Available online at www.sciencedirect.com
SCIENCE DIRECT®
Cognition 95 (2005) 129–173
COGNITION
www.elsevier.com/locate/COGNIT

Detecting continuity violations in infancy: a new account and new evidence from covering and tube events

Su-hua Wang[a,*], Renée Baillargeon[b], Sarah Paterson[c]

[a]Department of Psychology, University of California, 1156 High Street, Santa Cruz, CA 95064, USA
[b]University of Illinois at Urbana-Champaign, Champaign, IL, USA
[c]Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, USA

Received 16 January 2002; revised 29 September 2002; accepted 11 November 2002

Abstract

Recent research on infants' responses to occlusion and containment events indicates that, although some violations of the continuity principle are detected at an early age e.g. 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; Hespos, S. J., & Baillargeon, R. (2001). Knowledge about containment events in very young infants. Cognition 78, 207–245; Luo, Y., & Baillargeon, R. (in press). When the ordinary seems unexpected: Evidence for rule-based reasoning in young infants. Cognition; Wilcox, T., Nadel, L., & Rosser, R. (1996). Location memory in healthy preterm and full-term infants. Infant Behavior & Development 19, 309–323, others are not detected until much later e.g. Baillargeon, R., & DeVos, J. (1991). Object permanence in young infants: Further evidence. Child Development 62, 1227–1246; Hespos, S. J., & Baillargeon, R. (2001). Infants' knowledge about occlusion and containment events: A surprising discrepancy. Psychological Science 12, 140–147; Luo, Y., & Baillargeon, R. (2004). Infants' reasoning about events involving transparent occluders and containers. Manuscript in preparation; Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. Cognition 72, 125–166. The present research focused on events involving covers or tubes, and brought to light additional examples of early and late successes in infants' ability to detect continuity violations. In Experiment 1, 2.5- to 3-month-old infants were surprised (1) when a cover was lowered over an object, slid to the right, and lifted to reveal no object; and (2) when a cover was lowered over an object, slid behind

* Corresponding author.
E-mail address: suhua@ucsc.edu (S.-h. Wang).

0022-2860/$ - see front matter
doi:10.1016/j.cognition.2002.11.001

the left half of a screen, lifted above the screen, moved to the right, lowered behind the right half of the screen, slid past the screen, and finally lifted to reveal the object. In Experiments 2 and 3, 9- and 11-month-old infants were *not* surprised when a short cover was lowered over a tall object until it became fully hidden; only 12-month-old infants detected this violation. Finally, in Experiment 4, 9-, 12-, and 13-month-old infants were *not* surprised when a tall object was lowered inside a short tube until it became fully hidden; only 14-month-old infants detected this violation. A new account of infants' physical reasoning attempts to make sense of all of these results. New research directions suggested by the account are also discussed.

Keywords: Physical reasoning in infancy; Continuity violations; Event categories

1. Introduction

Spelke and her colleagues (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke, Breinlinger, Macomber, & Jacobson, 1992; Spelke, Phillips, & Woodward, 1995) have proposed that, from birth, infants interpret physical events in accord with a core principle of *continuity*, which states that objects exist continuously in time and space. This principle has many corollaries: for example, that stationary objects, whether visible or hidden, exist continuously in time; that moving objects, whether visible or hidden, travel along spatially continuous paths; and that two objects, whether visible or hidden and whether stationary or moving, cannot occupy the same space at the same time.[1]

Spelke's proposal (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke et al., 1992, 1995) might be taken to suggest that infants, at all ages, should detect all salient violations of the continuity principle. However, there is growing evidence that infants do not in fact detect all salient continuity violations: although some violations are detected at a young age, others are not detected until much later. Most of this evidence comes from violation-of-expectation tasks (e.g. Baillargeon, 1998; 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 they (1) possess the expectation under investigation; (2) detect the violation in the unexpected event; and (3) are surprised by this violation. The term "surprise" is used here simply as a short-hand descriptor, to denote a state of heightened attention or interest caused by an expectation violation.

To illustrate the claim that some continuity violations are detected at an early age but others not, consider the results of recent experiments on infants' responses to occlusion and containment events. On the one hand, (1) 2.5-month-old infants are surprised when an object disappears behind one occluder and reappears from behind another occluder

[1] This last corollary is often referred to as the *solidity* principle. Here we follow Spelke's lead (e.g. Carey & Spelke, 1994; Spelke et al., 1995) in viewing solidity as a consequence of a more general principle of continuity. Solidity comes into play whenever two objects are involved: for each object to exist continuously, the two cannot occupy the same space at the same time.

without appearing in the gap between them (Aguiar & Baillargeon, 1999; Luo & Baillargeon, in press; Wilcox, Nadel, & Rosser, 1996); (2) 2.5-month-old infants are surprised when an object is lowered inside a container through its closed top (Hespos & Baillargeon, 2001b); and (3) 2.5-month-old infants are surprised when an object that has been lowered inside a container is revealed when the container is slid aside (Hespos & Baillargeon, 2001b). On the other hand, (1) infants less than 3.5 months of age are not surprised when a tall object becomes hidden behind a short occluder (Aguiar & Baillargeon, 2002; Baillargeon & DeVos, 1991); (2) infants less than 7.5 months of age are not surprised when two objects, identical except for pattern (e.g. dots or stripes), emerge successively from behind a screen too narrow to hide them both (Wilcox, 1999; Wilcox & Chapa, 2004); (3) infants less than 7.5 months of age are not surprised when a tall object becomes hidden inside a short container (Hespos & Baillargeon, 2001a); (4) infants less than 9.5 months of age are not surprised when an object becomes hidden inside a transparent container (Luo & Baillargeon, 2005); and (5) infants less than 11.5 months of age are not surprised when two objects, identical except for color, emerge successively from behind a screen too narrow to hide them both (Wilcox, 1999; Wilcox & Chapa, 2004). Infants thus detect some continuity violations at a young age, but others only much later.

How can we reconcile these seemingly divergent results (for additional results, see Baillargeon, 1987, 1991, 1993, 1995; Newcombe, Huttenlocher, & Learmonth, 1999; Spelke et al., 1992; Wang et al., 2004)? To address this question, we have been developing a new account of infants' physical reasoning that builds on Spelke's proposal (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke et al., 1992, 1995) and attempts to make sense of infants' early and late successes at detecting different continuity violations. In this article, we briefly describe this new reasoning account and then report experiments that tested two of its predictions, one concerning younger infants and one older infants.

1.1. An account of infants' physical reasoning

Our account of infants' physical reasoning rests on four main assumptions (see Fig. 1). First, when watching a physical event, infants build a specialized *physical representation* that is used to predict and interpret the event's outcome. Second, the information in infants' physical representation of an event becomes subject to a few core principles, including that of *continuity* (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke et al., 1992, 1995).[2]

[2] 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 vivid 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 way, one might suggest that the principle of continuity bestows continued existence upon objects: all other things being equal, objects in physical representations are expected to persist, with all of their physical properties, through time and space. Although the events used in the present research involved forces (e.g. hands lifting, lowering, and sliding covers or objects), we focus here only on those aspects of the events that concerned the continuity principle.

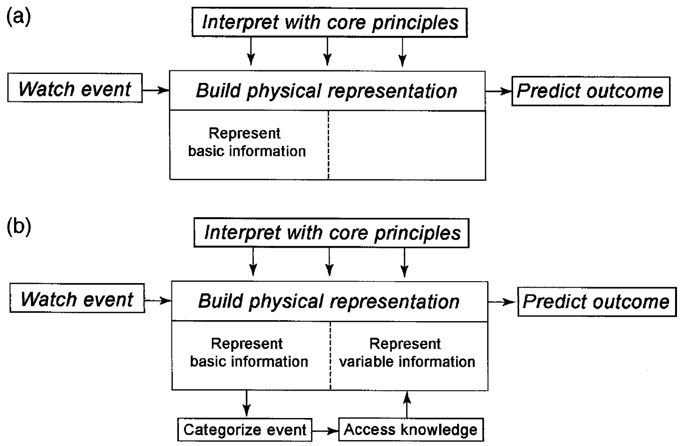

Fig. 1. Schematic presentation of the reasoning account for younger (a) and older (b) infants.

Third, in the first weeks of life, infants' physical representation of an event typically includes only basic spatial and temporal information about the event (see Fig. 1(a); e.g. Kestenbaum, Termine, & Spelke, 1987; Leslie, 1994; Needham, 2000; Slater, 1995; Spelke, 1982; Yonas & Granrud, 1984). This *basic information* specifies primarily (1) how many distinct objects are involved in the event (e.g. are there two objects present?); (2) what is the geometry of each object (e.g. is one object open at the top to form a container, open at the bottom to form a cover, or open at both ends to form a tube?); and (3) what is the spatial arrangement of the objects and how does it change over time (e.g. is one object being placed behind, inside, or under the other object?). The basic information thus captures essential aspects of the event, but leaves out most details: for example, it includes no information about the relative sizes of the objects (e.g. is one object taller or wider than the other object?), or about their surface appearance (e.g. are the objects transparent or opaque, dotted or striped, red or green?).

Fourth, with experience, infants include more and more of this detailed information, or *variable information*, in their physical representations (see Fig. 1(b)). Recent evidence indicates that, as they learn about physical events, infants form distinct event categories, such as occlusion and containment events (e.g. Aguiar & Baillargeon, 2003; Casasola & Cohen, 2002; Casasola, Cohen, & Chiarello, 2003; Hespos & Baillargeon, 2001a; Luo & Baillargeon, 2005; McDonough, Choi, & Mandler, 2003; Wilcox & Chapa, 2002); for each category, infants gradually identify a series of variables that enables them to predict outcomes within the category more and more accurately over time (e.g. Aguiar & Baillargeon, 2002; Baillargeon, 1991; Dan, Omori, & Tomiyasu, 2000; Huettel & Needham, 2000; Kotovsky & Baillargeon, 1998; Sitskoorn & Smitsman, 1995; Wang, Kaufman, & Baillargeon, 2003; Wilcox, 1999; for recent reviews, see Baillargeon, 2002, 2004; Baillargeon & Wang, 2002). This variable knowledge determines what detailed information infants include in their physical representations. 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 (e.g. the relative heights of the occluder and object in occlusion events, or the transparency of the container in containment events), and hence that should be included in the physical representation of the event. Variables not yet identified are typically not included in the representation.

To illustrate our reasoning account, consider first the finding that infants aged 7.5 months and older are surprised when a tall object is lowered inside a short container until it becomes hidden (Hespos & Baillargeon, 2001a). The account suggests that, when watching this event, infants represent the basic information about the event and interpret this information in accord with their continuity principle ("object being lowered inside container through open top"). Next, infants categorize the event as a containment event, and access their knowledge of this event category. Because at 7.5 months this knowledge includes the variable height, infants include information about the relative heights of the object and container in their physical representation of the event. This variable information then becomes subject to infants' continuity principle, making it possible for them to detect the continuity violation in the event: they recognize that the object is too tall to become hidden inside the short container. Infants younger than 7.5 months, who have not yet identified height as a containment variable, typically do not include height information in their physical representations of containment events. As a result, this information is not available and hence cannot be interpreted in accord with infants' continuity principle. Infants thus fail to detect continuity violations involving tall objects and short containers (Hespos & Baillargeon, 2001a).

Next, consider the finding that infants aged 2.5 months and older are surprised when an object is lowered inside a container through its closed top (Hespos & Baillargeon, 2001b). According to the reasoning account, infants again represent the basic information about the event and interpret this information in accord with their continuity principle ("object being lowered inside container through closed top"). In this case, no variable information (e.g. about the relative sizes or surface appearance of the object and container) is needed for infants to detect the continuity violation in the event: the interpretation of the basic information alone is sufficient for them to realize that the object should not have been able to pass through the closed top of the container.

1.2. Basic and variable continuity violations

Implicit in the reasoning account presented in the last section is the distinction between two types of continuity violations: those that can be detected through the inclusion and interpretation of purely basic information, or *basic continuity violations*; and those that can only be detected through the inclusion and interpretation of additional, variable information, or *variable continuity violations*. According to the account, basic violations should be detected early, because they involve information that even young infants typically include in their physical representations of events.[3] In contrast, variable

[3] We do not mean to imply that there are no improvements with age in infants' ability to represent the basic information about events; for example, the range of cues infants can use to represent depth relations between objects certainly increases with age (e.g. Held, Thorn, Gwiazda, & Bauer, 1996; Yonas & Granrud, 1984).

violations should be detected at somewhat later ages, after infants have identified the appropriate variables and include information about these variables in their physical representations.

In the opening paragraphs of this article, we reviewed evidence of early as well as late successes in infants' responses to continuity violations in occlusion and containment events (e.g. Aguiar & Baillargeon, 1999, 2002; Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a,b; Luo & Baillargeon, 2005, in press; Wilcox, 1999; Wilcox & Chapa, 2004; Wilcox et al., 1996). Our reasoning account helps explain these divergent results: it makes clear that the early successes involved *basic* violations, and the late successes *variable* violations. In order to detect the violation that occurs when an object fails to appear between two occluders (e.g. Aguiar & Baillargeon, 1999; Luo & Baillargeon, in press; Wilcox et al., 1996), or when an object that has been lowered inside a container is revealed when the container is slid aside (e.g. Hespos & Baillargeon, 2001a), infants need only represent the basic information about the objects present, their open/closed surfaces, and their spatial arrangement over time. However, to detect the violation that occurs when a tall object becomes hidden behind a short occluder (e.g. Baillargeon & DeVos, 1991; Baillargeon & Graber, 1987; Hespos & Baillargeon, 2001a), or when an object that has been lowered inside a transparent container is not visible through the front of the container (e.g. Luo & Baillargeon, 2005), infants must represent not only the basic information about the objects, but also more detailed–or variable information about their relative sizes or surface appearance. According to the reasoning account, infants typically do not include information about the variable height in their representations of occlusion events, or about the variable transparency in their representations of containment events, until they have identified these variables.

1.3. The present research

If the reasoning account introduced in the previous sections is correct, then one would expect the same general pattern of early successes with basic continuity violations, and later successes with variable continuity violations, to be observed not only in occlusion and containment events, but in other event categories as well. Experiments 1–3 examined whether this pattern would be found in infants' responses to *covering* events—events in which rigid covers are lowered over objects.

Experiment 1 asked whether 2.5- to 3-month-old infants would succeed in detecting *basic* continuity violations in covering events. The infants saw test events in which a cover was lowered over an object and then slid or lifted to a new location. In the violations, (1) the object failed to move with the cover when slid to a new location, or (2) the object moved with the cover when lifted to a new location. Although fairly complex, these violations could nevertheless be detected through the inclusion and interpretation of basic information about the presence and spatial arrangement of the cover and object; no variable information (e.g. about the relative heights or widths of the cover and object) was necessary for the violations to be detected.

Experiments 2 and 3 asked whether 9- to 12-month-old infants would succeed in detecting a *variable* continuity violation in a covering event. The infants saw test events in which a cover was lowered over a tall object until the object became fully hidden.

In the violation, the cover was considerably shorter than the object. Unlike the violations shown in Experiment 1, this violation could only be detected through the inclusion and interpretation of variable information, specifically, information about the relative heights of the cover and object.

Why did we choose to study covering events? It seemed plausible that infants might view these events as a distinct event category. Many of infants' initial event categories appear to capture relatively simple spatial relations between objects, such as "object *behind* other object" (occlusion), "object *inside* container" (containment), or "object *on* other object" (support) (e.g. Casasola et al., 2003; McDonough et al., 2003; Spelke & Hespos, 2001; see also Quinn & Eimas, 1996, for a review of spatial categorization in young infants).[4] Covering events would seem to capture yet another spatial relation, "object *under* cover". These different spatial relations often lead to different physical outcomes. Contrast, for example, events in which an object is placed behind an occluder, inside a container, or under a cover. We would expect the object to remain stationary when the occluder or the cover but not the container is lifted vertically. Conversely, we would expect the object to move when the container or the cover but not the occluder is slid to the side. These distinct clusters of spatial and mechanical relationships may well provide the bases for infants' early event categories (e.g. Keil, 1995; Leslie, 1995; Pauen, 1999).

We were aware that, despite the above differences, containment and covering events might initially form a single, indistinct category for infants; after all, a cover is often no more than an inverted container. If this were the case, then the present research could provide additional evidence of early and late successes in infants' ability to detect continuity violations in containment events—but it could not extend these results to a new event category. Fortunately, Experiments 2 and 3 could shed light on this issue. These experiments examined whether infants aged 9–12 months could reason about the variable height in covering events. Recall that the variable height is identified at about 7.5 months in containment events (Hespos & Baillargeon, 2001a). If infants view events involving containers and covers as a single, indistinct category, then they should also be able to reason about the variable height in covering events at about 7.5 months. Thus, evidence that infants do *not* attend to height information in covering events until later in the first year would suggest that infants view events involving containers and covers as distinct categories, and do not generalize variables identified in the context of containment events to covering events (see Baillargeon, 2002; Baillargeon & Wang, 2002, for discussion of other lags or décalages in infants' physical reasoning).

2. Experiment 1

The reasoning account presented in the Introduction suggests that young infants should detect *any* continuity violation, in *any* event category, as long as this violation involves

[4] Hespos and Baillargeon (2001a) found that 4.5-month-old infants are surprised when a tall object becomes almost fully hidden *behind* but not *inside* a short container. This finding makes clear that spatial relations play a key role in infants' categorization of events: an object being lowered behind a container is categorized as an occlusion event, not a containment event.

only the *basic* spatial and temporal information they can represent. Experiment 1 examined 2.5- to 3-month-old infants' ability to detect basic continuity violations in covering events (for related findings with older infants, see Baillargeon, Graber, DeVos, & Black, 1990; Leslie & Das Gupta, cited in Leslie, 1995). The infants were assigned to a slid- or a lifted-cover condition; all of the infants saw test events in which a rigid cover was lowered over a toy duck and then moved.

The infants in the *slid-cover* condition first received four familiarization trials (see Fig. 2). These trials were designed to acquaint the infants with the duck (duck event), the cover (cover event), and the cover's motion (cover-motion event); the first two events were shown on the first and second trials, respectively, and the third event was shown on the third and fourth trials. Next, the infants received a test trial in which they saw one of two test events (see Fig. 3). At the beginning of each event, the duck rested on the left end of a platform (from the infants' point of view); an experimenter's gloved hand held the cover above the duck. To start, the hand lowered the cover over the duck and then slid the cover to the right end of the platform. Next, the hand lifted

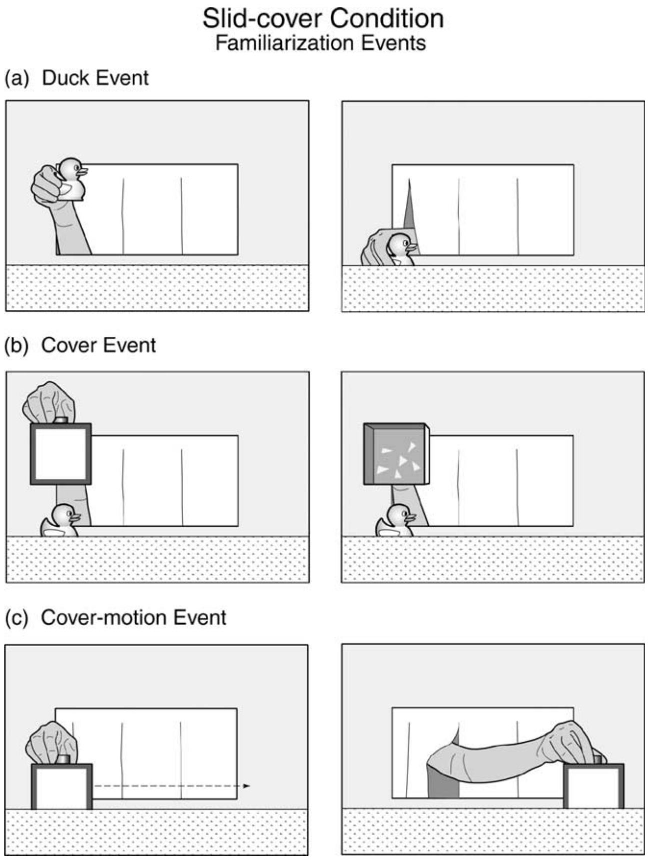

Fig. 2. Schematic drawing of the familiarization events shown in the slid-cover condition of Experiment 1.

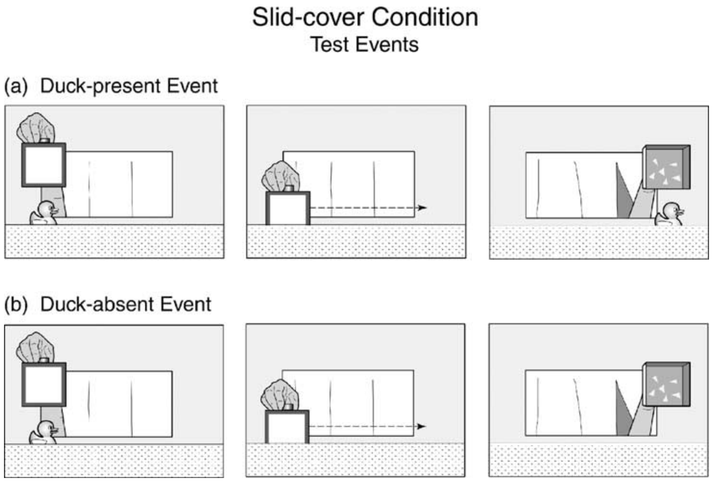

Fig. 3. Schematic drawing of the test events shown in the slid-cover condition of Experiment 1.

the cover to reveal either the duck (duck-present event) or no duck (duck-absent event). Finally, the hand rotated the cover 90° upward to show its empty interior. This entire event sequence was then repeated in reverse (the cover was rotated downward, lowered to the platform, slid to the left, and lifted), and the whole event cycle was repeated until the trial ended.

Our reasoning was as follows. If the infants represented the basic information in the test events and interpreted it in accord with their continuity principle, then they should (1) believe that the duck continued to exist after it became hidden under the cover; (2) assume that the duck was pushed by the cover along the platform; (3) expect the duck to be revealed when the cover was lifted at the right end of the platform; and (4) be surprised in the duck-absent event when this expectation was violated. The infants who saw the duck-absent event should thus look reliably longer than those who saw the duck-present event.

In the *lifted-cover* condition, a screen slightly taller than the duck hid the central portion of the platform. The infants again received four familiarization trials (see Fig. 4) followed by a single test trial in which they saw one of two test events (see Fig. 5). At the start of each test event, the cover and duck were visible to the left of the screen. 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 until it was above the right half of the screen, lowered it to the platform behind the screen, and finally slid it past the screen to the right end of the platform. As in the slid-cover condition, the hand lifted the cover to reveal either the duck (duck-present event) or no duck (duck-absent event), and then rotated the cover 90° upward to show its empty interior.

We reasoned that if the infants represented the basic information in the test events and interpreted it in accord with their continuity principle, then they should (1) believe that the duck continued to exist after it became hidden under the cover; (2) assume that the duck was pushed by the cover behind the left half of the screen; (3) infer that the duck

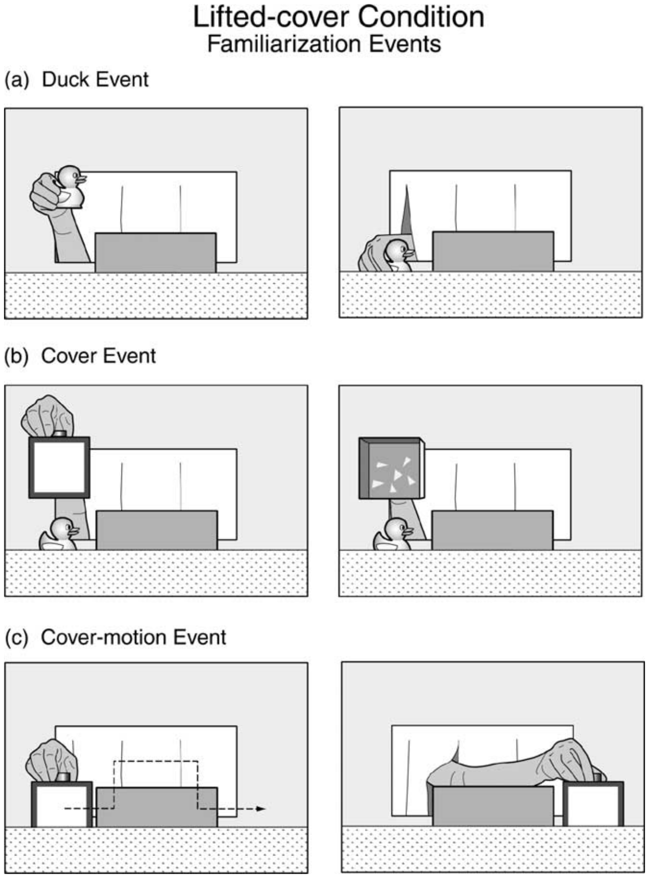

Fig. 4. Schematic drawing of the familiarization events shown in the lifted-cover condition of Experiment 1.

remained on the platform when the cover was lifted above the screen; (4) expect that no duck would be revealed when the cover was lifted at the right end of the platform; and (5) be surprised in the duck-present event when this expectation was violated. The infants who saw the duck-present event should thus look reliably longer than those who saw the duck-absent event. Opposite looking patterns were therefore predicted for the slid- and lifted-cover conditions.[5]

[5] One possible question about the design of Experiment 1 might concern our choice of a between-subjects as opposed to a within-subject design. Why not show the infants in each condition both the duck-present and duck-absent events on alternate trials? We have found in previous experiments that when very young infants are shown two complex events on successive trials, they sometimes produce differential responses to the first but not the second event (e.g. Hespos & Baillargeon, 2001b; see also Baillargeon, 1995, for a discussion of similar findings with older infants). It seems plausible that very young infants could predict the outcome of one complex event, but become either confused or fatigued when a (perceptually similar yet conceptually) different complex event is presented in the next trial. When piloting Experiment 1, it rapidly became obvious that the infants responded differentially only on the first test trial they received, so we adopted the between-subjects design described here.

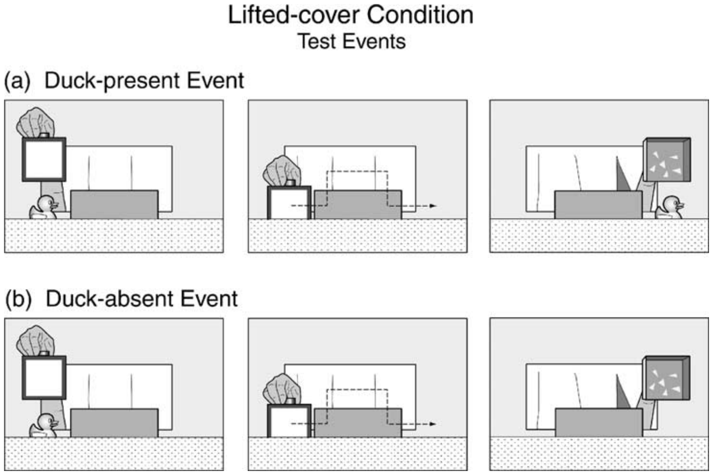

Fig. 5. Schematic drawing of the test events shown in the lifted-cover condition of Experiment 1.

2.1. Method

2.1.1. Participants

Participants were 36 healthy term infants, 18 male and 18 female, ranging in age from 2 months, 15 days to 3 months, 8 days ($M=2$ months, 28 days). An additional 22 infants were tested but eliminated, because they were fussy (10), inattentive (6), or active (3), or because they required substantial breaks for feeding or diaper change (3).[6] The infants were randomly assigned to the four experimental groups formed by crossing the two cover and the two event conditions: slid-cover/duck-absent (8 infants; $M=2$ months, 27 days), slid-cover/duck-present (10 infants; $M=2$ months, 26 days), lifted-cover/duck-absent (8 infants; $M=3$ months, 1 day), and lifted-cover/duck-present (10 infants; $M=2$ months, 29 days).

The infants' names in this and the following experiments were obtained from a file of birth announcements published in the local newspaper. Parents were contacted by letters and follow-up phone calls; they were offered reimbursement for their travel expenses but were not compensated for their participation.

2.1.2. Apparatus

The apparatus consisted of a wooden display box 106 cm high, 101 cm wide, and 35 cm deep that was mounted 76 cm above the room floor. The floor, back wall, and side walls of the apparatus were covered with pastel contact papers. Each infant faced an opening 41 cm high and 95 cm wide in the front wall of the apparatus. Between trials, a curtain consisting of a wooden frame 61 cm high, 100 cm wide, and covered with white muslin, was lowered

[6] A large proportion of eliminated participants is not uncommon in research with very young infants (e.g. Aguiar & Baillargeon, 1999, 2002; Baillargeon & DeVos, 1991; Canfield & Haith, 1991; Haith & McCarty, 1990; Hespos & Baillargeon, 2001b; Luo & Baillargeon, in press).

in front of this opening. An experimenter introduced her right hand (in a long black glove) into the apparatus through a window in the back wall. This window was 15 cm high, 30 cm wide, and filled with a white muslin fringe; it was located 20.5 cm above the apparatus floor, 34.5 cm from the left wall. Centered 15 cm above the window was a peephole 5 cm high and 13 cm wide that was used by the experimenter to monitor her actions; a flap prevented eye contact between the infant and experimenter.

A platform 16 cm high, 76 cm wide, and 31 cm deep rested on the apparatus floor against the back wall, 15 cm from the left wall. The entire platform was covered with a gray granite contact paper. In the platform's top surface were two openings, each 9 cm wide, 7 cm deep, and located 11 cm from the front of the platform. The openings were 23 cm apart, and the left opening was located 21 cm from the left edge of the platform. Running at the front and back of the openings, from the left edge of the left opening to the right edge of the right opening, were two parallel strips, 8 cm apart; each strip was made of light blue felt and was 44.5 cm long, 1.5 cm deep, and 0.2 cm thick. The cover was slid on these felt strips, smoothly and silently.

Under the platform, each opening had a set of two wooden boards hinged on its left and right sides. Each board was 12 cm wide and 9 cm deep and was covered with the same contact paper as the platform. One board in each set had a toy duck attached diagonally to its center (duck board); the other board was empty (empty board). Each board had a handle that protruded through the back of the apparatus. By means of these handles, a secondary experimenter could rotate one of the boards in each set upward 90° to close the opening. In the closed position, a board lay 0.2 cm below the top surface of the platform. The front of the platform extended 2.5 cm above its top surface to help hide this surface and the two openings from the infant's view.

Three identical toy ducks were used in the events; two were attached to boards, and the third was used in the duck familiarization event. Each duck was made of bright yellow plastic, had an orange bill and black eyes, and was 8 cm high, 6 cm wide, and 9 cm deep.

The cover used in the events was made of cardboard, cube-shaped, and 10 cm a side. Its exterior was white and its edges were outlined with red tape; its interior was green and decorated with small white triangles. A red knob 1.5 cm high and 3 cm in diameter was affixed to the top of the cover.

In the lifted-cover condition, a blue screen 9 cm high and 20 cm wide was placed on the top of the platform, against its front wall, centered between the two openings.

The infants were tested in a brightly lit room. One 150-W, one 60-W, and three 40-W lamps attached to the front and back wall of the apparatus provided additional light. Two frames, each 182 cm high, 71 cm wide, and covered with green cloth, stood at an angle on either side of the apparatus; these frames served to isolate the infants from the test room.

2.1.3. Events

Two experimenters worked in concert to produce the events: the primary experimenter wore the glove and manipulated the duck and cover; the secondary experimenter operated the boards for each opening in the platform. To help the experimenters adhere to the events' prescribed scripts, a metronome beat softly once per second. A camera mounted behind and next to the infant projected an image of the events onto a TV screen in a different part of the test room; a supervisor monitored the events to confirm that they followed the prescribed scripts. In the next sections, the numbers in parentheses indicate the time taken by the experimenters to perform the actions described.

2.1.3.1. Slid-cover condition.

Familiarization events. The infants saw three different familiarization events: a duck event, a cover event, and a cover-motion event.

At the start of the *duck event*, the primary experimenter's gloved right hand held the duck centered 11 cm above the closed opening (i.e. the empty board concealed the opening) in the left end of the platform. After a 1-s pause, the hand lowered the duck to the platform (2 s), paused for 1 s, and then lifted it back to its starting position (2 s). Each event cycle thus lasted about 6 s. For this and all other events, cycles were repeated until the trial ended (see below).

At the start of the *cover event*, the duck rested on the closed opening (i.e. the duck board concealed the opening) in the left end of the platform, and the hand held the cover centered 3 cm above the duck. After a 1-s pause, the hand rotated the cover 90° upward so that its opening faced the infant (2 s), paused for 1 s, and then returned the cover to its initial position (2 s). Each cycle thus lasted about 6 s.

At the start of the *cover-motion* event, the cover rested over the closed opening (empty board) in the left end of the platform. The hand slid the cover 32 cm to the right end of the platform (2 s), paused for 2 s, returned the cover to its original position (2 s), and then paused for another 2 s. Each event cycle thus lasted about 8 s.

Test events. The infants saw one of two test events: a duck-present or a duck-absent event.

At the start of the *duck-present* event, the duck rested on the closed opening (duck board) in the left end of the platform; the hand held the cover 3 cm above the duck. The hand lowered the cover over the duck (2 s) and then paused for 2 s, to allow the secondary experimenter to surreptitiously remove the duck (by substituting the empty board for the duck board). Next, the hand slid the cover 32 cm to the right (2 s) and again paused for 2 s, to allow the secondary experimenter to surreptitiously replace the duck (by substituting the duck board for the empty board). Finally, the hand lifted the cover above the duck (1 s) and rotated it 90° upward to show its empty interior (1 s). The entire sequence was then repeated in reverse. Each event cycle thus lasted about 20 s.

The *duck-absent* event was identical to the duck-present event except that the secondary experimenter did not replace the duck at the right end of the platform. The secondary experimenter simply opened and closed the empty board, so that whatever faint noise cues were present in the duck-present event were also present in the duck-absent event.

2.1.3.2. Lifted-cover condition.

Familiarization events. As in the slid-cover condition, the infants saw duck, cover, and cover-motion familiarization events. In all of these events, the small screen hid the center of the platform. The *duck* and *cover* events were otherwise identical to those shown in the slid-cover condition. In the *cover-motion* event, the hand slid the cover 11 cm to the right until it was hidden behind the left half of the screen (1 s). Next, the hand lifted the cover 11 cm (1 s), moved it 10 cm to the right (1 s), lowered it to the platform behind the right half of the screen (1 s), and finally slid it 11 cm to the right (1 s). After a 1-s pause, the entire sequence was repeated in reverse: the experimenter slid the cover behind the right half of the screen (1 s), lifted it above the screen (1 s), moved it to the left (1 s), lowered it behind the left half of the screen (1 s), slid it to its original position at the left end of the platform (1 s), and paused for 1 s. Each event cycle thus lasted about 12 s.

Test events. As in the slid-cover condition, the infants saw either a duck-present or a duck-absent test event. In both events, the small screen hid the center of the platform.

In the *duck-present* event, the hand lowered the cover over the duck (2 s) and then paused for 2 s, to allow the secondary experimenter to surreptitiously remove the duck. Next, the hand slid the cover behind the left half of the screen (1 s), lifted the cover (1 s), moved it to the right (1 s), and lowered it behind the right half of the screen (1 s). The hand then slid the cover to the right end of the platform (1 s) and again paused for 2 s, to allow the secondary experimenter to replace the duck. Finally, the hand lifted the cover above the duck (1 s) and rotated it 90° upward (1 s). The entire sequence was then repeated in reverse. Each event cycle thus lasted about 26 s.

The *duck-absent* event was identical to the duck-present event except that the secondary experimenter did not replace the duck at the right end of the platform. The secondary experimenter simply opened and closed the empty board to equate noise cues across the two test events.

2.1.4. Procedure

During the experiment, each infant sat on a parent's lap, centered in front of the apparatus. The infant's head was approximately 48 cm from the platform. Prior to the start of the experiment, the infant was shown, for a few seconds, the experimenter's black glove, the duck, and the cover. Parents were instructed to remain silent and neutral during the experiment and to close their eyes during the test trial.

The infant's looking behavior was monitored by two observers who viewed the infant through peepholes in the cloth-covered frames on either side of the apparatus. Each observer held a button linked to a computer and depressed the button when the infant looked at the event. The looking times recorded by the primary observer were used to determine when a trial had ended (see below).

The infants first received four familiarization trials. All the infants saw the duck event during the first trial and the cover event during the second trial. During the third and fourth trials, the infants saw the cover-motion event appropriate for their condition. Each familiarization trial ended when the infants either (1) looked away from the event for 2 consecutive seconds after having looked at it for at least 12 cumulative seconds or (2) looked for 60 cumulative seconds without looking away for 2 consecutive seconds. The 12-s minimal value was chosen to ensure that the infants had the opportunity to observe at least one cycle of each event (cycles ranged from 6 to 12 s).

Following the familiarization trials, the infants saw the test event appropriate for their cover and event condition for a single test trial. This trial ended when the infants either (1) looked away from the event for 0.5 consecutive second after having looked at it for at least 10 (slid-cover condition) or 13 (lifted-cover condition) cumulative seconds, or (2) looked for 60 cumulative seconds without looking away for 0.5 consecutive second. The 10- and 13-s minimal values were chosen to ensure that the infants had the opportunity to observe whether the duck was present or absent when the cover was lifted at the right end of the platform.

To assess interobserver agreement during the familiarization and test trials, each trial was divided into 100-ms intervals, and the computer determined in each interval whether the two observers agreed on whether the infant was or was not looking at the event. Percent agreement was calculated for each trial by dividing the number of intervals in which the observers agreed by the total number of intervals in the trial. Agreement was measured for 34 of the 36 infants (only one observer was present for two of the infants) and averaged 92% per trial per infant.

Preliminary analyses of the familiarization and test data revealed no significant interaction among cover condition, event condition, and sex, both $F\text{s}(1, 28)=0.01$; the data were therefore collapsed across sex in subsequent analyses.

2.2. Results

2.2.1. Familiarization trials

The infants' looking times during the four familiarization trials (see Fig. 6) were averaged and analyzed by means of a 2×2 analysis of variance (ANOVA) with cover condition (slid- or lifted-cover) and event condition (duck-present or duck-absent) as between-subjects factors. The interaction between cover condition and event condition was not significant, $F(1, 32)=1.27$, $p>.10$, suggesting that the infants in the four experimental groups looked about equally during the familiarization trials (slid-cover/duck-present: $M=50.5$, $SD=13.1$; slid-cover/duck-absent: $M=50.9$, $SD=17.0$; lifted-cover/duck-present: $M=53.1$, $SD=13.8$; and lifted-cover/duck-absent: $M=51.2$, $SD=14.0$). No other effect was significant.

2.2.2. Test trials

The infants' looking times during the test trial (see Fig. 6) were analyzed as above. The cover condition×event condition interaction was significant, $F(1, 32)=10.08$, $p<.005$. Planned comparisons indicated that (1) in the slid-cover condition, the infants who saw the duck-absent event ($M=43.9$, $SD=18.3$) looked reliably longer than those who saw the duck-present event ($M=24.3$, $SD=17.6$), $F(1, 32)=5.75$, $p<.025$; and (2) in the lifted-cover condition, the reverse looking pattern was found: the infants who saw the duck-present event ($M=39.8$, $SD=19.1$) looked reliably longer than those who saw the duck-absent event ($M=22.6$, $SD=12.7$), $F(1, 32)=4.38$, $p<.05$.[7]

Non-parametric Wilcoxon Rank-Sum tests confirmed the results of the slid-cover ($W_s=52.5$, $p<.01$) and lifted-cover ($W_s=52$, $p<.0025$) conditions.

[7] The test looking times of the infants in Experiment 1 were also compared by means of an analysis of covariance (ANCOVA), using the same factors as in the ANOVA and, as a covariate, the infants' mean looking times during the familiarization trials. The results of the ANCOVA were consistent with those of the ANOVA. The condition×event interaction was again significant, $F(1, 31)=8.44$, $p<.01$. Moreover, planned comparisons confirmed that (1) in the slid-cover condition, the infants who saw the duck-absent event looked reliably longer than those who saw the duck-present event, $F(1, 31)=5.82$, $p<.05$, and (2) in the lifted-cover condition, the infants produced the reverse pattern, $F(1, 31)=4.43$, $p<.05$.

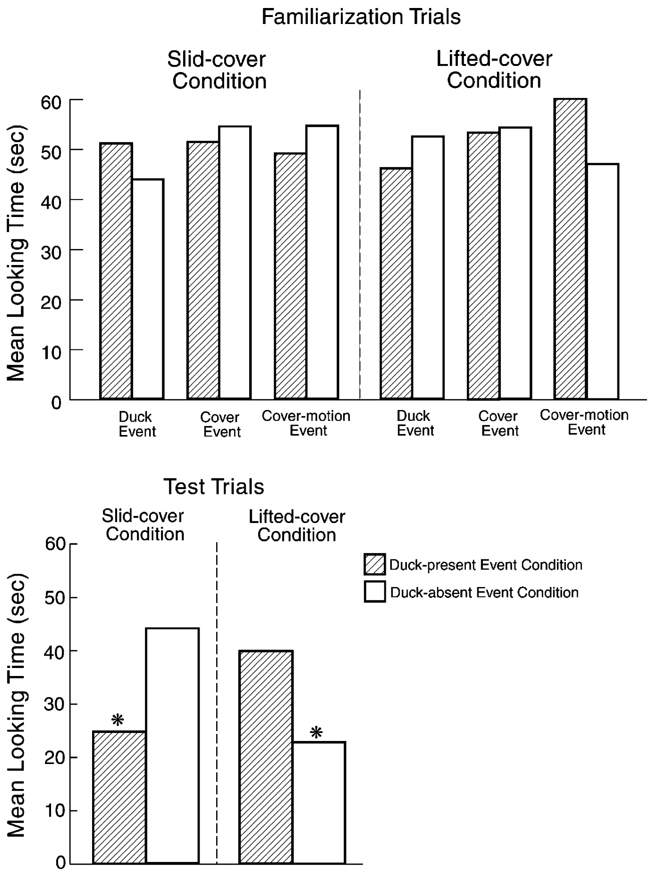

Fig. 6. Mean looking times of the infants in Experiment 1 at the familiarization and test events.

2.2.3. A perceptual-preference interpretation?

It has recently been suggested that some of the violation-of-expectation results taken to indicate rich cognitive abilities in young infants—such as the ability to represent and reason about hidden objects—might be more parsimoniously explained in terms of low-level perceptual preferences infants either bring to the laboratory or develop in the course of the familiarization or habituation trials they receive prior to the test trials (e.g. Bogartz, Shinskey, & Schilling, 2000; Bogartz, Shinskey, & Speaker, 1997; Cashon & Cohen, 2000; Rivera, Wakeley, & Langer, 1999; Roder, Bushnell, & Sasseville, 2000; Schilling, 2000; Thelen & Smith, 1994; see Wang et al., 2004, for discussion). Could the results of Experiment 1 be attributed to such preferences? It might be proposed that the infants in the slid-cover condition preferred the duck-absent event because they had formed an association between the duck and cover during the cover familiarization event, and now expected to see the duck whenever the cover was lifted above the platform. However, according to this interpretation, the infants in the lifted-cover condition, who had seen the same cover familiarization event, should also have preferred the duck-absent event, but they did not. It might also be proposed that the infants in the lifted-cover condition preferred the duck-present event because they had never seen the duck on the right side of the platform (recall that the duck was shown on the left side of the platform in the duck and cover events), and found this novel sight attractive. However, according to this interpretation, the infants in the slid-cover condition, who had seen the same duck and cover events, should also have preferred the duck-present event, but again they did not. It thus seems unlikely that low-level preferences could account for the results of Experiment 1.

2.3. Discussion

In the slid-cover condition, the infants who saw the duck-absent event looked reliably longer than those who saw the duck-present event. These results suggest that the infants (1) believed that the duck continued to exist after it became hidden under the cover; (2) assumed that the duck was pushed by the cover along the platform; (3) expected the duck to be revealed when the cover was lifted at the right end of the platform; and thus (4) were surprised in the duck-absent event when this expectation was violated.

In the lifted-cover condition, the infants who saw the duck-present event looked reliably longer than those who saw the duck-absent event. These results suggest that the infants (1) believed that the duck continued to exist after it became hidden under the cover; (2) assumed that the duck was pushed by the cover behind the left half of the screen; (3) inferred that the duck remained on the platform when the cover was lifted above the screen;[8] (4) expected that no duck would be revealed when the cover was lifted at the right end of the platform; and thus (5) were surprised in the duck-present event when this expectation was violated.

The results of Experiment 1 extend prior results by showing that 2.5- to 3-month-old infants can detect basic continuity violations in covering as well as in occlusion and containment events (e.g. Aguiar & Baillargeon, 1999; Hespos & Baillargeon, 2001b; Luo & Baillargeon, in press; Wilcox et al., 1996). As such, the present results provide support for the prediction, derived from the reasoning account presented in the Introduction, that young infants should be able to detect any continuity violation, in any event category, as long as it involves only the basic information they can represent.

The results of Experiment 1 also have implications for infants' ability to reason about invisible displacements. According to traditional accounts (e.g. Piaget, 1954), this ability does not emerge until late in the second year of life. However, the present results suggest that it may emerge much earlier. Recall that in the lifted-cover condition, the duck was left

[8] Why did the infants in the lifted-condition expect the duck to remain on the platform when the cover was lifted above the screen? At least two possibilities exist. One is that the infants noted at the end of the first event cycle that the duck remained on the left end of the platform when the cover was lifted above it. In subsequent cycles, the infants might have expected the duck to again remain on the platform when the cover was lifted above the screen. The other possibility (which we think more likely) is that the infants brought to bear their knowledge of support events: they realized that the duck could not move up with the cover. Previous research indicates that 3- and 4-month-old infants expect an object to fall when released in midair (e.g. Baillargeon, 1995; Needham & Baillargeon, 1993); perhaps 2.5- to 3-month-old infants already possess the same expectation (experiments are under way to test this possibility).

behind the screen when the cover was lifted, and was picked up again by the cover when the sequence was performed in reverse. Because the infants never observed these invisible displacements directly, they could only succeed by inferring them.

3. Experiment 2

According to the reasoning account described in the Introduction, even older infants should fail to detect a *variable* continuity 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. Experiment 2 examined 9-month-old infants' ability to detect a variable continuity violation in a covering event.

The infants received two test trials in which they saw two test events. At the start of each event, a rigid cover rested on the apparatus floor to the left of a tall object; an experimenter's gloved hand grasped the top of the cover (see Fig. 7). After a pause, the hand lifted the cover, moved it above the object, and then lowered it until it rested on the apparatus floor and thus fully hid the object. Finally, the hand released the cover, withdrew a short distance, and paused. The infants saw this paused scene—with the cover resting on the apparatus floor, hiding the object, and the hand poised a short distance away—until the trial ended. The only difference between the two test events had to do with the cover used. In the tall-cover event, the cover was slightly taller than the object; in the short-cover event, the cover was only about half as tall as the object, so that it should have been impossible for the cover to fully hide the object.

We reasoned that if the 9-month-old infants in Experiment 2 had already identified height as a covering variable, then they should include information about the relative heights of the object and cover in their physical representation of each test event, and they should interpret this information in accord with their continuity principle. As a result, the infants should expect the object to become fully hidden under the tall but not the short cover, and they should be surprised in the short-cover event when this last expectation was

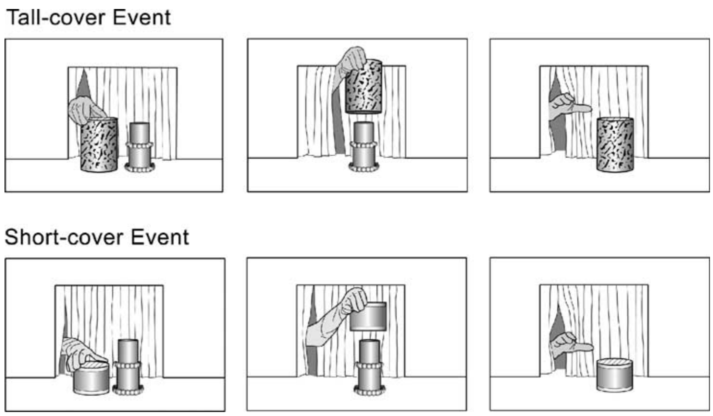

Fig. 7. Schematic drawing of the test events shown in Experiment 2.

violated. The infants should thus look reliably longer at the short- than at the tall-cover event. Conversely, if the infants had *not* yet identified height as a covering variable, then they should fail to include information about the relative heights of the object and cover when representing each test event, and they should therefore fail to notice that the object was too tall to become fully hidden under the short cover. The infants should thus look about equally at the short- and tall-cover events.

A final comment about the design of Experiment 2 might be in order. As noted above, the infants were given no familiarization or habituation trials: they received only test trials. In this we followed the example of Hespos and Baillargeon (2001a), who obtained positive results with 7.5-month-old infants in a violation-of-expectation task with test trials only. There were two main reasons for giving the infants in Experiment 2 only test trials. First, as shown by Hespos and Baillargeon, familiarization or habituation trials do not appear to be necessary when older infants are presented with simple events involving an experimenter's hand lowering an object behind an occluder, inside a container, and so on. If familiarization or habituation trials are not needed to help infants orient to a task, there is obviously no reason to provide them; indeed, without such trials, infants may be less likely to be bored or inattentive during the test trials. Second, as was mentioned earlier, some researchers have suggested that infants tested with violation-of-expectation tasks involving hidden objects may look reliably longer at the unexpected than at the expected events, not because they can represent and reason about hidden objects, but because the familiarization or habituation trials induce in them transient preferences for the unexpected events (e.g. Bogartz et al., 2000; Cashon & Cohen, 2000; Roder et al., 2000; Schilling, 2000; Thelen & Smith, 1994). Evidence that infants also succeed in tasks with test trials only can help alleviate these concerns (for additional evidence and discussion, see Wang et al., 2004).

3.1. Method

3.1.1. Participants

Participants were 12 healthy term infants, 5 male and 7 female, ranging in age from 8 months, 9 days to 9 months, 25 days ($M = 9$ months, 0 day). One additional infant was tested but eliminated because of fussiness.

3.1.2. Apparatus

The apparatus used in Experiment 2 was the same as in Experiment 1, except that the floor and back wall were changed. The new floor had a hole 7.5 cm in diameter located 56.5 cm from the left wall and 31 cm from the front of the apparatus. The new back wall stood 50 cm from the front of the apparatus and had a window 36 cm high and 42 cm wide extending from its lower edge, 32 cm from the left wall. As in Experiment 1, the window was filled with a white muslin fringe. An experimenter introduced her right hand (in a yellow rubber glove) into the apparatus through this window.

The two covers used in the events were made of plastic piping material 0.5 cm thick, were 11.5 cm in diameter, and were closed at the top with construction paper. The short cover was 8.5 cm tall; it was covered with green contact paper and decorated with horizontal yellow stripes at the top and bottom. The tall cover was 17 cm tall and was covered with gray granite contact paper.

The object used in the events was collapsible and composed of an outer and an inner cylinder made of cardboard 0.2 cm thick. The *outer* cylinder was 6.5 cm in diameter and 16 cm high. The lower half of the cylinder was not visible to the infants: it protruded through the hole in the apparatus floor. The upper half of the cylinder was covered with red contact paper and was decorated at the top and bottom with fuzzy bright orange pompoms. The top pompoms were about 1 cm in diameter, and the bottom pompoms about 1.5 cm; the bottom pompoms hid the hole in the apparatus floor and also prevented the outer cylinder from falling through the hole. The *inner* cylinder was closed at the top and was 6 cm in diameter and 15 cm tall. The bottom 8 cm portion of the cylinder was not visible to the infants: it stood hidden inside the outer cylinder and was covered with thin sponges. These sponges filled the space between the inner and outer cylinders and thus prevented the inner cylinder from falling through the outer cylinder (and the hole in the apparatus floor). The top 7 cm portion of the cylinder was again covered with red contact paper. The pompoms at the top of the outer cylinder hid the small gap between the inner and outer cylinders.

At the start of each test event, the object was 15 cm tall (7 cm inner cylinder and 8 cm outer cylinder). In the tall-cover event, the object remained fully extended under the tall cover. In the short-cover event, the inner cylinder was pressed downward by the interior top surface of the short cover as it was lowered over the object; the inner cylinder glided silently and smoothly inside the outer cylinder, making it possible for the object to become fully hidden, as in the tall-cover event.

Finally, a non-collapsible replica of the object, 15 cm tall, was shown to the infants prior to the experiment (see below).

3.1.3. Events

A single experimenter produced the test events. To help the experimenter adhere to the events' prescribed scripts, a metronome beat softly once per second. A camera mounted behind and next to the infant projected an image of the events onto a TV screen in a different part of the test room; a supervisor monitored the events to confirm that they followed the prescribed scripts.

Tall-cover event. At the start of the tall-cover event, the experimenter's gloved right hand grasped the top of the tall cover, which rested on the apparatus floor 42 cm from the left wall and 29 cm from the front of the apparatus; the object stood centered 2 cm to the right of the cover. The experimenter maintained her position until the computer signaled that the infant had looked at the display for 2 cumulative seconds. This pretrial gave the infants the opportunity to inspect the cover and object and to compare their heights.

Following the pretrial, the tall-cover event proper began; this event was composed of two distinct phases, an initial and a final phase. During the *initial* phase, the experimenter lifted the cover 17.5 cm (1 s) and moved it to the right 12.5 cm until it was centered 2.5 cm above the object (1 s). Next, the hand lowered the cover to the apparatus floor, thus fully hiding the object (2 s). Finally, the hand released the cover, moved about 12 cm to the left of the cover (1 s), and paused. The initial phase thus lasted about 5 s. During the *final* phase, the infants saw the same paused scene—with the cover resting on the apparatus floor and the hand poised a short distance away—until the trial ended (see below).

Short-cover event. The short-cover event was identical to the tall-cover event except that the short cover was substituted for the tall one.

3.1.4. Procedure

As in Experiment 1, each infant sat on a parent's lap centered in front of the apparatus. The infant's head was approximately 85 cm from the cover. Prior to the experiment, the experimenter knelt next to the parent's chair and showed the infant her gloved right hand and the non-collapsible object. To make clear to the infant that the object was not collapsible, the experimenter stood the object on the palm of her left hand and tapped its top firmly a few times with her right hand. Finally, the experimenter showed each cover, one at a time, calling the infant's attention to its closed top and hollow interior.

During the experiment, the infants saw the tall- and short-cover test events on two successive trials. Half of the infants saw the short-cover event on the first trial and the tall-cover event on the second trial; the other infants saw the two events in the reverse order. Looking times during the initial and final phases of each trial were computed separately. The infants' mean looking time during the initial, 5-s phase of each trial (beginning after the pretrial, when the experimenter lifted the cover), was 4.9 s; the infants thus looked almost continuously during the initial phase. The final phase of each trial ended when the infant either (1) looked away from the paused scene for 2 consecutive seconds after having looked at it for at least 3 cumulative seconds, or (2) looked at the paused scene for 60 cumulative seconds without looking away for 2 consecutive seconds. Interobserver agreement during the final phase of each test trial was calculated for all 12 infants and averaged 95% per trial per infant.

Preliminary analyses of the test data revealed no significant interaction involving sex and event, both $Fs(1, 8) < 1.19$, $p > .10$; the data were therefore collapsed across sex in subsequent analyses.

3.2. Results

The infants' looking times during the final phase of each test trial (see Fig. 8) were analyzed by means of a 2×2 ANOVA with order (short- or tall-cover event first) as a between-subjects factor and event (short- or tall-cover) as a within-subject factor. The main effect of event was not significant, $F(1, 10) = 0.31$, indicating that the infants tended to look equally at the short- ($M = 20.9$, $SD = 10.0$) and tall-cover ($M = 19.4$, $SD = 9.8$) events. This result was confirmed by a non-parametric Wilcoxon Signed-Ranks test, $T = 52$, $p > .10$.

3.2.1. Additional results: Experiment 2A

Why did the 9-month-old infants in Experiment 2 look about equally at the two test events? There were at least two possible interpretations for this negative result. A first interpretation, consistent with the reasoning account presented in the Introduction, was that the infants had *not* yet identified height as a variable in covering events, and hence did not include information about the relative heights of the cover and object in their physical

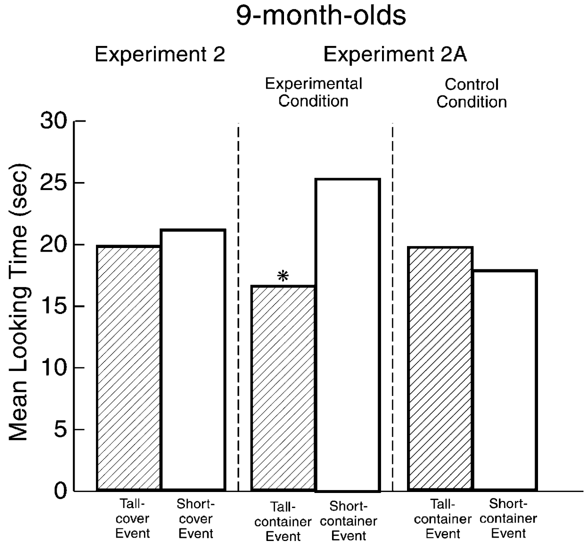

Fig. 8. Mean looking times of the infants in Experiments 2 and 2A at the test events.

representations of the test events. As a result, the infants failed to detect the continuity violation in the short-cover event: they did not notice that the object was too tall to become fully hidden under the short cover. This interpretation suggests that infants regard covering events as distinct from containment events, and do not generalize variables from containment to covering events: recall that 7.5-month-old infants attend to height information in containment events, and are surprised when a tall object becomes almost fully hidden inside a short container (Hespos & Baillargeon, 2001a).

However, another interpretation for the negative results of Experiment 2 was that the infants *had* identified height as a covering variable, but some aspect of our procedure made it difficult for them to reveal this knowledge. In Hespos and Baillargeon (2001a), for example, the event in each test trial was shown repeatedly until the trial ended; in Experiment 2, in contrast, each event was shown only once per trial. Perhaps our infants failed to detect the violation in the short-cover event because they did not have sufficient opportunity to observe this violation, or because they rapidly became bored with the paused scene in the final phase of the trial and so did not process it appropriately.

To examine this alternative interpretation, an additional group of 9-month-olds was tested in Experiment 2A using a procedure similar to that of Experiment 2 with two exceptions: the covers were turned upside-down and used as containers, and the object was lowered into the short or tall container until it became fully hidden. For half of the infants (experimental condition), the object was the tall, non-collapsible object used in Experiment 2 (see Fig. 9); for the other infants (control condition), the object was shorter and could be fully hidden inside either container.

We reasoned that if the infants in Experiment 2 failed because at 9 months infants have identified the variable height in containment but not in covering events, then the results of

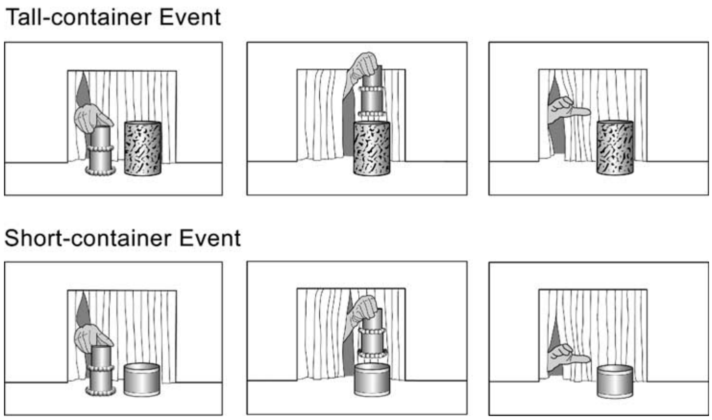

Fig. 9. Schematic drawing of the test events shown in the experimental condition of Experiment 2A.

Experiment 2A should replicate those of Hespos and Baillargeon (2001a): the infants in the experimental condition should look reliably longer at the short- than at the tall-container test event, whereas those in the control condition should look about equally at the two events. On the other hand, if the infants in Experiment 2 failed because our procedure was less optimal than that used by Hespos and Baillargeon, then the infants in the experimental and control conditions of Experiment 2A should all tend to look equally at the short- and tall-container events.

Participants were 24 9-month-old infants, 10 male and 14 female, ranging in age from 8 months, 7 days to 10 months, 1 day ($M=9$ months, 3 days). Two additional infants were tested but eliminated, one because a sneeze ended one of the infant's test trials, and the other because the infant showed an extreme bias (of over 40 s) for one of the test events. Half of the infants were assigned to the experimental condition ($M=9$ months, 3 days), and half to the control condition ($M=9$ months, 3 days). The apparatus and stimuli used in Experiment 2A were similar to those in Experiment 2 with the following exceptions. First, the tall and short covers were turned upside-down to serve as containers. Second, a new floor was used that had a hole 10 rather than 7.5 cm in diameter. Third, in the experimental condition, the tall object was used throughout the experiment (the bottom of the object was 9.5 cm in diameter, with its pompoms, making the larger hole necessary). Each container was placed over the hole, 2 cm to the right of the object. In the tall-container event, the object was simply lowered inside the tall container; it rested 2 cm below the top of the container and was not visible to the infants (recall that the object was 15 cm tall and the tall container 17 cm tall). In the short-container event, the bottom of the container was removed and the object was lowered through the hole in the apparatus floor until it again rested 2 cm below the top of the container and was no longer visible to the infants; a platform beneath the hole ensured that the object was always lowered to the same depth. Fourth, in the control condition, the tall, non-collapsible object was replaced with a similar but shorter object that could be fully hidden inside either container; this short object was 6.5 cm tall, 5.5 cm in diameter, covered with red contact paper, and decorated at the bottom with orange pompoms about 1.5 cm in diameter. In each event, the container (with its bottom intact) was placed over the hole, and the short object was simply lowered inside the container.

Prior to the experiment, the experimenter first showed the infant her gloved right hand and the tall or short object; she then showed each container, one at a time, calling the infant's attention to its hollow interior and closed bottom. Next, the infants saw the short- and tall-container events appropriate for their condition on two successive trials; half of the infants in each condition saw the short-container event first, and half saw the tall-container event first. As in Experiment 2, each trial began with a pretrial that allowed the infants to inspect the object and container and to compare their heights. When the computer signaled that the infant had looked for 2 cumulative seconds, the initial phase of the trial began. The experimenter's hand lifted the object (1 s) and moved it to the right 12.5 cm until it was centered 2.5 cm above the short or tall container (1 s). Next, the hand lowered the object until it became fully hidden inside the container (2 s). Finally, the hand released the object, moved about 12 cm to the left of the container (1 s), and paused. The initial phase thus lasted about 5 s, as in Experiment 2. The infants' mean looking time during this phase was 4.9 s, indicating that they tended to look continuously.

During the final phase of each test trial, the infants saw the paused scene—with the object hidden inside the short or tall container and the hand poised a short distance away—until the trial ended. The criteria for ending the trials were the same as in Experiment 2. Interobserver agreement was calculated for all 24 infants and averaged 95% per test trial per infant. Preliminary analyses of the test data revealed no significant interaction involving sex and event, all Fs$(1, 16) < 1.73$, $p > .10$; the data were therefore collapsed across sex in subsequent analyses.

The infants' looking times during the final phase of each test trial (see Fig. 8) were analyzed by means of a $2 \times 2 \times 2$ ANOVA with condition (experimental or control) and order (short- or tall-container event first) as between-subjects factors and event (short- or tall-container) as a within-subject factor. The analysis revealed a significant condition $\times$ event interaction, $F(1, 20) = 5.42$, $p < .05$. Planned comparisons indicated that the infants in the experimental condition looked reliably longer at the short- ($M = 25.2$, $SD = 16.3$) than at the tall-container ($M = 16.2$, $SD = 12.5$) event, $F(1, 20) = 6.67$, $p < .025$, whereas those in the control condition tended to look equally at the two events (short-container event: $M = 17.4$, $SD = 9.3$; tall-container event: $M = 19.9$, $SD = 10.7$), $F(1, 20) = 0.50$. Non-parametric Wilcoxon Signed-Ranks tests confirmed the results of the experimental ($T = 78$, $p < .0005$) and control ($T = 44.5$, $p > .10$) conditions.

In a final analysis, the looking times of the infants in Experiment 2 and in the experimental condition of Experiment 2A were compared by means of a $2 \times 2 \times 2$ ANOVA with experiment (2 or 2A) and order (short- or tall-cover/container event first) as between-subjects factors and event (short- or tall-cover/container) as a within-subject factor. The analysis revealed a significant main effect of event, $F(1, 20) = 10.05$, $p < .005$, and a significant experiment $\times$ event interaction, $F(1, 20) = 5.37$, $p < .05$. Planned comparisons confirmed that, whereas the infants in Experiment 2 tended to look equally at the short- and tall-cover events, $F(1, 20) = 0.36$, those in the experimental condition of Experiment 2A looked reliably longer at the short- than at the tall-container event, $F(1, 20) = 15.07$, $p < .001$.

3.3. Discussion

The infants in the experimental condition of Experiment 2A looked reliably longer at the short- than at the tall-container test event, whereas those in the control condition tended to look equally at the two events. Together, these results suggest that the infants (1) realized that the short object could become fully hidden inside either container, and that the tall object could become fully hidden inside the tall but not the short container, and (2) were surprised when this last expectation was violated. These results indicate that, at 9 months of age, infants include information about the relative heights of objects and containers in their physical representations of containment events. The results of Experiment 2A are thus consistent with those obtained by Hespos and Baillargeon (2001a) with 7.5-month-old infants.

In contrast to the infants in the experimental condition of Experiment 2A, those in Experiment 2 tended to look equally at the short- and tall-cover test events. Why did they do so? It does not seem likely that the infants were confused or misled by the procedure used in Experiment 2, since the infants in Experiment 2A, who were tested with a very similar procedure, succeeded in detecting the variable continuity violation they were shown. Nor does it seem likely that infants generally have more difficulty reasoning about covers than containers: recall that the 2.5- to 3-month-old infants in Experiment 1 were presented with covering events and readily detected the basic continuity violations in these events (see also Baillargeon et al., 1990; Leslie, 1995, for related findings with 5- to 6-month-old infants). Rather, it seems more likely, as suggested by the reasoning account presented in the Introduction, that the 9-month-old infants in Experiment 2 had not yet identified height as a variable relevant to covering events. As a result, the infants (1) did not include information about the relative heights of the cover and object when representing each test event; (2) could not interpret this (missing) information in accord with their continuity principle; and thus (3) failed to notice that the object was too tall to become fully hidden under the short cover.

At what age do infants typically identify height as a covering variable? To address this question, 11- and 12-month-old infants were tested in Experiment 3 using the same procedure as in Experiment 2.

4. Experiment 3

4.1. Method

4.1.1. Participants

Participants were 24 healthy term infants. There were 12 11-month-old infants, 6 male and 6 female, ranging in age from 11 months, 1 day to 11 months, 19 days ($M = 11$ months, 11 days), and 12 12-month-old infants, 6 male and 6 female, ranging in age from 11 months, 20 days to 12 months, 25 days ($M = 12$ months, 5 days). Three additional infants were tested but eliminated because they showed an extreme bias (of over 40 s) for one of the test events.

4.1.2. Apparatus, events, and procedure

The apparatus, events, and procedure used in Experiment 3 were identical to those in Experiment 2. The 11- and 12-month-old infants' mean looking times during the initial phase of each test trial were 5 and 4.9 s, respectively. Interobserver agreement during the final phase of each test trial was calculated for all 24 infants and averaged 95% per trial per infant. Preliminary analyses of the test data revealed no significant interaction involving sex and event, all $Fs(1, 16) < 1.17$, $p > .10$; the data were therefore collapsed across sex in subsequent analyses.

4.2. Results

The infants' looking times during the final phase of each test trial (see Fig. 10) were analyzed by means of a $2 \times 2 \times 2$ ANOVA with age (11 or 12 months) and order (short- or tall-cover event first) as between-subjects factors and with event (short- or tall-cover) as a within-subject factor. The analysis yielded a significant main effect of event, $F(1, 20) = 6.01$, $p < .025$, and a significant age × event interaction, $F(1, 20) = 5.79$, $p < .05$. Planned comparisons indicated that (1) the 12-month-old infants looked reliably longer at the short- ($M = 26.0$, $SD = 12.7$) than at the tall-cover ($M = 14.4$, $SD = 6.8$) event, $F(1, 20) = 11.81$, $p < .005$, but (2) the 11-month-old infants tended to look equally at the two events, $F(1, 20) = .00$ (short-cover event: $M = 17.5$, $SD = 12.3$; tall-cover event: $M = 17.4$, $SD = 13.0$). Non-parametric Wilcoxon Signed-Ranks tests confirmed the results obtained with the 12-month-olds ($T = 66$, $p < .05$) and 11-month-olds ($T = 40$, $p > .10$).

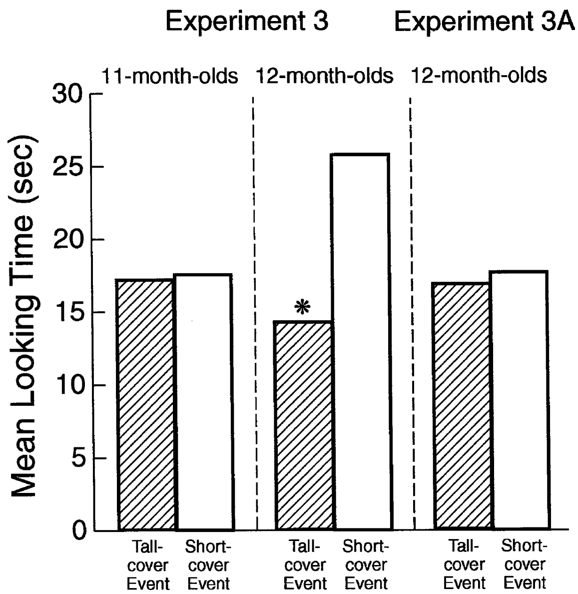

Fig. 10. Mean looking times of the infants in Experiments 3 and 3A at the test events.

The ANOVA also revealed a significant order × event interaction, $F(1, 20) = 20.02$, $p < .00025$. Examination of the data suggested that, overall, the infants in each order condition tended to look longer at whichever test event they saw first. Because this effect did not interact with age (the age × order × event interaction was not significant, $F(1, 20) = 0.07$), it does not bear on the present issues and will not be discussed further.

4.2.1. Additional results: Experiment 3A

The 12-month-old infants in Experiment 3 detected the variable continuity violation in the short-cover test event; in contrast, the 11-month-old infants, like the 9-month-old infants in Experiment 2, did not. Together, these results suggest that, at about 12 months of age, infants identify height as a covering variable, and begin including information about the relative heights of covers and objects in their physical representations of covering events. However, there was another, less interesting interpretation of the responses of the 12-month-old infants in Experiment 3: perhaps these older infants simply found the short cover intrinsically more attractive than the tall cover.

To examine this alternative, perceptual-preference interpretation, an additional group of 12-month-olds was tested in Experiment 3A using the same procedure as in Experiment 3 with one exception: a shorter object was used that could be fully hidden under either the short or the tall cover. We reasoned that if the 12-month-old infants in Experiment 3 simply preferred the short over the tall cover, then the infants in Experiment 3A should show the same preference and again look reliably longer at the short- than at the tall-cover event. On the other hand, if the 12-month-old infants in Experiment 3 realized that the tall object could be fully hidden under the tall but not the short cover, then the infants in Experiment 3A should recognize that the short object could be fully hidden under either cover, and they should thus look about equally at the two test events.[9]

Participants were 12 healthy term infants, 6 male and 6 female, ranging in age from 11 months, 22 days to 12 months, 23 days ($M = 12$ months, 10 days). One additional infant was tested but eliminated, because he showed an extreme bias (of over 40 s) for one of the test events. The apparatus, events, and procedure used in Experiment 3A were similar to those in Experiment 3 with two exceptions. First, the tall object was replaced with a similar but shorter object that could be fully hidden under either cover; this object was closed at the top, 6.5 cm in diameter, 8 cm tall, covered with red contact paper, and decorated at the bottom with orange pompoms about 1.5 cm in diameter. Second, the hole in the apparatus floor was filled with an Insert and the short object was placed over it and completely hid it. The short and tall covers were simply lowered over the short object in the test events.

[9] Experiment 3A also addressed another perceptual-preference interpretation of the results obtained with the 12-month-old infants in Experiment 3. According to this interpretation, the infants looked reliably longer at the short- than at the tall-cover event because they preferred seeing, in the initial phase of each test event, a cover and object that were different (short-cover event) as opposed to similar (tall-cover event) in height. Because the object used in Experiment 3A was similar in height to the short but not the tall cover, this alternative interpretation predicted that the infants should now look reliably longer at the tall- than at the short-cover event. Evidence that the infants tended to look equally at the events would thus also help rule out this interpretation.

The infants' mean looking time during the initial phase of each test trial was 4.8 s. Interobserver agreement during the final phase of each test trial was calculated for all 12 infants and averaged 95% per trial per infant. Preliminary analyses of the test data revealed no significant interaction involving sex and event, both $Fs(1, 8) < 1.52$, $p > .10$; the data were therefore collapsed across sex in subsequent analyses.

The infants' looking times during the final phase of each test trial (see Fig. 10) were analyzed as in Experiment 2. The main effect of event was not significant, $F(1, 10) = 0.02$, indicating that the infants tended to look equally at the short- ($M = 17.7$, $SD = 12.2$) and tall-cover ($M = 17.2$, $SD = 11.2$) events. A non-parametric Wilcoxon Signed-Ranks test confirmed this result, $T = 41$, $p > .10$.

In a further analysis, the test responses of the 12-month-old infants in Experiments 3 and 3A were compared by means of a $2 \times 2 \times 2$ ANOVA with experiment (Experiment 3 or 3A) and order (short- or tall-cover event first) as between-subjects factors and with event (short- or tall-cover) as a within-subject factor. The analysis yielded a significant main effect of event, $F(1, 20) = 5.71$, $p < .05$, and a significant experiment × event interaction, $F(1, 20) = 4.84$, $p < .05$. Planned comparisons confirmed that the 12-month-old infants in Experiment 3 looked reliably longer at the short- than at the tall-cover event, $F(1, 20) = 10.54$, $p < .005$, but that those in Experiment 3A did not, $F(1, 20) = 0.02$.

The ANOVA also revealed a significant order × event interaction, $F(1, 20) = 6.70$, $p < .025$. Examination of the data suggested that, overall, the infants who saw the short-cover event first looked longer at this event than at the tall-cover event, whereas those who saw the tall-cover event first tended to look equally at the two events. Because this effect did not interact with experiment (the experiment × order × event interaction was not significant, $F(1, 20) = 1.95$, $p > .10$), it has no bearing on the issues examined here and will not be discussed further.

4.3. Discussion

The 12-month-old infants in Experiment 3 looked reliably longer at the short- than at the tall-cover event, but those in Experiment 3A looked about equally at the two events. These results suggest that the infants (1) realized that the short object could be hidden under either cover, and that the tall object could be hidden under the tall but not the short cover, and (2) were surprised when this last expectation was violated.

Unlike the 12-month-old infants in Experiment 3, the 9- and 11-month-old infants in Experiments 2 and 3 tended to look equally at the short- and tall cover events. Together, these results suggest that infants identify height as a covering variable at about 12 months of age. When watching a covering event, infants now include information about the relative heights of the cover and object in their physical representation of the event; this variable information becomes subject to their continuity principle, making it possible for them to detect continuity violations involving short covers and tall objects.

McCall (2001) reported evidence consistent with the finding that infants attend to height information in covering events at 12 but not 9 months of age. In a manual search task, 12- and 9-month-old infants first repeatedly searched for a tall toy under a tall cover of a specific color (e.g. blue). Next, the infants were asked to search for the same tall toy under one of two covers: a short cover of the same color as before, and a tall cover of a different color than before (e.g. red). Only the 12-month-old infants consistently searched for the tall toy under the tall cover.

4.3.1. An explanation-based learning account

According to the reasoning account, infants who have identified a variable in an event category should succeed in detecting continuity violations involving the variable; in contrast, infants who have not yet identified the variable should fail to detect such violations. But what process leads infants to identify variables? And why should infants identify the same variable at different ages in different event categories? In particular, why should infants identify the variable height at about 7.5 months in containment events (Hespos & Baillargeon, 2001a), but only at about 12 months in covering events (Experiments 2 and 3)?[10] Why is there such a marked lag or décalage in infants' identification of the variable height in these two event categories?

One explanation is suggested by a recent account of the process by which infants typically identify a new variable in an event category. This process is thought to be one of *explanation-based learning* (EBL) and to involve three main steps (e.g. Baillargeon, 2002, 2004; Wang & Baillargeon, 2005b; 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 core principles (e.g. the continuity principle dictates 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 accepted 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 suggests a possible reason for why infants may identify height as a containment variable several months before they identify it as a covering variable (for a related discussion of the décalage in infants' identification of the variable height in occlusion and containment events, see Baillargeon, 2002; Hespos & Baillargeon, 2001a; Wang et al., 2004). This reason has to do with the first step described above, exposure to appropriate contrastive outcomes. It may be that, in everyday life, infants have more opportunities to notice that objects inside containers sometimes are fully and sometimes partly hidden, than to notice that objects placed under covers sometimes are fully

[10] We are not claiming that *every* infant identifies height first as a containment and only later as a covering variable. It may well be that some infants identify the two variables at about the same time, or in the reverse order. All that the available data indicate is that (1) most infants have identified height as a containment variable by 7.5 months of age, and (2) most infants have identified height as a covering variable by 12 months of age. For an experiment such as Experiment 2A or 3 to yield a positive result, the majority of the infants in the experiment must be able to detect the variable continuity violation in the experiment.

and sometimes partly hidden. Because exposure to appropriate contrastive outcomes is necessary to trigger the process of learning, it follows that variables will be learned later when exposure is less frequent.[11]

The preceding analysis leads to an interesting prediction: infants who are exposed in the laboratory to appropriate outcome and condition data for a variable should be able to identify this variable earlier than they would otherwise. To test this prediction, we recently conducted a series of experiments in which we attempted to "teach" 9-month-old infants the variable height in covering events (Wang & Baillargeon, 2005b). Our results were positive and as such support both the EBL account and the speculation above that the décalage in infants' identification of the variable height in containment and covering events stems from the fact that infants are typically exposed to appropriate observations for this variable at different ages in the two categories.

4.3.2. Alternative explanations

It might be suggested that there exist other, very different explanations for the décalage in infants' responses to height violations in containment and covering events.[12] For example, it might be that when asked to reason about events involving an object-to-be-hidden and a hiding-object, infants generally perform better when it is the object-to-be-hidden rather than the hiding-object that moves. Thus, infants are better at reasoning about events in which an object is lowered inside a container, as in Experiment 2A, as opposed to events in which a cover is lowered over an object, as in Experiments 2 and 3.

Another possible explanation is that when watching an object being lowered inside a container, or a cover being lowered over an object, infants tend to track whichever object is moving and to focus on its lower portion (i.e. the portion that leads the motion), even if it becomes invisible. This tracking strategy is helpful for detecting a continuity violation that involves a tall object being lowered inside a short container (the bottom of the object will appear to pass through the apparatus floor); however, it is less helpful for detecting a violation that involves a short cover being lowered over a tall object (the bottom of the cover will appear to simply come to rest against the apparatus floor).

Do infants detect height violations in containment events earlier than in covering events (1) because the variable height is generally identified earlier in containment than in covering events, as was suggested in the last section, or (2) because height violations are generally easier to detect in containment than in covering events, as suggested above? We speculated that one way to address this question might be to present infants with a height

---

[11] Another reason for the present décalage has to do with the second step in the EBL account, mapping the appropriate conditions onto the observed outcomes. To identify the variable height in containment and covering events, infants must compare the heights of objects and containers or covers. Prior research (e.g. Baillargeon, 1991, 1994, 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 containers or covers qualitatively, infants must compare them as they stand *side by side*. It may again be that in everyday life infants have more opportunities to perform such comparisons with containment than with covering events (e.g. there may be more instances in which objects standing next to containers are next lowered inside the containers, than instances in which covers standing next to objects are next lowered over the objects).

[12] We are indebted to an anonymous reviewer for pointing out these alternative interpretations.

violation in a *tube* event—an event in which a tall object is lowered inside a short tube until it becomes fully hidden.

Our reasoning was as follows. If infants detect height violations later in covering than in containment events because they are less often exposed to appropriate observations for identifying covering variables, then they might also detect height violations later in tube than in containment events. After all, it is likely that infants are exposed to fewer tube than containment events, and so again have fewer opportunities to gather appropriate observations for identifying tube variables. On the other hand, if infants detect height violations later in covering than in containment events because violations in which the hiding-object moves are generally harder to process than are violations in which the object-to-be-hidden moves, then infants should succeed at about the same age at detecting violations in containment and tube events. Experiment 4 examined 12- to 14-month-old infants' ability to detect a height violation in a tube event.

5. Experiment 4

Prior research suggests that young infants distinguish between tubes and containers, and can detect a *basic* continuity violation in a tube event (e.g. Kolstad and Baillargeon, cited in Baillargeon, 1995). In one experiment, 5.5-month-old infants were surprised when salt poured into a tube, but not a container, remained inside it (i.e. they expected the salt to fall out of the bottom of the tube). Experiment 4 examined 12-, 13-, and 14-month-old infants' ability to detect a *variable* continuity violation in a tube event.

The infants were tested with the same events and procedure as in the experimental condition of Experiment 2A, with one exception. Prior to the experiment, the bottoms of the short and tall containers were removed to form short and tall tubes. As before, the experimenter first showed the infants her gloved hand and the tall, non-collapsible object; she then showed each tube, one at a time, calling the infants' attention to its hollow interior and open top and bottom. From that point on, the infants never again saw the tubes' open bottoms: the short- and tall-tube test events were produced in the same manner as, and were perceptually identical to, the short- and tall-container events in Experiment 2A (see Fig. 9). The infants could only "know" that they were facing tubes rather than containers by remembering their visual inspection of the tubes prior to the experiment.

We reasoned that evidence that the younger infants in Experiment 4 tended to look equally at the short- and tall-tube test events would point to several conclusions. First, it would suggest that infants view containment and tube events as separate event categories, and do not generalize variables identified in one category to the other (e.g. Baillargeon, 2002; Hespos & Baillargeon, 2001a; Baillargeon & Wang, 2002). Second, such evidence would provide additional support for the reasoning account presented in the Introduction, and more specifically for the notion that infants typically cannot detect continuity violations that involve variables they have not yet identified. Third, such evidence would be consistent with the EBL claim that the same variable may be identified at different ages in different event categories, depending on the age at which infants are exposed to appropriate condition and outcome observations for identifying the variable in each category. By the same token, such evidence would also cast doubt on the alternative explanations described in the last section for the décalage in infants' identification of the variable height in containment and covering events. Finding that 12-month-old infants are *not* surprised when a tall object is lowered inside a short tube until it becomes fully hidden, even though they *are* surprised when a short cover is lowered over a tall object until it becomes fully hidden (Experiment 3), would argue against processing hypotheses that appeal to the salience of different object motions or to various tracking strategies. Finally, evidence that infants succeed in detecting a variable continuity violation when they believe that the objects before them are *containers* (Experiment 2A), but not *tubes* (Experiment 4), would provide strong evidence against perceptual-preference interpretations of the present findings.

5.1. Method

5.1.1. Participants

Participants were 40 healthy term infants. There were 12 12-month-old infants (11 months, 28 days to 12 months, 24 days, $M = 12$ months, 9 days); 14 13-month-old infants (12 months, 29 days to 14 months, 1 day, $M = 13$ months, 14 days), and 14 14-month-old infants (14 months, 2 days to 14 months, 23 days, $M = 14$ months, 11 days). Each age group had an equal number of male and female infants. Four additional infants were tested but eliminated, 2 because the observers had difficulty following the infants' gaze, 1 because of fussiness, and 1 because the infant showed an extreme bias (of over 40 s) for one of the test events.

5.1.2. Apparatus, events, and procedure

The apparatus, events, and procedure used in Experiment 4 were similar to those in the experimental condition of Experiment 2A, except that that the bottoms of the tall and short containers were removed prior to the experiment. The experimenter first showed each infant her gloved right hand and the tall, non-collapsible object; she then showed each tube, one at a time, calling the infant's attention to its hollow interior and open top and bottom. Next, the bottom of the tall container was replaced (out of the infants' view). The tall- and short-tube events were identical to the tall- and short-container events in the experimental condition of Experiment 2A.

In all three age groups, the infants' mean looking time during the initial phase of each test trial was 5 s. Interobserver agreement during the final phase of each test trial was calculated for all 40 infants and averaged 95% per trial per infant. Preliminary analyses of the test data revealed no significant interaction involving sex and event, all $F\text{s} < 1.73$, $p > .10$; the data were therefore collapsed across sex in subsequent analyses.

5.2. Results

The infants' looking times (see Fig. 11) were analyzed by means of a $3 \times 2 \times 2$ ANOVA with age (12, 13, or 14 months) and order (short- or tall-tube event first) as between-subjects factors and with event (short- or tall-tube) as a within-subject factor. The analysis yielded no significant effects, all $F\text{s} < 3.21$, $p > .05$. Planned comparisons indicated that (1) the 12-month-old infants tended to look equally at the short- ($M = 18.2$, $SD = 12.4$) and

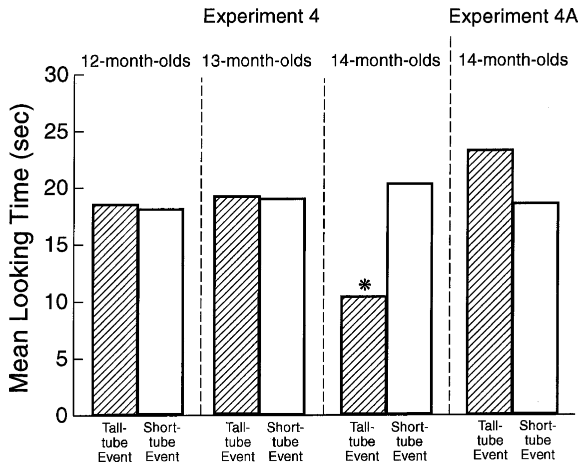

Fig. 11. Mean looking times of the infants in Experiments 4 and 4A at the test events.

tall-tube ($M=18.7$, $SD=16.8$) events, $F(1, 34)=0.01$; (2) the 13-month-old infants also looked about equally at the short- ($M=19.1$, $SD=15.0$) and tall-tube ($M=19.4$, $SD=9.4$) events, $F(1, 34)=0.01$; but (3) the 14-month-old infants looked reliably longer at the short- ($M=20.5$, $SD=10.1$) than at the tall-tube ($M=10.5$, $SD=5.3$) event, $F(1, 34)=6.86$, $p<.025$. Non-parametric Wilcoxon Signed-Ranks tests confirmed the results obtained with the 12-month-olds ($T=44$, $p>.10$), 13-month-olds ($T=56.5$, $p>.10$), and 14-month-olds ($T=96$, $p<.005$).

Further results. A group of 9-month-old infants was also tested with the short- and tall-tube events used in Experiment 4. Participants were 16 healthy term infants, 8 male and 8 female, ranging in age from 8 months, 11 days to 9 months, 23 days ($M=8$ months, 26 days). The infants' mean looking time during the initial phase of each test trial was 4.8 s. Interobserver agreement for all 16 infants was 95% per trial per infant. The infants' responses were compared to those of the 9-month-old infants in the experimental condition of Experiment 2A, who saw the (perceptually identical) short- and tall-container test events. Preliminary analyses of the test data revealed no significant interaction involving sex and event, all $Fs(1, 20)<3.85$, $p>.05$; the data were therefore collapsed across sex in subsequent analyses. The ANOVA yielded a significant experiment×event interaction, $F(1, 24)=5.18$, $p<.05$. Planned comparisons confirmed that, whereas the infants in Experiment 2A looked reliably longer at the short- ($M=25.2$, $SD=16.3$) than at the tall-container ($M=16.2$, $SD=12.5$) event, $F(1, 24)=6.66$, $p<.025$, those in the present experiment tended to look equally at the short- ($M=18.7$, $SD=11.2$) and tall-tube ($M=20.2$, $SD=13.2$) events, $F(1, 24)=0.25$ (a non-parametric Wilcoxon Signed-Ranks test supported this result, $T=75$, $p>.10$). The infants thus succeeded in detecting the variable violation they were shown when they believed that the objects before them were containers, but not tubes.

5.2.1. Additional results: Experiment 4A

The 14-month-old infants in Experiment 4 looked reliably longer at the short- than at the tall-tube event, but the 9-, 12-, and 13-month-old infants tended to look equally at the two events. These results suggest that it is not until infants are about 14 months of age that they identify height as a tube variable. Infants then begin to include information about the relative heights of objects and tubes in their physical representations of tube events, which allows them to detect violations involving tall objects and short tubes.

To rule out an alternative, perceptual-preference interpretation of the positive results obtained with the 14-month-old infants, an additional group of 14-month-olds was tested in Experiment 4A using the same procedure as in Experiment 4, except that a shorter object was used that could be fully hidden inside either tube; this short object was identical to that used in the control condition of Experiment 2A.

Participants were 14 healthy term infants, 7 male and 7 female, ranging in age from 14 months, 2 days to 14 months, 23 days ($M=14$ months, 9 days). The infants' mean looking time during the initial phase of each test trial was 5 s. Interobserver agreement during the final phase of each test trial was calculated for all 14 infants and averaged 96% per trial per infant. Preliminary analyses of the test data revealed no significant interaction involving sex and event, both $F\text{s}(1, 10)<0.11$; the data were therefore collapsed across sex in subsequent analyses.

The infants' looking times during the final phase of each test trial (see Fig. 11) were analyzed by means of a 2×2 ANOVA with order (short- or tall-tube event first) as a between-subjects factor and event (short- or tall-tube) as a within-subject factor. The main effect of event was not significant, $F(1, 12)=0.86$, indicating that the infants tended to look equally at the short- ($M=18.8$, $SD=15.1$) and tall-tube ($M=23.5$, $SD=15.5$) events. A non-parametric Wilcoxon Signed-Ranks test confirmed this result, $T=67$, $p>.10$.

In a final analysis, the looking times of the 14-month-old infants in Experiments 4 and 4A were compared by means of a $2\times2\times2$ ANOVA with experiment (4 or 4A) and order (short- or tall-tube event first) as between-subjects factors and with event (short- or tall-tube) as a within-subject factor. The analysis yielded a significant experiment$\times$event interaction, $F(1, 24)=6.74$, $p<.025$. Planned comparisons confirmed that the 14-month-old infants in Experiment 4 looked reliably longer at the short- than at the tall-tube event, $F(1, 24)=6.23$, $p<.025$, but that those in Experiment 4A did not, $F(1, 24)=1.37$, $p>.10$.

5.3. Discussion

The 14-month-old infants in Experiment 4 looked reliably longer at the short- than at the tall-tube event, but those in Experiment 4A did not. These results suggest that the infants (1) recognized that the short object could be hidden inside either tube, and that the tall object could be hidden inside the tall but not the short tube, and (2) were surprised when this last expectation was violated.

Unlike the 14-month-old infants in Experiment 4, the 9-, 12-, and 13-month-old infants tended to look equally at the short- and tall-tube events. These results suggest that infants identify height as a tube variable at about 14 months of age. When watching an object being lowered inside a tube, infants now include information about the relative heights of the object and tube in their physical representation of the event; this variable information then becomes subject to their continuity principle, making it possible for them to detect continuity violations involving tall objects and short tubes. Younger infants, who have not yet identified height as a tube variable, typically do not include height information in their physical representations of tube events, and hence cannot detect continuity violations involving this variable.

Together, the results of Experiments 2A, 4, and 4A suggest three conclusions. First, infants appear to view events involving containers and tubes as belonging to distinct event categories, and learn separately about each category: infants do not generalize variables identified in the context of containment events to tube events. When describing our account of physical reasoning in infancy (Section 1.1), we suggested that the basic information infants include in their physical representations of events includes simple information about the geometry of objects, and more particularly about their open/closed surfaces. From this perspective, it would make sense that infants would distinguish from an early age between objects that are open at the top (containers), at the bottom (covers), or both (tubes).

Second, the present results provide strong support for our reasoning account, and more specifically for the claim that whether infants detect a variable continuity violation in an event category depends primarily on whether they have identified the variable involved as relevant to the category. In general, only infants who have identified height as a tube variable will include information about the relative heights of the object and tube in their physical representations of tube events, and hence will be able to detect a continuity violation in which a tall object becomes fully hidden inside a short tube.

Third, the present results also provide support for the EBL account described earlier (Section 4.3.1). According to this account, the age at which infants identify a variable in an event category depends mainly on the age at which infants are exposed to appropriate condition and outcome observations from which to abstract the variable. From this perspective, the evidence that infants identify the variable height at about 7.5 months in containment events (Hespos & Baillargeon, 2001a), but only at about 12 months in covering events (Experiments 2, 3, and 3A), and even later, at about 14 months in tube events (Experiments 4 and 4A), simply reflects the fact that in their daily lives infants are generally exposed to more containment than covering or tube events. As noted earlier, this analysis predicts that infants exposed at a younger age to appropriate observations for the variable height in covering or in tube events should identify these variables earlier; and indeed, findings from experiments in which we attempted to “teach” the variable height in covering events to 9-month-old infants support this prediction (Wang & Baillargeon, 2005b)

5.3.1. Additional décalages

Hespos and Baillargeon (2001a) uncovered a striking discrepancy in 4.5-month-old infants’ responses to occlusion and containment events: they found that 4.5-month-old infants are surprised when a tall object becomes almost fully hidden behind but not inside a short container. The results of Experiments 2 and 2A revealed another discrepancy: 9-month-old infants are surprised when a tall object becomes fully hidden inside a short container, but not when the short container, turned upside-down to form a short cover, is lowered over the same tall object until it becomes fully hidden.

One explanation for these findings, consistent with the reasoning and EBL accounts presented in this article, is as follows: when infants happen to identify a variable in one event category several weeks or months before they identify the same variable in a different category, a marked décalage can be observed in their responses to similar events from the two categories. However, other explanations are possible. Because the events from the two categories differ (e.g. the object is lowered behind or inside the container; the object is lowered inside the container or the cover over the object), one could always appeal to subtle differences between the events to explain why infants reason correctly about one event but not the other. We saw examples of such alternative explanations earlier (Section 4.3.2), when discussing why 9-month-old infants can detect a height continuity violation in a containment but not a covering event. For example, one explanation was that infants may be generally better at reasoning about events in which the object-to-be-hidden, rather than the hiding-object, moves.

The results of Experiments 2A and 4 provide evidence against such alternative explanations: the 9-month-old infants in these experiments were surprised when a tall object became fully hidden inside what they believed to be a short container (Experiment 2A); however, they were *not* surprised when the same tall object became fully hidden inside what they believed to be a short tube (Experiment 4). In fact, the infants in the two experiments saw exactly the same short-container/tube event, produced in exactly the same manner. The only difference between the experiments was that the infants in Experiment 2A were shown a short and a tall container prior to the experiment, whereas those in Experiment 4 were shown a short and a tall tube. Because the infants in the two experiments saw exactly the same events, explanations based on subtle differences between the events are not possible in this case.

By the same token, the results of Experiments 2A and 4 also argue against low-level perceptual-preference explanations of the present and related findings. As was mentioned earlier, a few researchers have suggested that findings from violation-of-expectation tasks may reflect only perceptual preferences infants either bring to the laboratory or develop in the course of familiarization or habituation trials (e.g. Bogartz et al., 1997; Cashon & Cohen, 2000; Rivera et al., 1999; Schilling, 2000; Thelen & Smith, 1994; see Wang et al., 2004, for discussion). The fact that the infants in Experiments 2A and 4 succeeded in detecting the height continuity violation they were shown when they believed that the objects before them were containers, but not tubes, provides evidence against low-level interpretations of the present and related findings (e.g. Experiments 2, 2A, and 3; Hespos & Baillargeon, 2001a; Luo & Baillargeon, 2005).

6. General discussion

We began this article with examples of early and late successes in infants' ability to detect continuity violations in occlusion and containment events (e.g. Aguiar & Baillargeon, 1999, 2002; Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a,b; Luo & Baillargeon, in press, 2004; Wilcox, 1999; Wilcox & Chapa, 2004; Wilcox et al., 1996). The present research brought to light several additional examples involving other event categories. On the one hand, the 2.5- to 3-month-old infants in Experiment 1 *were* surprised (1) when a cover was lowered over an object, slid to the right, and lifted to reveal no object; and (2) when a cover was lowered over an object, slid behind the left half of a screen, lifted above the screen, moved to the right, lowered behind the right half of the screen, slid past the screen, and finally lifted to reveal the object. On the other hand, (1) the 9- and 11-month-old infants in Experiments 2 and 3 were *not* surprised when a short cover was lowered over a tall object until it became fully hidden–only the 12-month-old infants detected this violation; and (2) the 9-, 12-, and 13-month-old infants in Experiment 4 were *not* surprised when a tall object was lowered inside a short tube until it became fully hidden–only the 14-month-old infants detected this violation.

The reasoning account presented in the Introduction helps make sense of all of these results (see also Baillargeon, 1987, 1991, 1993, 1995; Newcombe et al., 1999; Spelke et al., 1992). The account distinguishes between two kinds of continuity violations: basic and variable violations. A *basic* violation in an event can be detected as long as the basic information about the event is adequately represented (this basic information becomes subject to infants' continuity principle, and the event is marked as violating the principle). Because young infants are often able to represent the basic information about events, the account predicts that they should succeed in detecting many basic continuity violations. Specifically, they should succeed in detecting any continuity violation that involves only the basic information they can represent.

A *variable* violation in an event can be detected as long as the information about the relevant variable in the event is represented (this variable information again becomes subject to infants' continuity principle, and the event is tagged as violating the principle). Because infants typically do not include information about a variable in their physical representation of an event from a category until they have identified the variable as relevant to the category, the account predicts that infants should fail to detect many variable continuity violations. Specifically, they should fail to detect any continuity violation that involves a variable they have not yet identified in an event category and hence do not yet include when representing events from the category

According to the account, infants' *early* successes at detecting continuity violations in occlusion, containment, covering, and other events thus typically involve *basic* violations, and their *later* successes *variable* violations.

6.1. New research directions

In order to confirm and extend our reasoning account, we have undertaken several new lines of research. Two are briefly described below; both focus on variable information. The *first* involves change blindness experiments and seeks to test more directly the claim that infants who fail to detect a variable continuity violation in an event typically do not include information about the variable in their physical representation of the event. The *second* line of research involves experiments that attempt to induce infants, through various contextual manipulations, to include information about a variable they have not yet identified when representing an event. According to the account, this variable information, once represented, should become subject to infants' continuity principle, allowing them to detect violations involving the variable earlier than they would otherwise.

6.1.1. Change blindness in infants

The 12-month-old infants in Experiment 3 looked reliably longer at the short- than at the tall-cover event, suggesting that they detected the continuity violation in the short-cover event. In contrast, the 11-month-old infants tended to look equally at the two events. According to the reasoning account presented in this article, the older infants had identified the variable height as relevant to covering events and hence included information about the relative heights of the cover and object in their physical representations of the short- and tall-cover events. In contrast, the younger infants had not yet identified height as a covering variable and hence did not include height information when representing the events.

Is this analysis correct? Are we right in assuming that the older but not the younger infants in Experiment 3 included information about the relative heights of the cover and object in their physical representations of the test events? To get at this question, we have recently undertaken a new series of experiments (Wang & Baillargeon, 2005a), inspired in part by the research on change detection and change blindness in the adult perception literature (e.g. Mitroff, Simons, & Levin, in press; Rensink, 2002; Simons, 1996, 2000). We reasoned that if infants include information about a variable in their physical representation of an event, then they should be able to detect surreptitious changes involving the variable. Conversely, if infants do not include information about a variable when representing an event, then they should be unable to detect surreptitious changes involving the variable: in other words, they should be *blind* to such changes.

To illustrate, in one experiment, 11-month-old infants were assigned to a covering or an occlusion condition. 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 tall cover; the experiment tested not whether infants could judge what object could 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.

Because 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 (Experiment 3), we predicted that the 11-month-old infants in the occlusion condition would detect the violation in the change event, but that those in the covering condition would not. 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. Thus, as predicted by our reasoning account, the 11-month-old infants in the covering condition were *blind* to the surreptitious change in the height of the object, suggesting that they did not include information about the relative heights of the object and cover in their physical representations of the test events. In a subsequent experiment, 12-month-old infants were tested in the covering condition;

as expected, these older infants detected the violation in the change event, suggesting that they did include height information when representing the test events.

Before leaving this section, we would like to offer one important caveat. We suppose that when infants watch a physical event, different computational systems form different representations simultaneously, for distinct purposes. In particular, infants' *object-recognition* system represents detailed information about the objects in the event, for recognition and categorization purposes. At the same time, infants' *physical-reasoning* system forms a physical representation of the event, to monitor it as it unfolds and to interpret and predict its outcome. Our intuition is that, when infants realize that they must include information about a variable in their physical representation of an event, they typically access the appropriate representations in their object-recognition system and retrieve the necessary information. On this view, the 9-, 11-, and 12-month-old infants in Experiments 2 and 3 *all* encoded information about the relative heights of the cover and object in their object-recognition system. However, *only* the 12-month-old infants accessed this system to retrieve this height information and include it in their physical representations of the short- and tall-cover events.

The preceding speculations suggest that even infants who fail to include variable information in their representation of an event could nevertheless have this information available in their object-recognition system. Tasks designed to tap this system directly should thus reveal this knowledge. This means, for example, that even 9-month-old infants should detect a change in the relative heights of a cover and object when given a task that taps their object-recognition rather than their physical-reasoning system. Designing such tasks should prove an interesting challenge.

6.1.2. Inducing infants to represent variable information

According to the reasoning account presented in this article, the 9- and 11-month-old infants in Experiments 2 and 3 failed to detect the continuity violation in the short-cover event because (1) they had not identified height as a covering variable; (2) they did not include information about the relative heights of the cover and object in their physical representation of the event; and (3) this (missing) information could not be interpreted in accord with their continuity principle. This analysis suggested an intriguing possibility: if infants could be *induced*, through some contextual manipulation, to include the relevant height information in their physical representation of the short-cover event, this information would then become subject to their continuity principle, making it possible for them to detect the violation in the event.

Prior research. To our knowledge, there have been two reports in the physical reasoning literature suggesting that infants can be induced to include key information in their physical representations of events; both reports made use of *priming* trials to highlight this information (Kotovsky, Mangione, and Baillargeon, cited in Baillargeon, 1995; Wilcox & Chapa, 2004). In one experiment, Kotovsky et al. showed 6.5- and 9.5-month-old infants an expected and an unexpected event. At the start of the unexpected event, a screen lay flat on the apparatus floor, away from the infants. An experimenter's gloved hand grasped the back right corner of the screen and rotated it forward through a 180° arc until it rested on the apparatus floor, toward the infants. As the screen was rotated forward, a large toy clown was revealed, standing on the apparatus floor behind the screen.

The expected event was similar except that the screen rested against the (hidden) clown at the start of the event, so that it was possible for the clown to be revealed when the screen was rotated forward. In each test event, the infants saw a single rotation: after rotating the screen forward, the experimenter paused until the trial ended. At both ages, the infants tended to look equally at the unexpected and expected events. One interpretation of this negative result was that, when representing the unexpected event, the infants did not include information about the orientation of the screen at the start of the event. As a result, they did not detect the continuity violation in the event: they did not realize that the clown could not have stood under the screen as it lay flat against the apparatus floor. In a subsequent experiment, Kotovsky and her colleagues attempted to prime the infants to attend to the orientation of the screen at the start of each event. Prior to the test trials, the infants saw two static displays with two screens standing side by side: one lay flat against the apparatus floor, away from the infants, and the other rested against the (hidden) clown. Both the 6.5- and the 9.5-month-old infants looked reliably longer at the unexpected than at the expected event, suggesting that simultaneously seeing the two different orientations of the screens had primed the infants to include information, in their physical representation of each test event, about the orientation of the screen at the start of the event. This information then became subject to the infants' continuity principle, allowing them to detect the violation in the unexpected event.

Wilcox and Chapa (2004) also attempted to prime infants to attend to key information in an event. Their research built on earlier findings (Wilcox, 1999) that at 9.5 months of age infants attend to shape but not color information in occlusion events: they are surprised when a green ball and a green box appear successively from behind a screen that is too narrow to hide them both simultaneously; however, they are *not* surprised when a green and a red ball appear successively from behind the same narrow screen. One interpretation of this negative result is that infants do not include information about the color of the balls in their physical representation of the event. As a result, they cannot detect the continuity violation in the event: they do not realize that there are two balls present, which cannot hide behind the narrow screen at the same time. Wilcox and Chapa set out to prime 9.5-month-old infants to attend to the color information in their narrow-screen event. Prior to the test trials, the infants received two pairs of priming trials. In each pair, the infants saw a pound event, in which an experimenter used a green cup to pound a peg, and a pour event, in which an experimenter used a red cup to pour salt. Different green and red cups were used in the two pairs of trials. Next, the infants saw a test event in which a green and a red ball appeared successively from behind a narrow (narrow-screen event) or a wide (wide-screen event) screen. Wilcox and Chapa reasoned that if the infants learned during the priming trials that color had functional significance in the situation (since it predicted which action the experimenter would perform with each cup), they might be more likely to include color information when representing the test events. The infants who saw the narrow-screen event looked reliably longer than those who saw the wide-screen event, suggesting that the priming trials had induced the infants to include color information in their physical representation of each test event. Additional experiments revealed that (1) the priming effect just described was eliminated when the infants were shown the same green and red cups on both pairs of priming trials; (2) 7.5-month-old infants could also be primed to include color information, but required three pairs of priming trials (with three different pairs of green and red cups) to do so; (3) 5.5-month-old infants could be primed to include pattern information, and again required three pairs of priming trials (with three different pairs of dotted and striped green cups) to do so; and finally (4) 4.5-month-old infants could also be primed to include pattern information, but both cups had to be present in each priming trial to allow simultaneous comparison of their patterns.

Our approach. In a recent series of experiments, we attempted, using a very different approach, to induce 8.5-month-old infants to include height information in their physical representations of covering events (Wang & Baillargeon, in press). Infants saw the same test events as in Experiments 2 and 3, with one exception. Prior to each test trial, the infants watched a pretrial intended to induce them to attend to the relative heights of the cover and object. In designing these pretrials, we took advantage of the fact that, although height is typically not identified until about 12 months in covering events (Experiments 2 and 3), it is identified much earlier, at about 3.5 months, in occlusion events (e.g. Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a; Baillargeon & Graber, 1987). In each pretrial, the cover was slid in *front* of the object, to create an occlusion event; after a few seconds, the cover was slid back to its original position, and then the trial proceeded as in Experiments 2 and 3, with the cover being lowered over the object.

We reasoned that the infants would categorize the event shown in each pretrial as an occlusion event, and then would access their knowledge of this event category; because at 8.5 months of age this knowledge includes the variable height, the infants would encode information about the relative heights of the cover and object in their representation of the event. We speculated that this information might still be available, or might again be encoded, when the infants next saw and represented the covering event. The information would then become subject to the infants' continuity principle, allowing them to detect the violation in the short-cover event.

The infants looked reliably longer at the short- than at the tall-cover event, suggesting that the occlusion pretrial did lead them to include height information in their physical representations of the test events. This conclusion was supported by a control experiment with a different pretrial. In this experiment, the cover was simply slid forward, next to the object, so that the infants saw a display rather than an occlusion event. The infants failed to detect the continuity violation in the short-cover event, suggesting that they did not include information about the heights of the cover and object in their representation of either the display or the covering event.

Additional remarks. The three series of experiments discussed in this section not only support our reasoning account, but also raise many exciting questions for future research. It seems likely that these experiments made use of very different mechanisms to induce infants to include information about variables in their physical representations; specifying the nature of these mechanisms should yield important insights about the nature of infants' physical reasoning.

As an example, consider the experiments we just described (Wang & Baillargeon, in press). One explanation for the induction effect observed in these experiments might be as follows. When infants see a sequence of two events from two distinct event categories, involving the same or similar objects, (1) they assign "indexes" to the objects in the first event (e.g. Leslie, Xu, Tremoulet, & Scholl, 1998; Scholl and Leslie, 1999; for related models of visual attention in adults, see, e.g. Kahneman, Treisman, & Gibbs, 1992;

Pylyshyn, 1989, 1994), and bind to these indexes information about the variables that have been identified as relevant to the event's category; and (2) they carry over (or keep on using) the same object indexes when representing the second event. This carry-over of object indexes with their bound variable information would explain why seeing first an occlusion and then a covering event with a short cover and a tall object would enhance 8.5-month-old infants' reasoning about the covering event: the height information bound to the object indexes in the occlusion event would enable the infants to detect the variable continuity violation in the covering event.

6.2. Origins of the continuity principle

In all of our discussions so far, we have assumed, following Spelke and her colleagues (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke et al., 1992, 1995) that infants' principle of continuity is innate. Could one interpret the findings discussed in this article without this assumption? We think not, for the following reasons.

Consider the results of Experiment 1. It does not seem likely that 2.5- to 3-month-old infants would have repeated opportunities to observe various covering events, and to learn to associate each event with its outcome. But could infants (1) develop a general expectation about continuity from observing more common events, and then (2) apply this expectation to other relevant events, such as covering events? As they watch their caretakers act on objects, young infants often observe occlusion and containment events. Could infants possess a learning mechanism capable of abstracting from such events a general expectation of continuity? Logically, it would certainly seem possible to design such a learning mechanism. The difficulty is to explain why infants, if they possess such a mechanism, fail to use it again after the first few months of life.

As we discussed earlier, when learning about event categories, infants acquire not event-general but event-specific expectations: variables identified in one event category are not generalized to other relevant categories. Thus, infants must learn separately about the variable height in occlusion, in containment, and, as shown in Experiments 2–4, in covering and in tube events (e.g. Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a). But if infants possess a learning mechanism capable of forming general expectations and applying them to all relevant events, how can we explain these piecemeal acquisitions? These and related speculations (see Baillargeon, 2002, 2004) led us to adopt the suggestion (e.g. Carey & Spelke, 1994; Spelke, 1994; Spelke et al., 1992, 1995) that infants' physical representations are interpreted from birth in accord with a continuity principle. The reasoning account proposed here builds on this suggestion, and helps make sense of infants' early and late successes in detecting continuity violations.

Acknowledgements

This research was supported by a dissertation fellowship from the Chiang Ching-kuo Foundation and by Faculty Research Funds from the University of California, Santa Cruz to the first author, and by a grant from NICHD (HD-21104) to the second author. We thank Jerry DeJong, Cindy Fisher, Gergely Csibra and Kris Onishi for helpful comments; Deepa Block, Laura Brueckner, Laura Glaser, Yuyan Luo, Venessa Nolen, Shana Sappington, Laurie Tondini, and the undergraduate assistants in the Infant Cognition Laboratory at the University of Illinois for their help with the data collection; and the parents and infants who participated in the research.

References

Aguiar, A., & Baillargeon, R. (1999). 2.5-month-old infants' reasoning about when objects should and should not be occluded. *Cognitive Psychology*, *39*, 116–157.

Aguiar, A., & Baillargeon, R. (2002). Developments in young infants' reasoning about occluded objects. *Cognitive Psychology*, *45*, 267–336.

Aguiar, A., & Baillargeon, R. (2003). Perseverative responding in a violation-of-expectation task in 6.5-month-old infants. *Cognition*, *88*, 277–316.

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

Baillargeon, R. (1991). Reasoning about the height and location of a hidden object in 4.5- and 6.5-month-old infants. *Cognition*, *38*, 13–42.

Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. E. Granrud (Ed.), *Visual perception and cognition in infancy* (pp. 265–315). Hillsdale, NJ: Erlbaum.

Baillargeon, R. (1994). How do infants learn about the physical world? *Current Directions in Psychological Science*, *3*, 133–140.

Baillargeon, R. (1995). A model of physical reasoning in infancy. In C. Rovee-Collier, & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 9) (pp. 305–371). Norwood, NJ: Ablex.

Baillargeon, R. (1998). Infants' understanding of the physical world. In M. Sabourin, & F. Craik (Eds.), *Advances in psychological science* (Vol. 2) (pp. 503–529). London, England: Psychology Press, 503–529.

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, England: Blackwell.

Baillargeon, R. (2004). Infants' reasoning about hidden objects: Evidence for event-general and event-specific expectations. *Developmental Science*, *7*, 391–424.

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

Baillargeon, R., & Graber, M. (1987). Where's the rabbit? 5.5-month-old infants' representation of the height of a hidden object. *Cognitive Development*, *2*, 375–392.

Baillargeon, R., Graber, M., DeVos, J., & Black, J. (1990). Why do young infants fail to search for hidden objects? *Cognition*, *36*, 225–284.

Baillargeon, R., & 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 × Event Set design. *Developmental Psychology*, *33*, 408–422.

Canfield, R. L., & Haith, M. M. (1991). Young infants' visual experiences for symmetric and asymmetric stimulus sequences. *Developmental Psychology*, *27*, 198–208.

Carey, S., & Spelke, E. (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. (2002). Infant categorization of containment, support, and tight-fit spatial relationships. *Developmental Science*, *5*, 247–264.

Casasola, M., Cohen, L. B., & 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' perception of possible and impossible events. *Infancy*, *1*, 429–446.

Dan, N., Omori, T., & Tomiyasu, Y. (2000). Development of infants' intuitions about support relations: Sensitivity to stability. *Developmental Science*, *3*, 171–180.

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., & McCarty, M. E. (1990). Stability of visual expectation at 3.0 months of age. *Developmental Psychology*, *26*, 68–74.

Held, R., Thorn, F., Gwiazda, J., & Bauer, J. (1996). Development of binocularity and its sex differentiation. In F. Vital-Durand, & J. Atkinson (Eds.), *The European Brain and Behaviour Society series* (Vol. 2) (pp. 265–274). Cambridge, MA: MIT Press.

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*, 207–245.

Huettel, S. A., & Needham, A. (2000). Effects of balance relations between objects on infants' object segregation. *Developmental Science*, *3*, 415–427.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, *24*, 174–219.

Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 234–262). Oxford, England: 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.

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

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, England: Clarendon Press.

Leslie, A. M., Xu, F., Tremoulet, P. D., & Scholl, B. J. (1998). Indexing and the object concept: developing 'what' and 'where' system. *Trends in Cognitive Sciences*, *2*, 10–18.

Luo, Y., & Baillargeon, R. (2005). *Infants' reasoning about events involving transparent occluders and containers*. Manuscript in preparation.

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

McCall, D. (2001, April). *Perseveration and infants' sensitivity to cues for containment*. Paper presented at the biennial meeting of the Society for Research in Child Development, Minneapolis, MN.

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

Mitroff, S.R., Simons, D.J., & Levin, D.T. (in press). Nothing compares two views: Change blindness can occur despite preserved access to the changed information. *Perception and Psychophysics*.

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., & Baillargeon, R. (1993). Intuitions about support in 4.5-month-old infants. *Cognition*, *47*, 121–148.

Newcombe, N., Huttenlocher, J., & Learmonth, A. (1999). Infant's coding of location in continuous space. *Infants Behavior and Development*, *22*, 483–510.

Pauen, S. (1999). The development of ontological categories: Stable dimensions and changing concepts. In W. Schnotz, S. Vosniadou, & M. Carretero (Eds.), *New perspectives on conceptual change* (pp. 15–31). Amsterdam: Elsevier.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.

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

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

Quinn, P., & Eimas, P. (1996). Perceptual organization and categorization in young infants. In C. Rovee-Collier, & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 10) (pp. 1–36). Westport, CT: Ablex.

Rensink, R. A. (2002). Change detection. *Annual Review of Psychology*, *53*, 245–277.

Rivera, S. M., Wakeley, A., & Langer, J. (1999). The drawbridge phenomenon: Representational reasoning or perceptual preference? *Developmental Psychology*, *35*, 427–435.

Roder, B. J., Bushnell, E. W., & Sasseville, A. M. (2000). Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy*, *1*, 491–507.

Schilling, T. H. (2000). Infants' looking at possible and impossible screen rotations: The role of familiarization. *Infancy*, *1*, 389–402.

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

Simons, D. J. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, *7*, 301–305.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, *7*, 1–15.

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. (2001). Continuity, competence, and the object concept. In E. Dupoux (Ed.), *Language, brain, and cognitive development: Essays in honor of Jacques Mehler* (pp. 325–340). Cambridge, MA: The MIT Press.

Spelke, E. S., Phillips, A., & Woodward, A. L. (1995). 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, England: 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. (2005a). *Change blindness in infancy: Event-specific effects*. Manuscript in preparation.

Wang, S., & Baillargeon, R. (in press). *Inducing infants to detect continuity violations: A new approach.* Manuscript under review.

Wang, S., & Baillargeon, R. (2005b). *Teaching infants 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., 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., & Chapa, C. (2002). Infants' reasoning about opaque and transparent occluders in an 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 full-term infants. *Infant Behavior and Development*, *19*, 309–323.

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

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.