowered into a jar), their widths can be compared qualitatively as one stands above the other. Furthermore, there should be little or no *décalage* in the identification of the variable width in occlusion and containment events: infants should be able to gather the needed qualitative width information as objects are lowered behind occluders or inside containers. The responses of the 4-month-old infants in the occlusion and containment experiments of Wang *et al.* (2004) support both of these predictions.

Testing transient-preference accounts


In some VOE tasks, it may not be possible to give infants only test trials (for fuller discussion, see Wang *et al.*, 2004). In particular, when shown events involving novel self-moving objects, unfamiliar motions, long event sequences and so on, infants may well require some habituation or familiarization trials in order to be able to focus in the test trials on the key manipulations of interest to the researchers. In such cases, how can we be certain that infants look reliably longer at the unexpected events because these events violate their physical knowledge, and *not* because the habituation or familiarization trials induced in them transient and superficial preferences for the events? The only recourse, as always, is to empirically test specific hypotheses about possible transient preferences.

To illustrate, consider a transient-preference account proposed by Bogartz *et al.* (1997) for one of our findings (Baillargeon & Graber, 1987; see also Baillargeon & DeVos, 1991; Luo, Baillargeon & Lécuyer, 2004). The infants in this experiment received familiarization trials in which they saw a tall or a short toy rabbit move back and forth behind a screen (see Figure 14). Next, a window was created in the screen's upper half, and the infants again saw the tall and the short rabbit move back and forth behind the screen. The short rabbit was shorter than the bottom of the window and did not become visible when passing behind the screen; the tall rabbit should have appeared in the window, but did not in fact

do so. The infants looked reliably longer at the tall- than at the short-rabbit test event, and we suggested that the infants had identified height as an occlusion variable and were surprised when the tall rabbit failed to appear in the window. Bogartz and his colleagues offered a very different, transient-preference account: they suggested that the infants focused on the rabbit's face in each familiarization event and, as they scanned horizontally back and forth, attended only to the portion of the screen that lay at the same height as the face. During test, the infants continued to scan the events in the same manner; as a result, they detected the novel window in the unexpected but not the expected event. The infants thus looked reliably longer at the unexpected event simply because they noticed the change in the screen in this event.

There is now a large body of experimental evidence inconsistent with the alternative account proposed by Bogartz *et al.* (1997). To begin with, our original experiment included a control condition that addressed the account (Baillargeon & Graber, 1987). The infants in this control condition received two pretest trials at the start of the test session in which they saw two tall or two short rabbits standing on either side of the familiarization screen. Like the 3-month-old infants in the two-mouse experiment described earlier (Aguiar & Baillargeon, 2002; see Section 2A), the infants in this control condition tended to look equally at the tall- and short-rabbit test events, suggesting that they were able to take advantage of the two-rabbit hint to make sense of the tall-rabbit event.⁷

Familiarization events



Test events




Figure 14 Familiarization and test events used by Baillargeon and Graber (1987).

⁷ If the 3.5-month-old infants in the mouse task described in Section 2A could posit a second, identical mouse to make sense of the occlusion violation they were shown (Aguiar & Baillargeon, 2002; see also Spelke & Kestenbaum, 1986; Spelke *et al.*, 1995a), why did the 5.5-month-old infants in the rabbit task not posit a second, identical rabbit (Baillargeon & Graber, 1987; see also Baillargeon & DeVos, 1991)? Recall that the infants in the mouse task faced a more flagrant violation than did those in the rabbit task: in the mouse task, the entire mouse failed to become visible as expected; in the rabbit task, only the top portion of the tall rabbit failed to become visible. It could be that the infants in the rabbit task assumed that the rabbit traveled the distance behind the screen, and were then puzzled as to why the top of the rabbit did not appear in the window. In the mouse task, the infants could not easily assume that the mouse traveled from one end of the screen to the other; this task may thus have been more conducive to the production of a two-object explanation.

The preceding interpretation makes an intriguing prediction: 5.5-month-old infants in the rabbit task should tend to look equally at the test events if tested with a *larger* high window in the tall-rabbit test event. With a larger window, the infants should be more likely (1) to realize that the tall rabbit did not travel the distance behind the screen and hence (2) to conclude that two identical tall rabbits must be involved in the event. Making the violation in the tall-rabbit test event more obvious (i.e. having a greater portion of the tall rabbit fail to appear in the window) should thus have the counterintuitive effect of eliminating infants' overall surprise at the violation.

In addition, several experiments have provided converging evidence that infants aged 4.5 to 6 months can reason about the variable height in occlusion events. First, as we saw in Section 2B, 4.5-month-old infants respond with increased attention when a tall object becomes hidden behind either a short occluder or a short container (Hespos & Baillargeon, 2001a; see Figure 7). Second, in a recent experiment, Yuyan Luo, Roger Lécuyer and I asked whether 5-month-old infants could predict, not *whether* a tall object should appear in a high window, but *how far up* the object should reach in the window (Luo *et al.*, 2004). The infants succeeded in detecting two violations: in one, a short cylinder that should have reached only half-way up the window actually reached the very top; in the other violation, a tall cylinder that should have reached the top of the window reached only half-way up.

Third, further converging evidence that young infants can reason about height in occlusion events comes from a recent action task Sue Hespos and I administered to 6- and 7.5-month-old infants (Hespos & Baillargeon, 2004). Each infant sat across from an experimenter at a table on which stood a large screen. The experimenter first brought out a tall frog from behind the screen, and set it on the table (pre-trial phase). After a few seconds, the experimenter returned the frog behind the screen, which was then removed to reveal a tall and a short occluder; two frog feet protruded from the bottom of each occluder, one on either side (main-trial phase). Our reasoning was that if the infants wanted to find the tall frog, and realized that it could be hidden behind the tall but not the short occluder, then they should be more likely to reach for the tall than for the short occluder. The infants received four trials, and the positions of the tall and short occluders were counterbalanced across trials; the infants were said to have succeeded at the task if they reached for the tall occluder on three or more trials. The results indicated that 77.8% of the 6-month-olds, and again 78% of the 7.5-month-olds, reached for the tall occluder on three or more trials. In a control condition, infants were not shown the tall frog, and tended to reach equally for the short and tall occluders.

There are thus several experiments, involving different methods, events and objects, which provide converging evidence that infants aged 4.5 to 6 months attend to height information in occlusion events. This evidence does not support the transient-preference account proposed by Bogartz *et al.* (1997).

Conclusions

We have seen in this section that even young infants can succeed at VOE tasks involving hidden objects

when given only test trials (Wang *et al.*, 2004). Such a demonstration does not mean, of course, that infants should succeed at all VOE tasks involving hidden objects in the absence of habituation or familiarization trials; such trials may sometimes be needed to acquaint infants with various aspects of the experimental situation and thus help them focus in the test trials on the key manipulations of interest to the investigators. In cases where habituation or familiarization trials are required, alternative accounts that appeal to transient preferences induced by these trials must be evaluated empirically. In this section, we considered one such transient-preference account (Bogartz *et al.*, 1997), and concluded that additional findings did not support it (for a discussion of other transient-preference accounts, see Wang *et al.*, 2004).

5B. Weak-representation accounts

A few researchers have suggested that, although young infants may be able to represent hidden objects, these representations are likely to be weak and short-lived – sufficient for success in most VOE tasks, which typically require infants to represent hidden objects for only a few seconds at a time, but not in more challenging tasks (e.g. Haith, 1998, 1999; Haith & Benson, 1998; Munakata, 2001; Munakata *et al.*, 1997).

In a recent experiment, Yuyan Luo, Laura Brueckner, Yuko Munakata and I tested whether 5-month-old infants in a VOE task could represent a hidden object for a substantial delay (Luo, Baillargeon, Brueckner & Munakata, 2003). To succeed in the experiment, the infants had to reason about an object hidden 3 or 4 minutes prior to the test trials.

The infants were assigned to a thin- or a thick-box condition (see Figure 15). The infants in the *thin-box* condition first received five familiarization trials. During the first trial (box-familiarization event), a screen lay flat on the apparatus floor, toward the infants, and an experimenter placed a thin box behind the screen, against the back wall of the apparatus. Because the box was thin (in depth), a substantial gap remained between the box and the screen. The second, third and fourth familiarization trials (screen-familiarization event) were all identical: at the start of each trial, the screen was raised to hide the box. After the screen was raised in the fourth familiarization trial, it remained upright for the rest of the experiment: the infants never again saw the box. During the fifth and final familiarization trial (cylinder-familiarization event), the experimenter placed a tall cylinder on the apparatus floor, next to the left wall. A timer was set for an interval of either 3 or 4 minutes when the screen was raised for the last time at the start of the fourth




Figure 15 Familiarization and test events used by Luo et al. (2003).

familiarization trial; during this interval, the infants completed their fourth and fifth familiarization trials and then interacted with their parent for whatever time remained in the interval. When the timer rang to signal the end of the interval, the test trials began. The infants received two blocks of three test trials in which they saw the cylinder move back and forth behind the screen. The infants in the *thick-box* condition received identical familiarization and test trials, except that the box was much thicker so that there was almost no gap between the box and the screen. Thus, it was possible for the cylinder to move back and forth behind the screen in the thin- but not the thick-box condition.

The same results were found for the infants who received a 3- or a 4-min delay. In the first block of test trials, the infants in the thick- and thin-box conditions

tended to look equally; their responses were essentially at ceiling, no doubt because they had never seen the cylinder move across the apparatus before. In the second block of trials, the infants in the thick-box condition looked reliably longer than those in the thin-box condition, suggesting that they (1) remembered the thick or thin box behind the screen after the delay, and (2) realized that the cylinder could pass behind the screen when the thin but not the thick box was present.

Conclusions

The results just summarized provide little support for the notion that young infants' representations of hidden objects are weak and short-lived. To the contrary, they suggest that, by 5 months of age, infants' representations of hidden objects are robust enough to withstand significant delays.

6. Final remarks

What do infants know about hidden objects? The research reviewed in this article supports six general conclusions. First, infants as young as 2.5 months of age realize that an object continues to exist after it becomes hidden behind an occluder, inside a container or under a cover. Second, infants are rather poor initially at predicting when an object behind an occluder, inside a container or under a cover should be hidden. Third, infants' predictions gradually improve as they identify the variables relevant for predicting outcomes in each event category. Fourth, infants who *have* identified a variable in an event category (1) can detect surreptitious changes and other violations involving the variable, and (2) can sometimes also generate explanations to make sense of these violations. Fifth, infants who have *not* yet identified a variable (1) cannot detect surreptitious changes and other violations involving the variable, but (2) can detect such violations after being taught the variable, both immediately and after a 24-hour delay. Finally, all of these findings are consistent with the new account of infants' physical reasoning presented here, which assumes that both event-general and event-specific expectations contribute to infants' responses to physical events.

Acknowledgements

This article is based on an invited address ('Infants' Physical World') presented at the biennial meeting of the Society for Research in Child Development in Tampa, Florida, in April 2003. The research reviewed here was

supported by a grant from the National Institute of Child Health and Human Development (HD-21104). I am grateful to Colette DeJong, Jerry DeJong, Cindy Fisher, Yuyan Luo, Kris Onishi, Hyun-joo Song and Suhua Wang for helpful discussions and comments; to Steve Holland, our graphic artist, for his patience and creativity in preparing our figures; to Laura Brueckner, Venessa Nolen and the staff of the Infant Cognition Laboratory for their help with the data collection; and last but not least to the many parents who kindly agreed to have their infants participate in our research.

References

- Aguiar, A., & Baillargeon, R. (1998). Eight-and-a-half-month-old infants' reasoning about containment events. *Child Development*, **69**, 636–653.
- 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.
- Aguiar, A., & Baillargeon, R. (2003). Perseverative responding in a violation-of-expectation task in 6.5-month-old infants. *Cognition*, **88**, 277–316.
- Arterberry, M.E. (1997). Perception of object properties over time. In C. Rovee-Collier & L.P. Lipsitt (Eds.), *Advances in infancy research*, Vol. 11 (pp. 219–268). Greenwich, CT: Ablex.
- Aslin, R.N. (2000). Why take the cog out of infant cognition? *Infancy*, **1**, 463–470.
- 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. (1994a). How do infants learn about the physical world? *Current Directions in Psychological Science*, **3**, 133–140.
- Baillargeon, R. (1994b). Physical reasoning in young infants: seeking explanations for unexpected events. *British Journal of Developmental Psychology*, **12**, 9–33.
- Baillargeon, R. (1995). A model of physical reasoning in infancy. In C. Rovee-Collier & L.P. Lipsitt (Eds.), *Advances in infant research*, Vol. 9 (pp. 305–371). Norwood, NJ: Ablex.
- Baillargeon, R. (1998). Infants' understanding of the physical world. In M. Sabourin, F. Craik, & M. Robert (Eds.), *Advances in psychological science*, Vol. 2 (pp. 503–529). London: Psychology Press.
- Baillargeon, R. (1999). Young infants' expectations about hidden objects: a reply to three challenges (article with peer commentaries and response). *Developmental Science*, **2**, 115–163.
- Baillargeon, R. (2000). Reply to Bogartz, Shinskey, and Schilling; Schilling; and Cashon and Cohen. *Infancy*, **1**, 447–462.
- Baillargeon, R. (2001). Infants' physical knowledge: of acquired expectations and core principles. In E. Dupoux (Ed.), *Language, brain, and cognitive development: Essays in honor of Jacques Mehler* (pp. 341–361). Cambridge, MA: MIT Press.
- Baillargeon, R. (2002). The acquisition of physical knowledge in infancy: a summary in eight lessons. In U. Goswami (Ed.), *Handbook of childhood cognitive development* (pp. 47–83). Oxford: Blackwell.
- Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: further evidence. *Child Development*, **62**, 1227–1246.
- Baillargeon, R., Fisher, C., & DeJong, G.F. (2000). Teaching infants about support: What data must they see? Paper presented at the biennial International Conference on Infant Studies, Brighton, England, July.
- 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., & Luo, Y. (2002). Development of object concept. *Encyclopedia of Cognitive Science*, Vol. 3 (pp. 387–391). London: Nature Publishing Group.
- Baillargeon, R., Needham, A., & DeVos, J. (1992). The development of young infants' intuitions about support. *Early Development and Parenting*, **1**, 69–78.
- Baillargeon, R., & Wang, S. (2002). Event categorization in infancy. *Trends in Cognitive Sciences*, **6**, 85–93.
- Bogartz, R.S., Shinskey, J.L., & Schilling, T.H. (2000). Object permanence in five-and-a-half-month-old infants? *Infancy*, **1**, 403–428.
- Bogartz, R.S., Shinskey, J.L., & Speaker, C.J. (1997). Interpreting infant looking: the event set \times event set design. *Developmental Psychology*, **33**, 408–422.
- Bremner, J., & Mareschal, D. (2004). Reasoning . . . what reasoning? *Developmental Science*, **7**, 419–421.
- Bruner, J.S. (1973). *Beyond the information given*. New York: Norton.
- 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: Cambridge University Press.
- Casasola, M., Cohen, L., & Chiarello, E. (2003). Six-month-old infants' categorization of containment spatial relations. *Child Development*, **74**, 679–693.
- Cashon, C.H., & Cohen, L.B. (2000). Eight-month-old infants' perceptions 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.
- DeJong, G.F. (1993). *Investigating explanation-based learning*. Boston, MA: 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.
- Haith, M.M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior and Development*, **21**, 167–179.

- Haith, M.M. (1999). Some thoughts about claims for innate knowledge and infant physical reasoning. *Developmental Science*, **2**, 153–156.
- Haith, M.M., & Benson, J.B. (1998). Infant cognition. In W. Damon (Series Ed.) & D. Kuhn & R. Siegler (Vol. Eds.), *Handbook of child psychology*, Vol. 2 (pp. 199–254). New York: Wiley.
- 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. (2004). Décalage in infants' knowledge about occlusion and containment events: converging evidence from action tasks. Under review.
- Hood, B. (2004). Is looking good enough or does it beggar belief? *Developmental Science*, **7**, 415–417.
- Huettel, S.A., & Needham, A. (2000). Effects of balance relations between objects on infants' object segregation. *Developmental Science*, **3**, 415–427.
- Johnson, S.P., & Aslin, R.N. (2000). Infants' perception of transparency. *Developmental Psychology*, **36**, 808–816.
- 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.
- 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.
- Lécuyer, R. (2001). Rien n'est jamais acquis. De la permanence de l'objet . . . de polémiques. *Enfance*, **55**, 35–65.
- Lécuyer, R., & Durand, K. (1998). Bi-dimensional representations of the third dimension and their perception by infants. *Perception*, **27**, 465–472.
- 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. (2004). Who's for learning? *Developmental Science*, **7**, 417–419.
- Luo, Y., & Baillargeon, R. (2004a). Development of infants' reasoning about transparent occluders. Manuscript in preparation.
- Luo, Y., & Baillargeon, R. (2004b). Infants' reasoning about transparent occluders and containers. Manuscript in preparation.
- Luo, Y., & Baillargeon, R. (in press). When the ordinary seems unexpected: evidence for rule-based physical reasoning in young infants. *Cognition*.
- Luo, Y., Baillargeon, R., Brueckner, L., & Munakata, Y. (2003). Reasoning about a hidden object after a delay: evidence for robust representations in 5-month-old infants. *Cognition*, **88**, B23–B32.
- Luo, Y., Baillargeon, R., & Lécuyer, R. (2004). Young infants' reasoning about height in occlusion events. Manuscript in preparation.
- McCall, D. (2001). Perseveration and infants' sensitivity to cues for containment. Paper presented at the biennial meeting of the Society for Research in Child Development, Minneapolis, MN, April.
- McDonough, L., Choi, S., & Mandler, J.M. (2003). Understanding spatial relations: flexible infants, lexical adults. *Cognitive Psychology*, **46**, 229–259.
- Munakata, Y. (1997). Perseverative reaching in infancy: the roles of hidden toys and motor history in the A \bar{B} task. *Infant Behavior and Development*, **20**, 405–416.
- Munakata, Y. (2000). Challenges to the violation-of-expectation paradigm: throwing the conceptual baby out with the perceptual processing bathwater? *Infancy*, **1**, 471–490.
- Munakata, Y. (2001). Task-dependency in infant behavior: toward an understanding of the processes underlying cognitive development. In F. Lacerda, C. von Hofsten, & M. Heimann (Eds.), *Emerging cognitive abilities in early infancy* (pp. 29–52). Mahwah, NJ: Erlbaum.
- Munakata, Y., McClelland, J.L., Johnson, M.H., & Siegler, R. (1997). Rethinking infant knowledge: toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, **104**, 686–713.
- Needham, A. (2000). Improvements in object exploration skills may facilitate the development of object segregation in early infancy. *Journal of Cognition and Development*, **1**, 131–156.
- 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.
- Onishi, K.H. (2000). Infants can reason about the support of 2 but not 3 stacked boxes. Paper presented at the biennial International Conference on Infant Studies, Brighton, England, July.
- 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.
- Schilling, T.H. (2000). Infants' looking at possible and impossible screen rotations: the role of familiarization. *Infancy*, **1**, 389–402.
- Simons, D.J. (1996). In sight, out of mind: when object representations fail. *Psychological Science*, **7**, 301–305.
- Simons, D.J., Franconeri, S.L., & Reimer, R.L. (2000). Change blindness in the absence of a visual disruption. *Perception*, **29**, 1143–1154.
- Sitskoorn, S.M., & Smitsman, A.W. (1995). Infants' perception of dynamic relations between objects: passing through or support? *Developmental Psychology*, **31**, 437–447.

- 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, NJ: Ablex.
- Spelke, E.S. (1982). Perceptual knowledge of objects in infancy. In J. Mehler, E. Walker, & M. Garrett (Eds.), *Perspectives on mental representation* (pp. 409–430). Hillsdale, NJ: Erlbaum.
- Spelke, E.S. (1994). Initial knowledge: six suggestions. *Cognition*, **50**, 431–445.
- Spelke, E.S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, **99**, 605–632.
- Spelke, E.S., & Hespos, S.J. (2002). Conceptual development in infancy: the case of containment. In N.L. Stein, P. Bauer, & M. Rabinowitch (Eds.), *Representation, memory, and development: Essays in honor of Jean Mandler* (pp. 223–246). Mahwah, NJ: Erlbaum.
- Spelke, E.S., & Kestenbaum, R. (1986). Les origines du concept d'objet. *Psychologie Française*, **31**, 67–72.
- 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.
- Thelen, E., & Smith, L.B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Wang, S., & Baillargeon, R. (2004a). Change blindness in infancy: event-category effects. Manuscript in preparation.
- Wang, S., & Baillargeon, R. (2004b). Inducing infants to detect continuity violations: a new approach. Manuscript under review.
- Wang, S., & Baillargeon, R. (2004c). Teaching infants to attend to a variable in an event category: the case of height in covering events. Manuscript in preparation.
- 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. (in press). Detecting continuity violations in infancy: a new account and new evidence from covering and tube events. *Cognition*.
- Wang, S., Kaufman, L., & Baillargeon, R. (2003). Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events (Special issue). *Infant Behavior and Development*, **26**, 529–568.
- Wilcox, T. (1999). Object individuation: infants' use of shape, size, pattern, and color. *Cognition*, **72**, 125–166.
- Wilcox, T., & Baillargeon, R. (1998a). Object individuation in infancy: the use of featural information in reasoning about occlusion events. *Cognitive Psychology*, **17**, 97–155.
- Wilcox, T., & Baillargeon, R. (1998b). Object individuation in young infants: further evidence with an event-monitoring task. *Developmental Science*, **1**, 127–142.
- Wilcox, T., & Chapa, C. (2002). Infants' reasoning about opaque and transparent occluders in an object individuation task. *Cognition*, **85**, B1–B10.
- 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 fullterm infants. *Infant Behavior and Development*, **19**, 309–323.
- Wilcox, T., & Schweinle, A. (2003). Infants' use of speed information to individuate objects in occlusion events. *Infant Behavior and Development*, **26**, 253–282.
- 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.
- Xu, F., & Carey, S. (1996). Infants' metaphysics: the case of numerical identity. *Cognitive Psychology*, **30**, 111–153.
- 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.

Received: 28 November 2003

Accepted: 23 April 2004