

Blackwell Handbook of Childhood Cognitive Development

Edited by

Usha Goswami

© 2002 by Blackwell Publishers Ltd
a Blackwell Publishing company
except for editorial arrangement and introduction © 2002 by Usha Goswami

350 Main Street, Malden, MA 02148-5018, USA
108 Cowley Road, Oxford OX4 1JF, UK
550 Swanston Street, Carlton, Victoria 3053, Australia
Kurfürstendamm 57, 10707 Berlin, Germany

The right of Usha Goswami to be identified as the Author of the Editorial Material in this Work has been asserted in accordance with the UK Copyright, Designs, and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

First published 2002 by Blackwell Publishers Ltd

Library of Congress Cataloging-in-Publication Data

Blackwell handbook of childhood cognitive development / edited by Usha Goswami.

p. cm.—(Blackwell handbooks of developmental psychology)

Includes bibliographical references and index.

ISBN 0-631-21840-8

1. Cognition in children. I. Goswami, Usha. II. Series.

BF723.C5 B5 2002

155.4'13—dc21

2002023054

A catalogue record for this title is available from the British Library.

Set in 10.5 on 12.5 pt Adobe Garamond
by SNP Best-set Typesetter Ltd., Hong Kong
Printed and bound in the United Kingdom by
TJ International, Padstow, Cornwall

For further information on
Blackwell Publishing, visit our website:
<http://www.blackwellpublishing.com>

CHAPTER THREE

The Acquisition of Physical Knowledge in Infancy: A Summary in Eight Lessons

Renée Baillargeon

As adults we possess a great deal of knowledge about the physical world. For example, we realize that an object continues to exist when placed behind a nearer object, that a wide object can be lowered inside a wide but not a narrow container, and that an object typically falls when released in midair. Piaget (1954) was the first researcher to examine whether infants, like adults, hold expectations about physical events. Analyses of infants' responses in various object-manipulation tasks led him to conclude that, during the first year of life, infants possess very little physical knowledge. For the next several decades, this conclusion was generally accepted (for reviews of this early research, see Bremner, 1985; Gratch, 1976; Harris, 1987; and Schubert, 1983). This state of affairs began to change in the 1980s, however, when evidence obtained with novel, more sensitive tasks revealed that even young infants hold at least limited expectations about physical events (e.g., Baillargeon, 1986; Baillargeon, Spelke, & Wasserman, 1985; Baillargeon & Graber, 1987; Diamond, 1985; Hood & Willatts, 1986; Leslie, 1982, 1984; Pieraut-Le Bonniec, 1985; Spelke & Kestenbaum, 1986).

In subsequent years, researchers began to explore many new facets of infants' physical knowledge, bringing to light new competences and developments (e.g., Arterberry, 1993; Clifton, Rochat, Litovsky, & Perris, 1991; Diamond, 1991; Goubet & Clifton, 1998; Kotovsky & Baillargeon, 1994; Lécuyer, 1993; Needham & Baillargeon, 1993; Oakes & Cohen, 1990; Spelke, Breinlinger, Macomber, & Jacobson, 1992). Today, investigators generally agree (with a few notable exceptions: e.g., Bogartz, Shinsky, & Speaker, 1997;

The preparation of this chapter was supported by a grant from the National Institute of Child Health and Human Development (HD-21104). I would like to thank Jerry DeJong, Cindy Fisher, Yuyan Luo, Kris Onishi, and Su-Hua Wang for many helpful discussions, comments, and suggestions. Correspondence concerning this chapter should be sent to: Renée Baillargeon, Department of Psychology, University of Illinois, 603 East Daniel, Champaign, IL 61820, USA (email: <rbail-lar@s.psych.uiuc.edu>).

Haith & Benson, 1998; Rivera, Wakeley, & Langer, 1999) that young infants' physical world is far more sophisticated than Piaget (1954) – with the limited methodological tools at his disposal – would ever have thought possible.

In keeping with these advances, for the past 10 years my collaborators and I have been investigating the development of infants' physical knowledge. Our research has focused on two main questions: first, what knowledge do infants possess, at each age, about different physical events (e.g., occlusion, support, collision, and containment events);¹ and second, how do infants attain this knowledge? In this chapter, I summarize some of the main findings we have obtained to date. For ease of communication, I have organized this summary in eight "lessons." These lessons are of course still preliminary. Nevertheless, they are useful in providing a framework for what has been learned so far, and in making clear what needs to be studied next.

To give a brief overview, the first three lessons are concerned with the nature of the expectations infants acquire about physical events. The next three lessons deal with some of the factors and processes involved in the acquisition of these physical expectations. Finally, the last two lessons address the possible contributions of innate concepts to infants' physical reasoning.

Lesson 1: Infants Acquire Rules about Physical Events

The first lesson suggested by our research and that of other investigators is that infants acquire expectations or rules about physical events; these rules specify for them what are the likely outcomes of events (e.g., Aguiar & Baillargeon, 1999; Baillargeon, Graber, DeVos, & Black, 1990; Hespos & Baillargeon, 2001b; Kotovsky & Baillargeon, 2000; Lécuyer & Durand, 1996; Needham, 1998; Newcombe, Huttenlocher, & Learmonth, 1999; Wilcox, Nadel, & Rosser, 1996). When faced with events inconsistent with their rules, as in violation-of-expectation experiments (e.g., Baillargeon, 1995, 1998, 2000b), infants typically are surprised or puzzled, as evidenced by increased attention: under most circumstances, infants look reliably longer at events that violate, as opposed to confirm, their physical expectations.

Not surprisingly, in the initial stages of learning, infants' rules about physical events tend to be rather primitive or incomplete, so that they often err in determining what are the likely outcomes of events. Two types of errors have been documented to date. First, infants sometimes fail to view as unexpected events that adults perceive to be physically impossible (I will refer to such events as violation events). Second, infants sometimes view as unexpected events that adults perceive to be physically possible and indeed commonplace (non-violation events). Their lack of physical knowledge thus leads infants both (1) to respond to violation events as though they were expected and (2) to respond to non-violation events as though they were unexpected. To illustrate these two types of errors, I briefly describe new findings on 2.5-month-old infants' knowledge of occlusion events.

Occlusion events

Recent evidence suggests that, although 2.5-month-old infants recognize that an object continues to exist *after* it becomes occluded, they are rather poor at predicting *when* it

should be occluded (e.g., Aguiar & Baillargeon, 1999; Hespos & Baillargeon, 2001b; Luo, 2000; Luo & Baillargeon, 2001a; Spelke et al., 1992; Wilcox et al., 1996). At 2.5 months of age, infants appear to follow a simple “behind/not-behind” rule when predicting the outcomes of occlusion events: they expect an object to be hidden when behind an occluder, and to be visible otherwise. At this stage, infants do not take into account information about the relative sizes of the object and occluder, or about the presence of openings in the occluder; *any* object is expected to be hidden when behind *any* occluder.

Because their knowledge of the conditions under which objects should and should not be occluded is very limited, 2.5-month-old infants often err in distinguishing between violation and non-violation occlusion events. A recent experiment by Yuyan Luo and myself clearly illustrates this point (Luo & Baillargeon, 2001a). This experiment built on prior findings by Aguiar and Baillargeon (1999) and examined infants’ ability to determine whether an object should remain continuously hidden or become temporarily visible when passing behind a screen with a large opening in its midsection.

The infants were assigned to a cylinder-appears or a cylinder-does-not-appear condition (see figure 3.1). In both conditions, the infants first saw a familiarization event in which an upright cylinder moved back and forth along a track whose center was hidden by a screen; the cylinder disappeared at one end of the screen and reappeared, after an appropriate interval, at the other end. Next, the infants saw two test events. In one (separate-screens event), the entire midsection of the screen was removed to create two separate screens. In the other event (connected-screens event), the two screens remained connected at the top by a short strip. In both events, the cylinder moved back and forth along the track, as in the familiarization event. For the infants in the cylinder-appears condition, the cylinder appeared in the gap between the screens in each test event. For the infants in the cylinder-does-not-appear condition, the cylinder disappeared behind one screen and reappeared from behind the other screen *without* appearing in the gap between them.


As adults, we would expect the cylinder to appear both between the separate and the connected screens. What of the 2.5-month-old infants in the experiment, with their simple behind/not-behind rule? We predicted that, like adults, the infants should expect the cylinder to appear between the separate screens, because at that point the cylinder did not lie behind any occluder. Unlike adults, however, the infants should expect the cylinder *not* to appear between the connected screens: the infants should view these as a single occluder, and they should expect the cylinder to remain hidden when passing behind it.

The results supported our predictions: the infants in the cylinder-appears condition looked reliably longer at the connected- than at the separate-screens test event, whereas those in the cylinder-does-not-appear condition showed the reverse looking pattern. Together, these results suggested that the infants expected the cylinder to appear between the separate but not the connected screens, and were surprised when each of these expectations was violated.


To return to our first lesson: the infants’ limited knowledge about occlusion events led them to err, in both conditions, in their response to the connected-screens event. The infants were not surprised when the cylinder failed to appear between the screens (a violation event), and they were surprised when it did appear (a non-violation event).

Cylinder-appears condition

Connected-screens event




Separate-screens event




Cylinder-does-not-appear condition

Connected-screens event



Separate-screens event

**Figure 3.1** Schematic drawing of the test events in Luo and Baillargeon (2001a)*Additional remarks*

It is not very surprising that infants detect far fewer physical violations than do adults. One would expect that, as infants' physical knowledge grows, the set of violation events they recognize gradually approximates that of adults. What is more intriguing is the fact that infants sometimes perceive physically possible, non-violation events as violation events (e.g., Baillargeon, DeJong, & Sheehan, 2001; Luo & Baillargeon, 2001a; Wang & Baillargeon, 2001a). Such findings provide strong evidence that infants are acquiring rules about physical events, rules that are initially limited and incomplete and as such can lead to false predictions. One is reminded here of the young child who produces such words as "goed" and "eated" in the course of acquiring the past-tense rule of English (e.g., Marcus, Pinker, Ullman, Hollander, Rosen, & Xu, 1992).

Some investigators have recently questioned the notion that infants acquire rules about physical events (e.g., Bogartz et al., 1997; Haith & Benson, 1998; Rivera et al., 1999; Thelen & Smith, 1994). For example, Bogartz et al. (1997) argued that the process of knowledge acquisition is essentially one of data collection. Infants collect and store “videotapes” of physical events. When faced with an event, infants search through their “library” of videotapes, retrieve the most relevant, and compare the current and stored events; mismatches engage infants’ attention and cause them to update the existing videotape or create a new one.

This approach can address why infants respond to violation events with increased attention – no videotape in their library would be likely to match these physically impossible events. However, what of the findings that infants also respond to physically possible, non-violation events with increased attention? Why would infants detect mismatches where there are none?

A more parsimonious explanation for the evidence currently available is that, in their attempts to make sense of physical events, infants formulate rules about how the events might operate. Initially, these rules tend to be rather primitive, with the result that infants often err in determining what are violation and non-violation events. Of course, these errors themselves may play a powerful motivating role in the development of infants’ physical knowledge. The 2.5-month-old infant who notices that, contrary to his expectations, objects do become temporarily visible when passing behind occluders with central openings, is taking the first step toward improving his knowledge of the conditions under which objects should and should not be occluded. We return to this issue in Lesson 4.

Lesson 2: Infants’ Rules Become More Sophisticated Over Time

The second lesson suggested by the research from our and other laboratories is that infants’ rules become more sophisticated over time (e.g., Aguiar & Baillargeon, in press; Baillargeon, 1991; Dan, Omori, & Tomiyasu, 2000; Hespos & Baillargeon, 2001a; Kotovsky & Baillargeon, 1998; Needham, 1999; Sitskoorn & Smitsman, 1995; Wilcox, 1999). In fact, infants’ rules about different physical events all seem to develop according to the same general pattern. Specifically, when learning about events such as occlusion, support, collision, and other events, infants typically first form an initial concept centered on a primitive, all-or-none distinction. With further experience, infants identify a sequence of variables – some discrete and others continuous – that revise and elaborate this initial concept, resulting in increasingly accurate predictions and interpretations over time. To illustrate this developmental pattern, I summarize the results of experiments from our laboratory on the development of infants’ knowledge about occlusion and support events.

Occlusion events

In our experiments on the development of infants’ expectations about occlusion events (e.g., Aguiar & Baillargeon, 1999, in press; Baillargeon & DeVos, 1991; Luo, 2000; Luo

Infants' knowledge about occlusion events
violation detected
at each stage

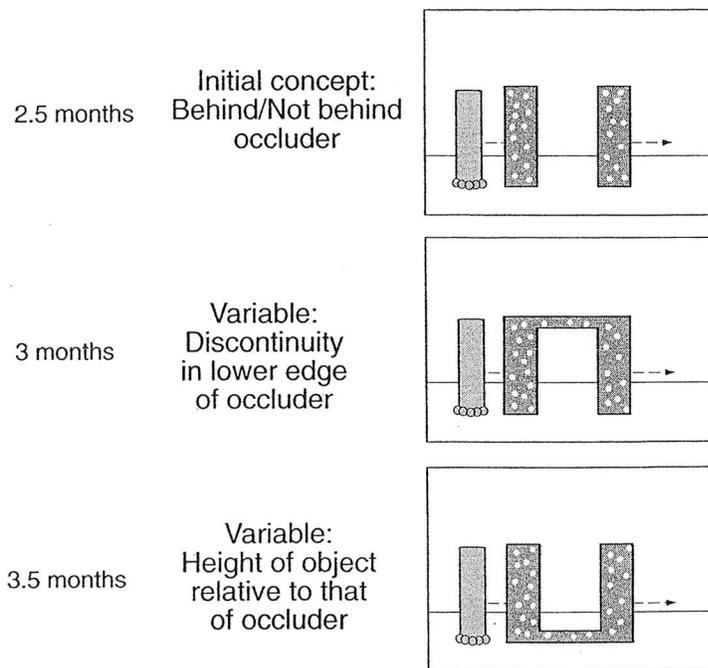


Figure 3.2 Schematic description of the development of infants' knowledge about occlusion events between 2.5 and 3.5 months of age

& Baillargeon, 2001a; for reviews, see Baillargeon, 1998, 1999), infants aged 2.5 to 3.5 months watched an object (e.g., a cylinder, a toy mouse, or a toy carrot) move back and forth behind a screen; next, a portion of the screen was removed, and the infants judged whether the object should remain hidden or become (at least partly) visible when passing behind the screen. The results of these experiments are illustrated in figure 3.2. By 2.5 months of age, as was discussed earlier, infants have formed an initial concept of occlusion centered on a simple *behind/not-behind* distinction. When the entire mid-section of the screen is removed to form two separate screens, infants expect the object to become visible in the gap between them. However, if the screens remain connected either at the top or at the bottom by a short strip, infants no longer expect the object to become visible: they view the connected screens as a single screen and expect the object to be hidden when behind it. Over the course of the next month, infants rapidly progress beyond their initial concept. At about 3 months of age, infants begin to consider the presence of a *discontinuity in the lower edge* of the screen. Although infants still expect the object to remain hidden when passing behind two screens that are connected at the bottom by a short strip, they now expect the object to become visible when passing behind two screens that are connected at the top by a short strip. Finally, at about 3.5 months of age, infants begin to consider the relative *heights* of the object and screen. When the object passes behind two screens that are connected at the bottom by a strip, infants expect the object to become partly visible if it is taller but not shorter than the strip.

Infants' knowledge about support events
violation detected
at each stage




Figure 3.3 Schematic description of the development of infants' knowledge about support events between 3 and 12.5 months of age

Support events

In our experiments on the development of infants' knowledge about support events (e.g., Baillargeon, Needham, & DeVos, 1992; Needham & Baillargeon, 1993; for reviews, see Baillargeon, 1995, 1998, and Baillargeon, Kotovsky, & Needham, 1995), infants aged 3 to 12.5 months were presented with support problems involving a box and a platform; the box was held in one of several positions relative to the platform, and the infants judged whether the box should remain stable when released. The results of these experiments are summarized in figure 3.3. By 3 months of age, infants have formed an initial concept of support centered on a simple *contact/no-contact* distinction: they expect the box to remain stable if released in contact with the platform, and to fall otherwise. At this stage, any contact with the platform is deemed sufficient to ensure the box's stability. In the months

that follow, infants identify a sequence of variables that progressively revise and elaborate their initial concept. At about 4.5 to 5.5 months of age (females precede males by a few weeks in this acquisition), infants begin to take into account the *type of contact* between the box and the platform. Infants now expect the box to remain stable when released on but not against the platform. At about 6.5 months of age, infants begin to consider the *amount of contact* between the box and the platform. Infants now expect the box to remain stable if a large but not a small portion of its bottom surface rests on the platform. Finally, at about 12.5 months of age, infants begin to attend to the *proportional distribution* of the box; they realize that an asymmetrical box can be stable only if the proportion of the box that rests on the platform is greater than that off the platform.

Additional remarks

The results summarized above all focus on infants' responses to *violation* events – what new violation infants become able to detect at each age, and how the total set of violations they can detect increases steadily with age. However, based on these same results, one can also predict how infants should respond to the converse *non-violation* events. In the case of occlusion events, it should be the case that 2.5-month-old infants, for example, are surprised when an object appears between two screens that are connected either at the top or at the bottom by a short strip. Similarly, in the case of support events, it should be the case that 3-month-old infants, for example, are surprised to see an object fall when deposited against a wall or on the edge of a table. Until now, most of the research on infants' physical reasoning has tended to focus on infants' responses to violation as opposed to non-violation events. However, as we saw in Lesson 1, evidence that infants perceive non-violation events as surprising is extremely useful in that it helps make clear the nature of the rules that underlie their interpretations of events.

The focus on violation as opposed to non-violation events has also tended to obscure certain developments in infants' physical knowledge: those in which the addition of a new variable does not lead to the detection of a new violation (as in figures 3.2 and 3.3). Recent research by Dan et al. (2000), Huettel and Needham (2000), and Wang and Baillargeon (2001a) suggests that, at about 8.5 months of age, infants add a new variable to their understanding of support events. Specifically, infants begin to distinguish between situations in which the *side or middle portion* of a box's bottom surface rests on a platform; they recognize that, in the latter case, when the box is balanced on a narrower platform, stability is possible even if less than half of the box's bottom surface is in contact with the platform. This development is interesting in that it does not allow infants to detect a novel violation (and hence does not easily fit into figure 3.3). Rather, it leads infants to regard events previously viewed as violation events (e.g., a box that remains stable when balanced on a narrow platform, with, say, only the middle 33 percent of its bottom surface supported) as non-violation events.

The findings summarized to this point do not, of course, represent all that infants learn about occlusion and support events. First, many more variables must be identified. For example, in the case of occlusion events, infants must also learn that wide objects can be fully hidden behind wide but not narrow occluders (e.g., Baillargeon & Brueckner,

2000; Wilcox, 1999; Wilcox & Baillargeon, 1998a, 1998b), that faster objects reappear sooner than slower objects from behind occluders (e.g., Spelke, Kestenbaum, Simons, & Wein, 1995; Wilcox & Schweinlee, 2001), and that opaque occluders function differently than do transparent occluders (e.g., Luo, 2001; Luo & Baillargeon, 2001b, 2001d).

Second, infants' reasoning about several of the variables they identify must undergo considerable refinement. For example, in the case of support events, preliminary data collected by Su-Hua Wang and myself suggest that when 6.5-month-old infants first identify amount of contact as a variable, they assume that an object will be stable if 66 percent but not 50 percent of its bottom surface is supported; over time, infants come to realize that 50 percent is typically sufficient to ensure stability. In the same vein, Dan et al. (2000) reported that after infants come to realize that an object can be stable when balanced on a narrower platform, they must still learn how wide the platform needs to be for the object to be stable.

Finally, other developments involve the transition from qualitative to quantitative reasoning. The distinction between quantitative and qualitative reasoning strategies is derived from computational models of everyday physical reasoning (e.g., Forbus, 1984). A strategy is said to be quantitative if it requires one to encode and use information about absolute quantities (e.g., object A is "this" tall, where "this" stands for some absolute measure of A's height). In contrast, a strategy is said to be qualitative if it requires one to encode and use information about only relative quantities (e.g., object A is taller than object B). There is now considerable evidence (for reviews, see Baillargeon, 1994, 1995) that, when infants first identify a continuous variable, they can reason about the variable qualitatively but not quantitatively: they are not able at first to encode and remember absolute information about the variable.

A recent experiment by Yuyan Luo and myself clearly illustrates this point (Luo & Baillargeon, 2001c). This experiment examined whether 5-month-old infants realize that the height of an object relative to that of an occluder determines not only (1) *whether* the object should appear above the occluder, but also (2) *how much* of the object should appear above the occluder. The infants were assigned to a qualitative or a quantitative condition (see figure 3.4). The infants in the qualitative condition first saw familiarization events in which a tall or a short cylinder moved back and forth along a track whose center was hidden by a screen; the tall cylinder was the same height as the screen (30 cm), and the short cylinder was 8 cm shorter (22 cm). The tall and short cylinders were shown on alternate trials. Next, a window was cut into the top midsection of the screen; the bottom of the window was located 14 cm above the screen's lower edge. The infants saw two test events in which the tall and short cylinders again moved back and forth along the track. In both events, the cylinder appeared in the window when passing behind the screen; and in both events, the cylinder extended to the top of the window. This last outcome was expected for the tall cylinder, which was as tall as the screen, but not for the short cylinder, which was 8 cm shorter and should have reached only the midpoint of the window. The infants in the quantitative condition saw similar familiarization and test events except that the screen was 8 cm taller. In this condition, both the tall and the short cylinders were shorter than the screen, and both reached a point 8 cm below the top of the window when passing behind the screen (this outcome was expected for

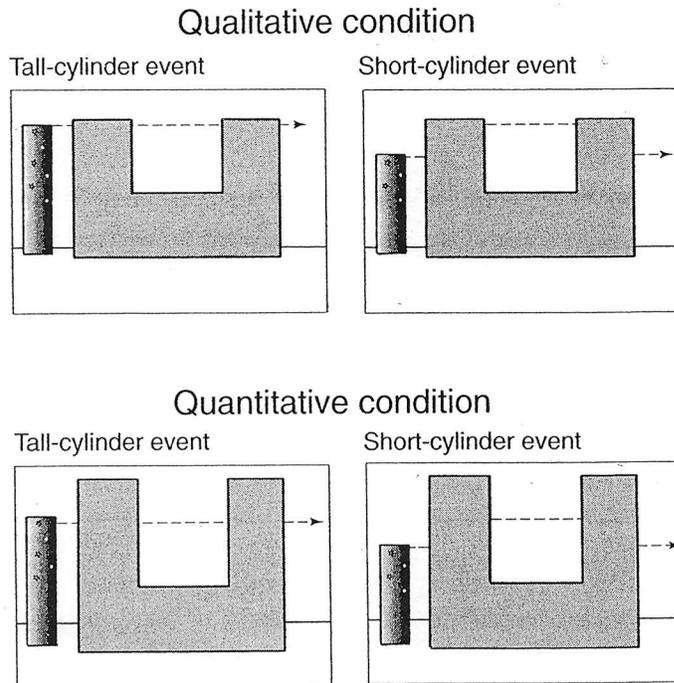


Figure 3.4 Schematic drawing of the test events in Luo and Baillargeon (2001c)

the tall cylinder, which was 8 cm shorter than the screen, but not for the short cylinder, which was 16 cm shorter).

The infants in the qualitative condition could succeed at detecting the violation in the short-cylinder test event by reasoning qualitatively about the event. As the short cylinder approached the screen, the infants could visually compare its height to that of the bottom and top of the window. Based on these comparisons, they could conclude that (1) since the short cylinder was taller than the bottom of the window, it would appear in the window; and (2) since the short cylinder was shorter than the top of the window, it would not reach the top of the window. This last expectation was violated when the short cylinder extended all the way to the top of the window.

In contrast, the infants in the quantitative condition could detect the violation in the short-cylinder test event only by engaging in quantitative reasoning. A qualitative comparison of the heights of the short cylinder and window could establish only that (1) the cylinder would appear in the window and (2) the cylinder would not reach the top of the window. To detect the violation in the short-cylinder event – that is, to detect that the short cylinder extended higher than it should have in the window – the infants had to engage in quantitative reasoning: they needed to encode the absolute height of the short cylinder as it approached the screen, and to compare this (represented) height to that of the cylinder in the window.

The infants in the qualitative condition looked reliably longer at the short- than at the tall-cylinder test event, whereas those in the quantitative condition tended to look equally at the two events. Together, these results suggested that 5-month-old infants can reason qualitatively but not quantitatively about height information in occlusion events. Further

research is needed to determine at what age infants become able to engage in quantitative reasoning about this variable.

Lesson 3: Infants' Rules are Narrow in Scope

So far we have learned that infants form rules about physical events, and that these rules become richer and more complex with the identification of additional variables. In Lesson 3, we consider the following question: How general or specific are infants' rules about physical events? Do infants acquire *general* rules that are applied broadly to all relevant physical events, or more *specific* rules that remain tied to the events where they are first acquired?

Recent research suggests that the second possibility is more likely (e.g., Hespos, 1998, 2000; Hespos & Baillargeon, 2001a; Luo, 2001; Luo & Baillargeon, 2001b; Wang, Baillargeon, & Paterson, 2001; Wang & Paterson, 2000; Wilcox & Baillargeon, 1998a). Specifically, it appears that infants "sort" physical events into event categories, and learn separately how each category operates. A variable acquired in the context of one event category is not generalized to other relevant categories; it is kept tied to the specific category where it is first identified. As a result, infants must sometimes "relearn" in one event category a variable they have already acquired in another category. When weeks or months separate these two acquisitions, striking lags (or, to borrow a Piagetian term, *décalages*; e.g., Flavell, 1963) can be observed in infants' responses to events from the two categories. To illustrate such lags, I briefly describe the results of recent experiments on infants' reasoning about height and transparency information in occlusion, containment, and other events.

Height information

In a recent experiment, Sue Hespos and I compared 4.5-month-old infants' ability to reason about height information in containment and in occlusion events (Hespos, 1998; Hespos & Baillargeon, 2001a). Specifically, we asked whether infants this age realize that a tall object cannot be fully hidden when placed inside a short container or behind a short occluder.

The infants were assigned to a container or an occluder condition and saw two test events (see figure 3.5). At the start of each event shown in the container condition, an experimenter's gloved hand grasped a knob at the top of a tall cylindrical object; next to the object was a cylindrical container. The hand lifted the object and lowered it inside the container until only the knob protruded above the rim. The container used in the tall-container event was as tall as the cylindrical portion of the object; the container used in the short-container event was only half as tall, so that it should have been impossible for the cylindrical portion of the object to become fully hidden inside the container. Prior to the test trials, the infants received familiarization trials in which the containers were rotated forward so that the infants could inspect them. The infants in the occluder

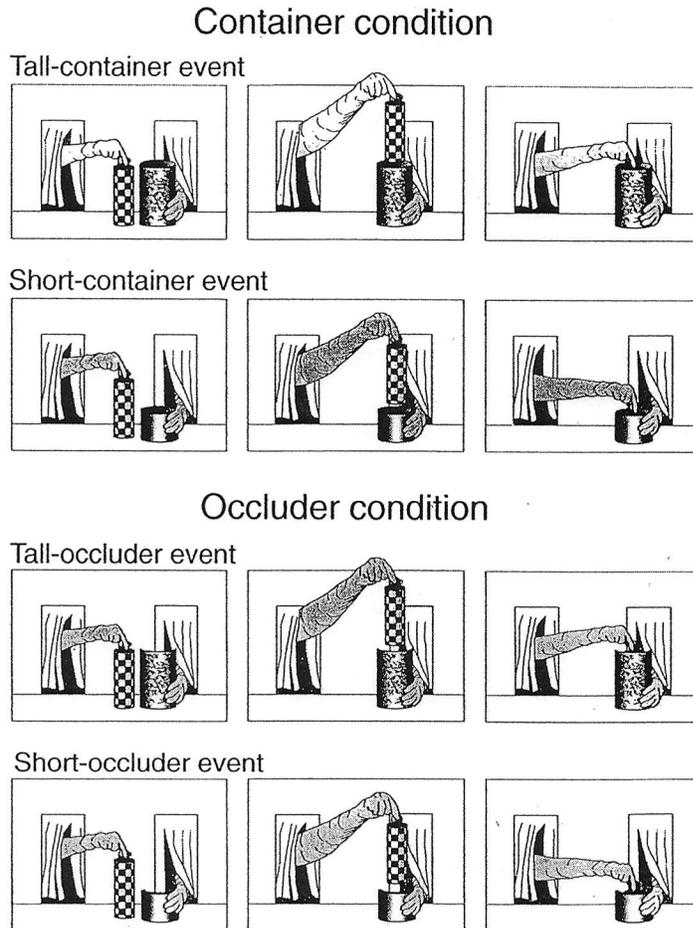


Figure 3.5 Schematic drawing of the test events in Hespos and Baillargeon (2001a)

condition saw similar familiarization and test events with one exception: the bottom and back half of each container were removed to create a rounded occluder.

The infants in the occluder condition looked reliably longer at the short- than at the tall-occluder test event, whereas those in the container condition looked about equally at the two events. Our interpretation of these results was that 4.5-month-old infants view occlusion and containment as two distinct event categories and do not generalize rules or variables acquired about occlusion to containment. Infants realize that the height of an object relative to that of an occluder determines whether the object can be fully or only partly hidden when behind the occluder, but they do *not* yet appreciate that the height of an object relative to that of a container determines whether the object can be fully or only partly hidden when inside the container.

This interpretation led to a striking prediction: infants shown the same test events as in the container condition but with the object lowered *behind* rather than *inside* each container should be able to detect the violation in the short-container test event. With the containers serving as mere occluders, infants' performance should mirror that of the infants in the occluder condition. This prediction was confirmed: when the object was lowered behind rather than inside each container, infants looked reliably longer at the

short- than at the tall-container event. These and additional control results (Hespos & Baillargeon, 2001a) provided converging evidence that 4.5-month-old infants view occlusion and containment as distinct event categories and learn separately about each category. Infants consider the height of an object relative to that of a container when the object is lowered *behind* but not *inside* the container.

At what age do infants become able to reason about height in containment events? In an additional experiment, 5.5-, 6.5-, and 7.5-month-old infants were shown the container condition test events (Hespos & Baillargeon, 2001a). Only the 7.5-month-old infants looked reliably longer at the short- than at the tall-container event, suggesting that it is not until this age that infants begin to consider height information when reasoning about containment events.

It might be suggested that the 7.5-month-old infants were perhaps successful, not because they had – at long last – identified height as a containment variable, but because they had achieved a general ability to reason about height in different physical events. However, results obtained with other events cast doubts on such a possibility (e.g., Baillargeon, 1993; Wang & Paterson, 2000; Wang et al., 2001). For example, Wang and Paterson (2000), building upon the present results, compared 9-month-old infants' ability to reason about height information in containment and in covering events (in these events, the containers were turned upside down and lowered over the object). The infants succeeded in the container but not the cover condition: they were surprised when a tall object was fully lowered inside a short container, but not when a short cover was fully lowered over the same tall object. Recent results (Wang & Baillargeon, 2001b) indicate that it is not until infants are about 12 months of age that they begin to reason about height information in covering events. These findings make clear that infants do not acquire at 7.5 months of age a generalized ability to reason about height in physical events.

The results reported in this section (and converging results obtained with object-retrieval paradigms: Hespos, 1998; McCall, 2001) suggest that infants view events involving occluders, containers, and covers as belonging to separate categories, and do not generalize information acquired about one category to the others. Infants begin to consider height information in occlusion events at about 3.5 months of age (e.g., Baillargeon & DeVos, 1991), in containment events at about 7.5 months of age (e.g., Hespos, 1998; Hespos & Baillargeon, 2001a), and in events involving covers at about 12 months of age (e.g., McCall, 2001; Wang & Baillargeon, 2001b; Wang & Paterson, 2000; Wang et al., 2001).

Transparency information

In a recent experiment, Yuyan Luo and I compared 8.5-month-old infants' ability to reason about transparency information in containment and in occlusion events (Luo, 2001; Luo & Baillargeon, 2001b). More specifically, the experiment asked whether infants expect an object to be visible when lowered inside a transparent container or behind a transparent occluder.

The infants were assigned to a container or an occluder condition and saw two test events (see figure 3.6). At the start of each event shown in the container condition, a

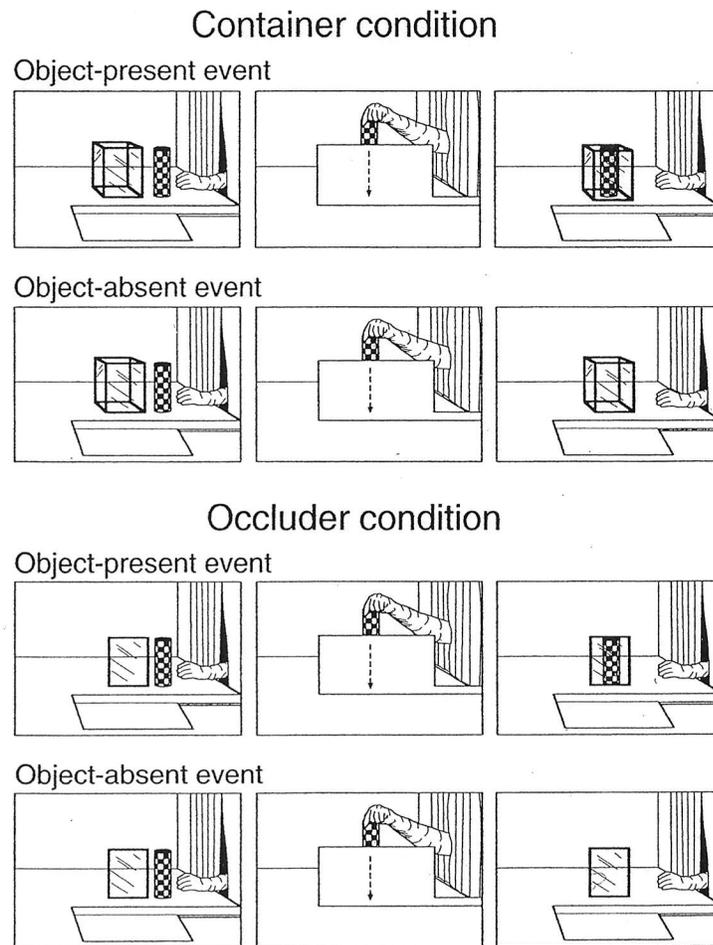


Figure 3.6 Schematic drawing of the test events in Luo and Baillargeon (2001b)

cylindrical object stood next to a rectangular container of the same height; the container was made of transparent Plexiglass and its edges were outlined with red tape (prior to the test session, an experimenter showed the container to the infants for a few seconds, so that they had the opportunity to inspect its surfaces). To start, a screen was raised to hide the container, and an experimenter's gloved hand grasped the object and lowered it inside the hidden container. Next, the screen was lowered to reveal the container with either the object standing inside it (object-present event), or no object inside it (object-absent event). The infants in the occluder condition saw similar test events except that the transparent container was replaced with a transparent occluder. This occluder was identical to the front of the container; a small Plexiglass base behind the occluder allowed it to stand upright.

The infants in the occluder condition looked reliably longer at the object-absent than at the object-present test event, whereas those in the container condition tended to look equally at the two events. These results suggested that 8.5-month-old infants expect an object to be visible when lowered behind a transparent occluder, but not when lowered inside a transparent container. In subsequent experiments (Luo & Baillargeon, 2001b), we found that it is not until infants are about 10 months of age that they begin to attend to transparency information in containment events.

These results (and additional results, some of which are discussed later; Luo & Baillargeon, 2001d) provide converging evidence that infants view occlusion and containment as distinct event categories, and learn separately how each category operates. As was the case with height (e.g., Hespos & Baillargeon, 2001a), infants identify transparency as a variable first in occlusion events, and only after some time in containment events.

Additional remarks

The findings reviewed in this section suggest that infants sort physical events into distinct categories, and learn separately about each category. Such a learning strategy must, overall, greatly facilitate infants' acquisition of physical knowledge; after all, breaking down the task of learning into smaller, more manageable components is a time-honored solution to the difficulties of knowledge acquisition.

The notion that infants acquire event-specific rather than event-general rules or expectations is consistent with an emerging theme in the developmental literature that infants' knowledge tends to be highly context-specific. For example, Adolph (1997) reported that infants learn to navigate steep slopes with caution in the first weeks of crawling – and again in the first weeks of walking; the knowledge that steep slopes can lead to falling is not generalized from crawling to walking but must be learned all over again. Similarly, Needham (in press) found that infants use featural information to segregate objects placed side by side several months before they succeed in doing the same with objects (even the same objects) placed one on top of the other. Finally, Onishi (2000) reported that 13-month-old infants correctly judge whether a stack of boxes should be stable when they are led to believe that the stack is composed of two but not three boxes; apparently, infants do not generalize the support rules they use when reasoning about stacks of two boxes to stacks of three boxes.

The evidence that infants form event categories raises many questions for future research. In particular, on what basis are categories generated? Why are occlusion and containment, for example, viewed as distinct categories? In many cases (and contrary to those studied here), occlusion and containment outcomes are different: for example, a wide object can be lowered behind a narrow occluder, but not inside a narrow container; and an object that has been lowered inside a container typically moves with it when displaced, but an object that has been lowered behind an occluder does not. It seems plausible that these clusters of interrelated causal relationships underlie infants' distinct event categories (e.g., Keil, 1991, 1995; Leslie, 1994, 1995; Pauen, 1999). We return to this issue in Lesson 6.

Lesson 4: The Acquisition of Rules is Triggered by Exposure to Unpredicted Outcomes

The evidence reviewed in the previous lessons suggests that infants sort physical events into distinct categories and, for each category, identify a sequence of variables that specify

(rightly or wrongly) expected outcomes. One important issue that has not been discussed so far is infants' identification of variables in each event category. How does this process occur?

Before addressing this question, we need to consider more closely what variables are. My collaborators and I (e.g., Aguiar & Baillargeon, 1999; Baillargeon, in press; Hespos & Baillargeon, 2001b) have suggested that variables are akin to condition–outcome rules. A variable specifies, for a set of contrastive outcomes, what condition produces each outcome. For example, the containment variable *width* specifies that an object *can* be inserted into a container if it is *narrower* than the opening of the container, but *cannot* be inserted if it is *wider* than the opening. For each of the two contrastive outcomes (can or cannot be inserted into the container), the variable identifies the condition responsible for the outcome (narrower or wider than the opening of the container).²

How, then, do infants go about identifying variables? We have proposed that what typically triggers the identification of a variable in an event category is exposure to contrastive outcomes that are not predicted by infants' current knowledge of the category (e.g., Aguiar & Baillargeon, 1999; Baillargeon, in press; Hespos & Baillargeon, 2001b). When infants register these contrastive outcomes, they begin to seek out the conditions that are responsible for them. The identification of these condition–outcome relations signals the identification of a new variable.

To illustrate, consider once again the containment variable *width*. When reasoning about containment events, infants initially assume that any object can be inserted into any container with an open – as opposed to a closed – top (e.g., Hespos & Baillargeon, 2001b). After some time, however, infants begin to notice that objects sometimes can and sometimes cannot be inserted into containers with open tops. At that point, infants begin to search for the conditions that map onto these contrastive outcomes. Eventually, infants come to realize that objects can be inserted into containers whose openings are wider but not narrower than the objects (e.g., Aguiar & Baillargeon, 1998, 2000b, 2001; Baillargeon & Brueckner, 2000; Sitskoorn & Smitsman, 1995).

In many instances, as in the preceding example, contrastive outcomes will involve contradictions of infants' current knowledge. Infants will realize that, contrary to their expectations, objects do not always remain hidden when passing behind occluders, do not always move when hit, do not always remain stable when placed on supports, and so on. In each case, noticing these contradictions will lead to a revision of infants' rules – until the noticing of further contradictions brings forth further revisions. In other instances, however, contrastive outcomes may simply involve facets of events that had hitherto gone unnoticed. For example, infants seem to realize at some point in their development that objects not only move when hit, but move different distances depending on the sizes (or masses) of the objects hitting them (e.g., Kotovsky & Baillargeon, 1994, 1998). Such a realization does not contradict infants' prior knowledge about collision events, but still adds a new variable to this knowledge.

From the present perspective, what triggers the identification of new rules is thus exposure to some unpredicted variation in outcome, whether or not this variation conflicts with pre-existing rules. Is there any evidence for such a notion? Recent findings by Yuyan Luo and myself could be taken to support it (Luo & Baillargeon, 2001d). This experiment built on the results of the transparency experiment reported in Lesson 3 (Luo, 2001;

Luo & Baillargeon, 2001b), and asked at what age infants identify transparency as an occlusion variable. As we saw earlier, 8.5-month-old infants expect an object that is placed behind a transparent occluder to be visible through the occluder. We next tested 7.5-, 7-, and 6.5-month-old infants, using the same object-present and object-absent test events as before (see figure 3.6).

The results were unexpected. Like the 8.5-month-old infants in our previous experiment (Luo, 2001; Luo & Baillargeon, 2001b), both the 7.5- and the 6.5-month-old infants looked reliably longer at the object-absent than at the object-present test event. In contrast, the 7-month-old infants showed the *reverse* pattern: they looked reliably longer at the object-present than at the object-absent event.

Our interpretation for these results (and control results; Luo, 2001; Luo & Baillargeon, 2001b) is as follows. Prior to about 7 months of age, infants do not perceive the clear surface of the occluder: they see only an empty frame (remember that the edges of the occluder are outlined with red tape). Based on their knowledge of occlusion events, infants expect the object to be visible in this frame, and they are surprised in the object-absent test event when it is not. At about 7 months of age, infants' vision is sufficiently improved that they can now perceive the clear surface of the occluder. At this stage, infants are puzzled by the object-present event. Their knowledge of occlusion specifies that an object should be hidden, not visible, when placed behind a larger occluder (e.g., as when a small cup is placed behind a large cereal box); and yet infants can clearly see the object behind the occluder. How is this possible? In their daily lives, 7-month-old infants no doubt experience a range of similar experiences (e.g., when they notice people and objects through car or house windows). Infants then begin to search for the conditions that explain why objects behind larger occluders are sometimes visible and sometimes hidden. In essence, infants realize that an additional variable must be taken into account to explain these observations, and they set about discovering what it is. By 7.5 months of age (and thus in a very short time), infants have formed a new condition–outcome rule: they now expect objects to be hidden when behind larger occluders that are opaque but not transparent. Armed with this new knowledge, infants can now detect violations such as that in the object-absent event.

Support for the notion that infants' perception of transparent surfaces improves at about 7 months of age comes from an experiment by Johnson and Aslin (2000). They reported improvements at this age in infants' responses to a computer-animated display consisting of a transparent box suspended in front of an opaque rod. Johnson and Aslin suggested that this change might reflect gains in contrast sensitivity, which might in turn be tied to the maturation of the magnocellular system.

Additional remarks

The main claim made in Lesson 4 is that the identification of new variables or rules is typically triggered by exposure to unpredicted outcome variation. This claim makes strong predictions about the conditions under which learning about event categories should and should not occur. In particular, it predicts that where no variation is experienced, no learning should occur.

To illustrate, consider once again the variable width in containment events. What is being proposed is that in order to identify this variable, infants must see *both* objects being inserted into wider containers, and objects failing to be inserted into narrower containers. On the present view, seeing only objects being inserted into wider containers would not be sufficient for infants to abstract the variable width. Infants would not be able to reflect on these observations and detect the underlying regularity that the object inserted into the container is always narrower than its opening. For width to be identified as a variable, infants must experience variation in outcome – they must notice that objects sometimes can and sometimes cannot be inserted into containers – and thus be induced to search for the conditions responsible for these different outcomes.

One important caveat needs to be introduced here. The claim being made is *not* that infants can never learn facts about objects in the absence of outcome variation. Such a claim is blatantly false. As an illustration, consider the results of an experiment conducted by Laura Kotovsky and myself on 5.5-month-old infants' responses to collision events (Kotovsky & Baillargeon, 1998). The infants were first habituated to the following event: a cylinder rolled down a ramp and hit a wheeled toy bug, causing it to roll to the middle of a track. Following habituation, the infants saw two test events in which the same cylinder hit the bug, causing it to roll to either the middle (same event) or the end (different event) of the track. The infants looked reliably longer at the different than at the same event, suggesting that they had learned how far the bug rolled when hit by the cylinder and noticed when a change was introduced. Because the bug always traveled the same distance during the habituation trials, learning obviously occurred in the absence of any outcome variation, through some associative process.

What is being suggested, then, is that infants acquire their knowledge about objects through several learning mechanisms, each with its own requirements for learning. Facts about individual objects and events can be learned through repetition; but general and abstract facts that apply to entire event categories cannot. To return to our collision example above (Kotovsky & Baillargeon, 1998), infants can readily learn, through simple repetition of the cylinder-bug event, that the bug rolls to the middle of the track when hit by the cylinder. But infants could not learn, even with a million repetitions of the cylinder-bug event, the general rule that objects move farther when hit by larger (or heavier) than by smaller (or lighter) objects. Indeed, infants could not learn this rule even if they saw different objects hit different wheeled toys, all of which then rolled to the middle of the track. In order to learn such a rule, we believe, infants would need to see objects roll *different distances* when hit by different objects. Exposure to these contrastive outcomes would induce infants to search for the conditions responsible for them, resulting in the acquisition of a new condition–outcome rule.

To put the preceding arguments another way: infants are not designed to reflect on a collection of similar observations (e.g., a spoon being lowered into a cup, an apple being placed into a bowl, a toy car being dropped into a bucket), and abstract from them a general rule (e.g., “in every case, the object being lowered into the container is smaller than its opening”). Infants do not gratuitously compare events, in search of abstract truths. Rather, infants are designed to solve concrete problems: they look for abstract truths when challenged to do so through exposure to unpredicted contrastive outcomes: why sometimes this outcome, and sometimes not?

Lesson 5: The Ages at which Rules are Identified Depend in Part on Exposure to Appropriate Outcome and Condition Data

It was suggested in Lesson 4 that the acquisition of a variable typically begins with infants noticing contrastive outcomes they cannot predict based on their current knowledge, and then searching for the conditions that map onto these outcomes. This description leaves many questions unanswered about the processes involved in the identification of variables. Despite its limitations, however, this description does make clear some of the factors likely to affect the *ages* at which variables are identified in different event categories. Two such factors are discussed below.

Exposure to relevant outcomes

If it is true that infants begin the process of identifying a variable when they become aware of contrastive outcomes for the variable, then it follows that the age at which a variable is identified will depend in part on the age at which infants are exposed to and register contrastive outcomes for the variable.

We saw an example of this in Lesson 4 when discussing young infants' responses to transparent occluders (e.g., Johnson & Aslin, 2000; Luo & Baillargeon, 2001d). Infants younger than 7 months of age cannot register the contrastive outcomes for the variable transparency in occlusion events because they cannot *see* clear or transparent occluders. Only when the infants' visual system has matured sufficiently to enable them to detect such surfaces, at about 7 months of age, do they realize that objects can be visible *through* the surfaces.

In the preceding case, infants are exposed to the contrastive outcomes for a variable in their everyday life, but cannot register these outcomes because of visual limitations. In many other cases, however, infants possess sufficient visual ability to detect the contrastive outcomes for the variables – but it just so happens that they are rarely exposed to these outcomes in their daily lives.

To illustrate, consider the finding, discussed in Lesson 2, that infants do not identify the support variable amount of contact until about 6.5 months of age (e.g., Baillargeon et al., 1992). We have suggested that infants do not acquire this variable sooner because they are not exposed sooner to appropriate contrastive outcomes. In everyday life, infants often see their caretakers place objects on supports (e.g., plates on tables, pots on burners, or bottles on counters). However, in most instances, the objects remain stable when released; only in rare accidental cases do they fall. Hence, it is typically not until infants themselves begin to deposit objects on supports (after about 6 months of age, when they learn to sit independently; e.g., Rochat, 1992) that they have the opportunity to notice that objects placed on supports sometimes remain stable and sometimes do not. At that point, infants begin to seek out the conditions that are responsible for these different outcomes, and eventually come to the realization that an object on a support can be stable when a large but not a small portion of its bottom surface rests on the support.

Availability of data on relevant conditions

Another factor likely to affect the age at which infants identify a variable is how easy it is for them, after they are exposed to the relevant contrastive outcomes, to uncover the conditions that map onto these outcomes.

To illustrate, consider the finding, discussed in Lesson 3, that infants do not identify height as a containment variable until about 7.5 months of age (e.g., Hespos, 1998; Hespos & Baillargeon, 2001a). In order to identify this variable, infants must be able to encode information about the heights of objects and containers. As we saw in Lesson 2, prior research (e.g., Baillargeon, 1994, 1995) suggests that, when infants begin to reason about a continuous variable in an event category, they can do so qualitatively, but not quantitatively: they cannot encode and remember information about absolute amounts. To encode information about the heights of objects and containers qualitatively, infants must compare them as they stand side by side. Unfortunately, infants may witness relatively few instances in which objects are placed first next to and then inside containers; caretakers will more often insert objects directly into containers, allowing infants no opportunity to compare their heights.

In the scenario outlined here, infants would thus notice that objects placed inside containers sometimes do and sometimes do not protrude above the containers. However, infants would have few opportunities to gather data about the relative heights of the objects and containers, because they would rarely see (perhaps until they themselves begin placing objects inside containers) the objects standing next to the containers.

The preceding speculations suggest possible explanations for the *décalages* described in Lesson 3 in infants' identification of similar variables across event categories. Consider, for example, the findings that infants identify height as an occlusion variable at about 3.5 months of age (e.g., Aguiar & Baillargeon, in press; Baillargeon & DeVos, 1991), and as a containment variable at about 7.5 months of age (e.g., Hespos, 1998; Hespos & Baillargeon, 2001a). One possibility is, of course, that infants observe many more occlusion than containment events in their daily lives, and hence learn about occlusion events earlier. However, another possibility is that infants can more easily collect qualitative data about the relative heights of objects and occluders than of objects and containers. In the case of occlusion, infants will not only see objects being lowered from above behind occluders – they will also see objects being pushed from the side behind occluders (e.g., as when a parent slides a cup behind a teapot, or a sibling pushes a toy car behind a box). In these side occlusions, it will typically be possible for infants to qualitatively compare the heights of the objects and their occluders; infants will then be in a position to begin mapping conditions onto outcomes.

The importance placed here on the availability of qualitative observations for the identification of continuous variables makes a number of interesting predictions. For example, this approach suggests that, in containment events, infants should learn the variable *width* before height, because each time an object is lowered inside a container infants can compare their relative widths. And indeed, recent findings (e.g., Aguiar & Baillargeon, 2000, 2001; Baillargeon & Brueckner, 2000; Sitskoorn & Smitsman, 1995) indicate that

infants do identify width before height as a containment variable, at about 4 months of age.

Additional remarks

The general approach adopted here makes a strong experimental prediction: in cases where infants cannot acquire a variable because they are rarely exposed in their daily lives to appropriate outcome or condition data for the variable, then deliberate exposure to such data in the laboratory should result in infants' acquisition of the variable. In other words, it should be possible to "teach" infants variables they have not yet had the opportunity to identify by showing them observations from which to abstract the variables.

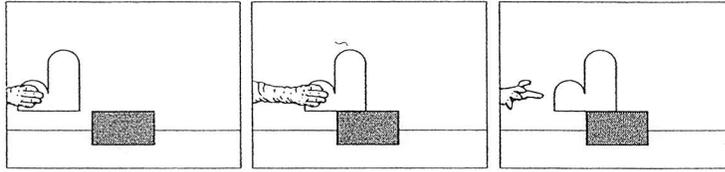
To test this approach, my colleagues and I have undertaken a number of "teaching" experiments. In some of these experiments (Baillargeon et al., 2001; Baillargeon, Fisher, & DeJong, 2000; for reviews, see Baillargeon, 1998, 1999), we attempted to teach 11.5- and 11-month-old infants the variable proportional distribution in support events; as was discussed in Lesson 2, this variable is typically not identified until about 12.5 months of age (e.g., Baillargeon, 1995). In other, more recent experiments (Wang & Baillargeon, 2001e), we attempted to teach 9-month-old infants the variable height in covering events; as was mentioned in Lesson 3, this variable is usually not identified until about 12 months of age (McCall, 2001; Wang & Baillargeon, 2001b; Wang & Paterson, 2000; Wang et al., 2001). All of these experiments have yielded positive results; due to space constraints, however, only one of the support teaching experiments is described here.

In this experiment (Baillargeon et al., 2001), 11-month-old infants received three pairs of teaching trials; each pair of trials involved a box-stays and a box-falls event (see figure 3.7). In both events, an experimenter's gloved hand deposited an asymmetrical box on a platform in such a way that half of the box's bottom surface rested on the platform. In the box-stays event, the larger end of the box was placed on the platform, and the box remained stable when released. In the box-falls event, the smaller end of the box was placed on the platform, and the box fell when released. Three different asymmetrical boxes were used in the three pairs of teaching trials. Following these trials, the infants saw two static test displays involving a novel, L-shaped box on a platform. In each display, half of the box's bottom surface rested on the platform; either the larger end (adequate-support display) or the smaller end (inadequate-support display) of the box was supported. The experiment thus examined whether the infants could form a new condition-outcome rule during the teaching trials that would enable them to detect the violation in the inadequate-support test display.

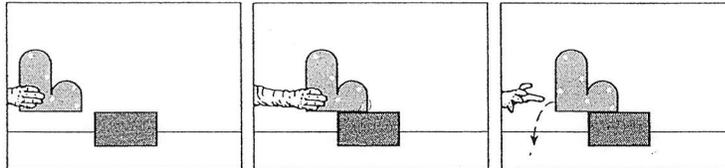
After observing the six teaching trials, the infants looked reliably longer at the inadequate- than at the adequate-support test display. These results suggested that the teaching trials presented the infants with appropriate outcome and condition data from which to abstract the new variable proportional distribution. In each pair of trials, the infants saw contrastive outcomes they could not explain based on their current knowledge. They expected the box to remain stable in both the box-stays and box-falls events, because in each case half of the box's bottom surface was supported; and yet the box fell

Teaching events

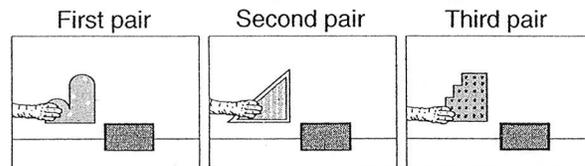
Box-stays event



Box-falls event

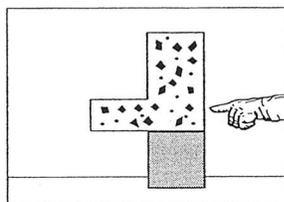


Boxes used in each pair of teaching events



Test displays

Adequate-support display



Inadequate-support display

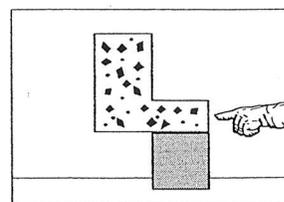


Figure 3.7 Schematic drawing of the teaching events and test displays in Baillargeon et al. (2001)

in the box-falls event. How could this be? By comparing the box-stays and box-falls events, the infants could come to the realization that the box remained stable when the proportion of the entire box (not just of the box's bottom surface) resting on the platform was greater than that off the platform.

Much additional research is needed to flesh out these initial findings (for additional findings, see Baillargeon, 1998, 1999; and Baillargeon et al., 2000, 2001). For example, we need to find out whether the rules infants form in the laboratory are permanent or transitory. Would the same infants tested a week later still show an understanding of proportional distribution in support events? Furthermore, under what teaching conditions do infants show evidence of learning, and under what conditions do they not? Are we correct in saying that infants must be exposed to appropriate outcome and condition data in order to learn? Would infants fail in our support teaching experiment if shown, for example, only the box-stays or the box-falls events during the teaching trials? The answers

to these and related questions should shed light on the conditions under which infants can and cannot add to their physical knowledge.

Lesson 6: New Rules are Integrated with Infants' Prior Knowledge

Up to this point, we have suggested that infants form narrow event categories and, for each category, identify a sequence of condition–outcome rules that enable them to predict and interpret outcomes more and more accurately over time. We have also proposed that the identification of rules depends in part on exposure to appropriate outcome data (learning is typically triggered by the detection of outcome variation not predicted by infants' current knowledge) and appropriate condition data (in order to identify the conditions that map onto the contrastive outcomes they have observed, infants must have access to qualitative condition data, for both discrete and continuous variables).

But is this all there is to infants' learning? Are there no limits or constraints on the condition–outcome rules they are capable of forming? Could infants acquire *any* new rule, as long as they were exposed to appropriate condition and outcome data? Could one successfully teach infants, for example, that tall objects can be hidden behind short but not tall occluders, that wide objects can be lowered inside narrow but not wide containers, or that asymmetrical objects can be stable when their smaller but not their larger ends are supported?

A recent experiment addressed this last question. In this experiment, we attempted to teach 11-month-old infants the *reverse* of the proportional distribution rule (Baillargeon et al., 2001). As before, the infants received three pairs of teaching trials; however, the box now fell when its larger end rested on the platform, and remained stable when its smaller end rested on the platform. Had the infants abstracted the reverse rule “an object on a support is stable when the proportion of the object off the support is greater than that on the support,” they should have looked reliably longer at the adequate- than at the inadequate-support test display. However, the infants tended to look equally at the two displays, suggesting that they had *not* learned this reverse rule. Negative results were also obtained in a similar experiment with older, 11.5-month-old infants (Baillargeon et al., 2001).

What should one make of these results? It seems likely that the infants could not learn the reverse proportional distribution rule because they could not integrate it with the knowledge they brought to the laboratory. But how should this prior knowledge be characterized? Two very different interpretations are possible, depending on one's model of infants' knowledge acquisition process and the constraints that limit it.

A *first* model is that infants' knowledge acquisition process consists primarily in the detection of *statistical regularities* in the environment. On this view, there would be few constraints on the condition–outcome rules infants can learn, and these would involve mainly basic limitations in infants' perception, memory, and information-processing abilities (e.g., infants cannot learn regularities they cannot see, or see too infrequently to remember). According to this first, *correlation-based* model, it should in principle be possible for infants to learn that an asymmetrical object typically falls when its larger end is

supported, or any such reverse rule, as long as they were exposed to the necessary condition and outcome data to learn it. The fact that the infants in our experiment failed to learn the reverse proportional distribution rule (Baillargeon et al., 2001) could be attributed to their having already begun accumulating observations about asymmetrical objects prior to coming to the laboratory. When shown novel observations *consistent* with their own (in our initial, successful teaching experiment), infants were able to abstract a new rule, presumably because they now possessed sufficient data to do so. However, when shown novel observations *inconsistent* with their own (in our reverse teaching experiment), infants could learn nothing – they could only note the inconsistency between their past and present observations.

A *second*, radically different model of infants' knowledge acquisition process is that it consists mainly in the detection of *causal regularities* (e.g., Keil, 1991, 1995; Leslie, 1994, 1995; Pauen, 1999; Wilson & Keil, 2000). On this view, only regularities for which infants could construct explanations based on their prior knowledge would be accepted as condition–outcome rules. These explanations would obviously tend to be shallow and incomplete; nevertheless, they would require some degree of causal analysis, which would place severe limits on the condition–outcome rules infants could learn. According to this second, *explanation-based* model, infants would be able to learn the proportional distribution rule because they could construct an explanation for it; and they would *not* be able to learn the reverse rule because they could not make sense of it in terms of their prior knowledge (for a discussion of explanation-based machine learning, see DeJong, 1988, 1993, 1997).

What causal analysis might have enabled the infants in our initial teaching experiment to identify proportional distribution as an acceptable rule? One possibility is that the infants brought to bear their knowledge of *weight*. Upon seeing the teaching trials, the infants might have reasoned along the following lines: (1) the larger end of each asymmetrical box very likely weighed more than the smaller end; (2) when off the platform, the heavier end of the box pulled downward, and pulled the lighter end downward as well, causing the box to fall; and (3) when off the platform, the lighter end of the box pulled downward, but could not pull the heavier end downward as well, so that the box remained stable. For the infants in the reverse teaching experiment, the fact that each asymmetrical box fell when its heavier end was supported and remained stable when not, would have been impossible to reconcile with their knowledge of weight and of the forces exerted by heavier and lighter objects, so that no learning could occur.

How could one decide between the correlation- and explanation-based interpretations of the results of our teaching experiments? As a first step, Su-Hua Wang, Cindy Fisher, Jerry DeJong, and I are planning to test the hypothesis, derived from the explanation-based interpretation, that the infants in our initial teaching experiment used their knowledge of weight to construct an explanation for the teaching trials. This new experiment will be identical to our initial experiment with one exception: prior to the test session, the infants will hold, one at a time, the three asymmetrical boxes used in the three pairs of teaching trials. For half of the infants (consistent condition), the larger end of each box will be heavier than the smaller end; for the other infants (inconsistent condition),

the reverse will be true. Next, all of the infants will receive exactly the same teaching and test trials as in our initial teaching experiment.

Evidence that the infants in the consistent but not the inconsistent condition learn the proportional distribution rule (i.e., look reliably longer at the inadequate- than at the adequate-support test display) would be important for two reasons. First, it would suggest that infants do bring to bear weight information when attempting to make sense of the teaching trials. Second, and more generally, it would suggest that the acquisition of physical knowledge in infancy involves the detection of causal, rather than merely statistical, regularities. Such evidence would thus support an explanation-based, rather than a correlation-based, model of learning.

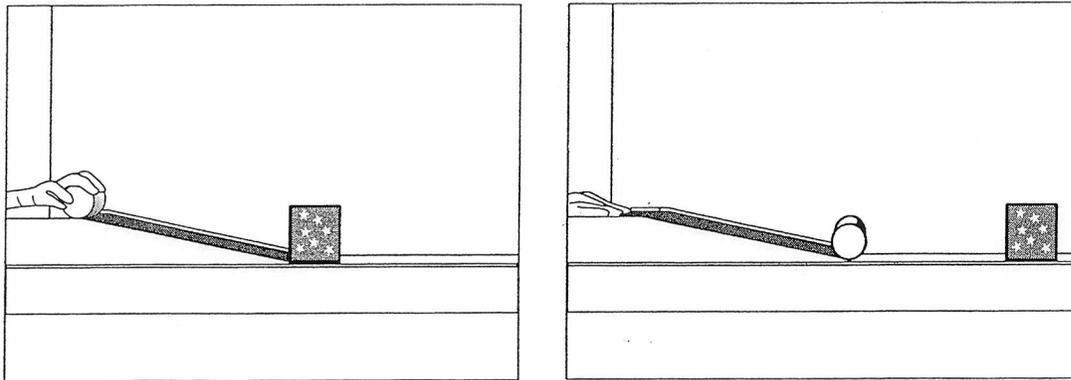
Additional remarks

What evidence is there that infants attend to the weights of objects, and hold different expectations for events involving objects of different weights? In a recent series of experiments, Su-Hua Wang and I examined 10-month-old infants' ability to reason about collision events involving heavier and lighter objects (Wang, 2001; Wang & Baillargeon, 2001d). The infants in one experiment were first given two cylinders to hold, one at a time. The cylinders were identical in size but differed in color: one was blue and one was yellow (see figure 3.8). For half of the infants (same-weight condition), the two cylinders were equally light; for the other infants (different-weight condition), the blue cylinder was again light but the yellow cylinder was much heavier (it was simply laid on the infants' lap, as they typically could not hold it up). All of the infants first saw a familiarization event in which the yellow cylinder rolled down a ramp and hit a box, causing it to move a short distance. Next, the infants saw a test event in which the blue cylinder rolled down the ramp and hit the same box, which now remained stationary.

The infants in the same-weight condition looked reliably longer during the test event than did those in the different-weight condition. This result suggested two conclusions. First, the infants in the same-weight condition (1) remembered that the yellow cylinder weighed about the same as the blue one and (2) reasoned that, if the yellow cylinder could displace the box, then the blue cylinder should also be able to do so. Second, the infants in the different-weight condition (1) remembered that the yellow cylinder was heavier than the blue one and (2) appreciated that the heavier yellow cylinder might be able to displace the box, and the lighter blue cylinder fail to do so.

Parallel results were obtained in a second experiment that was similar to our initial experiment with two exceptions: first, the infants no longer felt the cylinders, they could only inspect them visually; and second, the yellow cylinder was now either the same size as (same-size condition) or much larger than (different-size condition) the blue cylinder. In this experiment, the knowledge that the yellow cylinder was as light as (same-size condition) or heavier than (different-size condition) the blue cylinder thus had to be *inferred* – it was no longer available through direct proprioceptive feedback. Nevertheless, the results were analogous to those of the first experiment: the infants in the same-size condition looked reliably longer than did those in the different-size condition. This result

Familiarization event



Test event

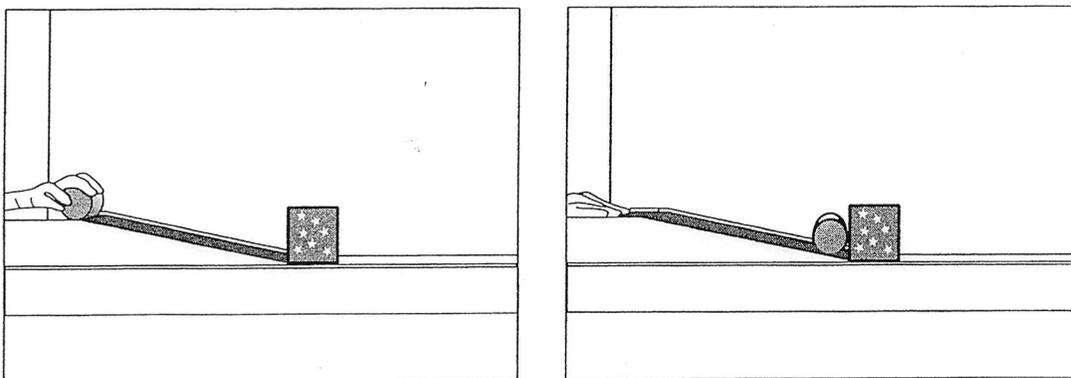


Figure 3.8 Schematic drawing of the test events shown in the same- and different-weight conditions in Wang and Baillargeon (2001d)

suggested that the infants in the same-size condition (1) inferred that the two cylinders were about the same weight and (2) reasoned that, if the yellow cylinder could displace the box, then the blue cylinder should be able to do so as well. For their part, the infants in the different-size condition (1) inferred that the large yellow cylinder was heavier than the smaller blue one and (2) recognized that the heavier yellow cylinder might be able to displace the box, and the lighter blue cylinder might not.

Consistent results were obtained in further experiments in which the weight of the box, rather than that of the cylinder, was manipulated (Wang, 2001; Wang & Baillargeon, 2001d). Together, the results of these experiments suggest that 10-month-old infants attend to and remember the weights of objects, and recognize that events may have different outcomes depending on the weights of the objects involved. In particular, infants appear to possess an expectation that heavier objects can both exert and resist greater forces than can lighter objects. In light of these findings, it does not seem implausible that older, 11-month-old infants should extend this expectation to make sense of support events involving asymmetrical objects, such as those in the teaching trials of our initial teaching experiment (Baillargeon et al., 2001).

Lesson 7: Innate Concepts Guide Infants' Interpretation of their Physical Representations

In our previous lessons, it was suggested that infants sort events into distinct categories and, for each category, identify a sequence of variables (or condition–outcome rules) that result in more and more accurate predictions and interpretations over time. Furthermore, the identification of a new variable may depend on (1) infants' exposure to appropriate outcome and condition data and (2) infants' ability to use their prior knowledge to construct a causal explanation for the variable. In Lesson 7, we begin to consider the possible contributions of *innate physical knowledge* to this acquisition process.

Several researchers (e.g., Carey & Spelke, 1994, 1996; Gelman, 1990; Gopnik & Wellman, 1994; Keil, 1991, 1995; Leslie, 1994, 1995; Spelke, 1994; Spelke et al., 1992; Wellman & Gelman, 1992, 1998) have proposed that infants' naïve physics is a foundational or core domain, and that as in other such domains (e.g., language, number, and naïve psychology), reasoning is facilitated by innate concepts. For example, Spelke (1994; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995) has proposed that core principles of *continuity* (objects exist and move continuously in time and space) and *solidity* (two objects cannot exist in the same space at the same time) constrain from birth infants' interpretations of physical events.

It might be argued that much of the evidence discussed in our previous lessons calls into question Spelke's (1994; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995) proposal. For if infants possessed core principles of continuity and solidity, shouldn't they detect all salient violations of these principles? And yet that is clearly not the case: as we have seen, infants often fail to detect what appear to adults to be marked continuity and solidity violations. To give a few examples, recall that 3-month-old infants are not surprised when a tall object becomes fully hidden behind a short occluder (Aguiar & Baillargeon, in press; Baillargeon & DeVos, 1991); that 6.5-month-old infants are not surprised when a tall object is fully lowered inside a short container (Hespos & Baillargeon, 2001a); that 8.5-month-old infants are not surprised when an object lowered inside a transparent container is not visible through the front of the container (Luo, 2001; Luo & Baillargeon, 2001b); and that 9-month-old infants are not surprised when a tall object becomes fully hidden under a short cover (Wang & Paterson, 2000; Wang et al., 2001). How could these negative findings be reconciled with the claim that core principles of continuity and solidity constrain infants' interpretations of events from birth?

My colleagues and I have suggested that such a reconciliation is in fact possible, and depends on a number of assumptions about the nature and development of infants' representations of physical events (e.g., Aguiar & Baillargeon, in press; Baillargeon, in press; Hespos & Baillargeon, 2001b). Below, I describe these assumptions, and then return to the issue of how core principles might contribute to infants' interpretations of events.

Three Assumptions about Infants' Physical Representations

My collaborators and I have arrived at three main assumptions about infants' representations of events (e.g., Aguiar & Baillargeon, in press; Baillargeon, in press; Hespos & Baillargeon, 2001b). First, we assume that, when observing physical events, infants build representations – which we call *physical representations* – that focus on the physical properties, displacements, and interactions of the objects and surfaces within the events.³ Second, we assume that infants' physical representations of events are by no means faithful copies of the events: they are abstract descriptions that include some but not all of the physical information in the events. Third, we assume that how much information infants include in their physical representation of an event depends in large part on their knowledge of the variables likely to affect the event.

To put these assumptions more concretely, we suppose that, when watching an event, infants first categorize it, and then access their knowledge of the event category selected. This knowledge specifies what variables should be attended to, and thus what information should be included in the physical representation of the event. To illustrate, this means that 3.5-month-old infants who see an object being lowered behind a container (occlusion event) will include information about the relative heights of the object and container in their physical representation of the event, because they have already identified height as an occlusion variable (Aguiar & Baillargeon, in press; Baillargeon & DeVos, 1991). In contrast, 3.5-month-old infants who see an object being lowered inside a container (containment event) will not encode the relative heights of the object and container, because they have not yet identified height as a containment variable (Hespos & Baillargeon, 2001a).

According to the preceding assumptions, infants' physical representations would thus be initially very sparse. As infants discover the importance of different variables in predicting outcomes, they would begin to include this variable information, thus achieving richer and more detailed physical representations over time.

A Case of Impoverished Physical Representations

If one accepts the assumptions described in the previous section, then it becomes clear how infants might possess core continuity and solidity principles and still fail to detect salient violations of these principles. Infants' core principles, like all of their physical knowledge, can operate only at the level of their representations (infants, like adults, do not apply their knowledge directly to events, only to their representations of the events). It follows that infants can succeed in detecting violations of their continuity and solidity principles *only* when the key information necessary to detect these violations is included in their physical representations. Infants' principles can only guide the interpretation of information that has been included in their representations; information that has not been represented cannot be interpreted.

To illustrate how incomplete physical representations could lead infants to ignore violations of their continuity and solidity principles, consider one of the findings discussed in Lesson 2, that 3-month-old infants are not surprised when a tall object fails to appear between two screens connected at the bottom by a short strip (Aguiar & Baillargeon, in press; Baillargeon & DeVos, 1991). What is being suggested is that, when observing such an event, 3-month-old infants typically do not include information about the relative heights of the object and occluder in their physical representation of the event. Thus, when infants apply their continuity principle to their incomplete physical representation of the event, they have no basis for predicting that a portion of the object should be visible above the short strip between the screens.

Additional remarks

It was suggested above that young infants might possess core continuity and solidity principles and still fail to detect violations of these principles, because of sparse or incomplete physical representations. This approach makes a number of intriguing predictions. In particular, it suggests that if infants could be temporarily induced or “primed” to include key variable information in their representations of physical events, they should then be in a position to detect continuity or solidity violations they would otherwise have been unable to detect. On the present view, it should not matter whether infants include variable information in their physical representation of an event because (1) they have been primed to do so by some contextual manipulation or (2) they have already identified the variable as relevant to the event category. In either case, the information, once represented, should be subject to infants’ continuity and solidity principles.

To test these ideas, Su-Hua Wang and I recently conducted a “priming” experiment with 8-month-old infants, focusing on the variable height in covering events (Wang & Baillargeon, 2001c). As we saw in Lesson 3, this variable is typically not identified until about 12 months of age (McCall, 2001; Wang & Baillargeon, 2001b; Wang & Paterson, 2000; Wang et al., 2001). In the experiment of Wang et al. (2001), for example, 9-month-old infants saw a tall- and short-cover test event. At the start of the tall-cover event, a tall cylindrical cover stood next to a tall cylindrical object. An experimenter’s gloved hand grasped the cover, lifted it over the object, and lowered it to the apparatus floor, thereby fully hiding the object. The short-cover event was identical except that the cover was only half as tall, so that it should have been impossible for the cover to fully hide the object. The infants tended to look equally at the events, suggesting that they had not yet identified height as a covering variable.

The 8-month-old infants in our priming experiment (Wang & Baillargeon, 2001c) saw the same test events as in Wang et al. (2001) with one exception: each test event was preceded by a brief pretrial designed to prime the infants to include information about the relative heights of the cover and object in their physical representation of the covering event (see figure 3.9). Because infants aged 3.5 months and older attend to height information in occlusion events (Baillargeon & DeVos, 1991; Hespos & Baillargeon, 2001a), the pretrial presented the infants with an *occlusion* event involving the cover and object. At the start of each pretrial, the cover rested next to the object; after a pause, the

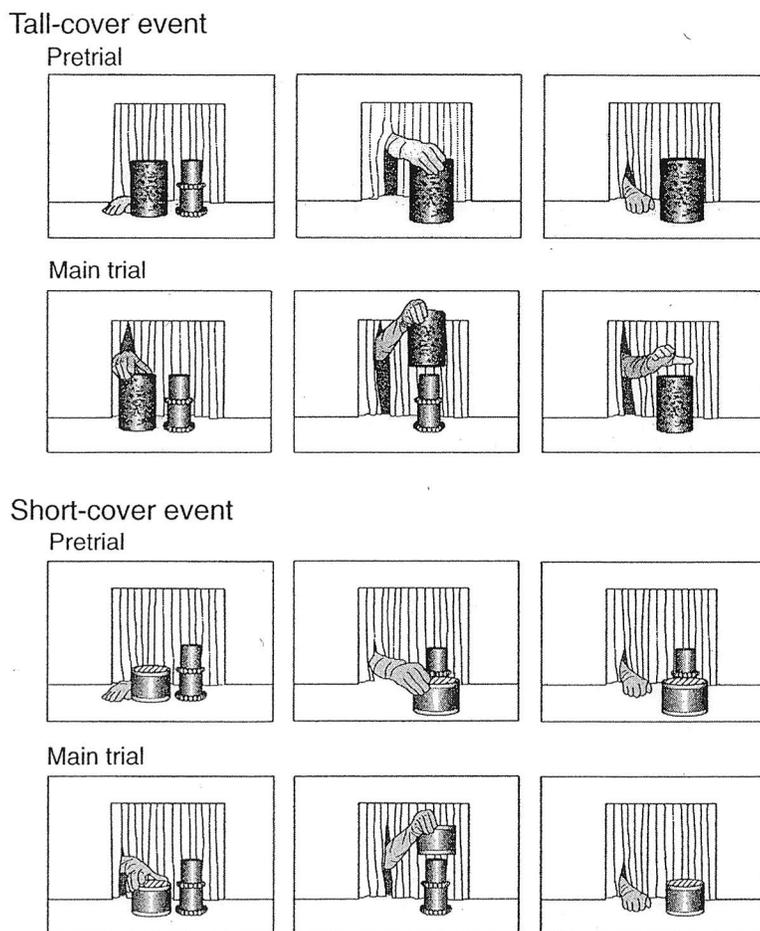


Figure 3.9 Schematic drawing of the test events shown in the experimental condition in Wang and Baillargeon (2001c)

hand grasped the cover, slid it in front of the object, and then returned to the apparatus floor. The tall cover occluded all of the object, the short cover only its bottom portion. The cover remained in its new position until the computer signaled that the infant had looked at the display for five cumulative seconds. The hand then returned the cover to its original position next to the object, and the test trial proceeded exactly as in Wang et al. (2001). Infants in a control condition received a similar pretrial at the start of each test event, except that the cover was simply slid forward and thus never occluded the object.

Our reasoning in designing the pretrials was as follows: after the infants in the experimental condition included the relative heights of the cover and object in their physical representation of the occlusion event, they might be inclined to do the same – or they might have this information still available – when forming their physical representation of the covering event. This information would then be subject to the infants' core principles of continuity and solidity, making it possible for them to detect – at a much younger age than they would otherwise – the violation in the short-cover test event.

As expected, the infants in the experimental condition looked reliably longer at the short- than at the tall-cover test event, whereas those in the control condition looked about

equally at the events, like the 9-month-old infants in our original experiment (Wang & Paterson, 2000; Wang et al., 2001). These results are important for three reasons. First, they provide strong support for the notion that infants' failures to detect continuity and solidity violations reflect impoverished physical representations: what is not represented cannot be interpreted. Second, the results give weight to our analysis of infants' physical representations, and more specifically to our proposal that infants do not routinely include information about a variable in a physical representation of an event until they have identified the variable as pertinent to the event category. Finally, our results (see also Chapa & Wilcox, 1999, for a different type of priming experiment) raise many new questions about priming issues: for example, what manipulations are helpful for priming variables and what manipulations are not; what are the long-term effects of successful priming experiences; and finally, what are the similarities and differences between priming and teaching experiences (as in Baillargeon et al., 2001, which was discussed in Lesson 5).

Lesson 8: Innate Concepts Guide Even Very Young Infants' Interpretations of their Physical Representations

It was suggested in Lesson 7 that (1) principles of continuity and solidity constrain from birth infants' interpretations of physical events (e.g., Spelke, 1994; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995); (2) these principles operate at the level of infants' physical representations; and (3) because these representations are at first very limited – variables not yet identified as relevant to an event category are typically ignored – infants cannot detect continuity and solidity violations involving these variables. This approach makes one strong prediction: even very young infants should be able to detect continuity and solidity violations, as long as these do not involve variable information.

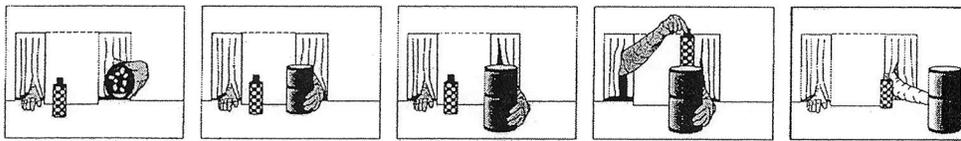
In the first months of life, infants' physical representations most likely include only basic spatial, temporal, and possibly mechanical (Leslie, 1994, 1995) information. As they grow older, infants certainly become better at representing this information (e.g., Arterberry, 1997; Yonas & Granrud, 1984). But the key point here is this: even very young infants should detect continuity and solidity violations that involve only the basic information they can represent.

Over the past ten years, evidence has been slowly but steadily accumulating that 2.5-month-old infants (the youngest that have been successfully tested to date with violation-of-expectation tasks) interpret a wide range of physical events in accord with continuity and solidity principles (e.g., Aguiar & Baillargeon, 1999, in press; Hespos & Baillargeon, 2001b; Spelke et al., 1992; Wilcox et al., 1996). Due to space constraints, only one experiment, on containment events, is described here.

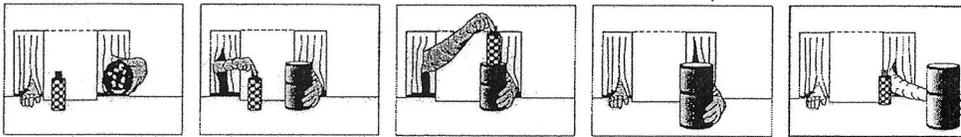
Containment events

Sue Hespos and I recently asked whether 2.5-month-old infants realize that an object that has been lowered inside a container should move with the container when the latter is slid to a new location (Hespos & Baillargeon, 2001b).

Behind-container condition



Inside-container condition

**Figure 3.10** Schematic drawing of the test events in Hespos and Baillargeon (2001b)

The infants were assigned to a behind- or an inside-container condition and saw a single test event (see figure 3.10). At the start of the event shown in the *behind-container* condition, an experimenter's gloved right hand rested on an apparatus floor next to a cylindrical object; the same experimenter's gloved left hand grasped the midsection of a tall cylindrical container standing to the right of the object. To start, the left hand rotated the container forward so that the infants could see its open top and hollow interior. After a few seconds, the container was placed upright next to the object and then slid forward. Next, the right hand grasped the object, moved it above and behind the container, and lowered it until it disappeared behind the container. The left hand then moved the container to the right, revealing the object standing behind it. The infants in the *inside-container* condition saw the same test event except that the object was lowered inside the container before the latter was moved forward; hence, it should have been impossible for the object to be revealed when the container was moved to the right.

Prior to the test trials, the infants in the behind- and inside-container conditions received baseline trials identical to the test trials with one exception: the container was never moved to the right to reveal the object behind it. The baseline trials provided an assessment of whether the infants had an intrinsic preference for seeing the object being lowered behind or inside the container. The test events shown in the two conditions were perceptually identical except for the fact that in one condition the container was moved forward and the object was then lowered behind it, whereas in the other condition the object was lowered inside the container which was then moved forward. Because the baseline events also were perceptually identical except for this difference, the data from the baseline trials could be used to assess whether seeing the object lowered behind or inside the container was intrinsically more attractive to the infants.

During the baseline trials, the infants in the inside- and behind-container conditions tended to look equally; during the test trials, however, the infants in the inside-container condition looked reliably longer than did those in the behind-container condition. Together, these results suggested that the infants (1) believed that the object continued to exist after it disappeared from sight; (2) remembered whether it had been lowered inside or behind the container; (3) realized that the object, when lowered inside the container, could not pass through its closed sides and thus had to move with it when

displaced; and therefore (4) expected the object to be revealed when the container was moved to the right in the behind- but not the inside-container condition.

It is certainly remarkable that infants as young as 2.5 months of age can detect the violation in the inside-container test event. However, to return to our earlier discussion, such a success is precisely what should be possible to young infants whose physical representations include only basic, pre-variable information, but whose interpretations of this information are guided by continuity and solidity principles. Detecting the violation in the inside-container test event required only keeping track of the locations of the object and container over time, and realizing that the object continued to exist when inside the container (an assumption that would be suggested by the infants' continuity principle), and moved with it when displaced (an assumption that would be suggested by the infants' solidity principle).

Additional remarks

In Lessons 7 and 8, we saw that whether infants succeed in detecting continuity and solidity violations depends on two main factors: first, what information infants include in their physical representations of events; and second, what information must be attended to for the violations to be detected. Thus, even very young infants may detect violations if these involve only the basic, pre-variable information they encode about events; and much older infants may fail to detect violations if these involve variables they have not yet identified and thus tend to ignore when forming their physical representations.⁴

In the preceding discussion, we have focused primarily on how infants' core principles of continuity and solidity might guide their interpretations of events. Here, I would like to consider another role for these and other innate concepts, one which is related to our discussion in Lesson 6 of the explanation-based approach to learning. According to this approach, infants only learn rules they can make sense of or explain in terms of their prior physical knowledge. To the extent that this prior knowledge includes innate concepts, then these concepts could also contribute to infants' acquisition of rules, by helping them build acceptable explanations for novel variables.

It is easy to see how infants' principles of continuity and solidity might contribute to their causal analyses of variables in occlusion, containment, covering, and other physical events. For example, it might be suggested that infants readily accept height as an occlusion variable (e.g., Aguiar & Baillargeon, in press; Baillargeon & DeVos, 1991), because it is highly consistent with their notion of continuity: if an object continues to exist when behind an occluder, and is taller than the occluder, then it should indeed extend above the occluder. Similarly, consider the variable width in containment events (e.g., Aguiar & Baillargeon, 1998, 2000b, 2001; Baillargeon & Brueckner, 2000; Sitskoorn & Smitsman, 1995): infants' notion of solidity would suggest that if an object is wider than the opening of a container, then it should not be possible for the object to be lowered through this opening, for to do so would mean passing through the solid body of the container.

Although we have focused here on Spelke's (1994; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995) core principles of continuity and solidity, other innate notions might

also contribute to infants' explanations for novel variables. Consider, for example, Leslie's (1994, 1995) proposal that infants are born with a primitive notion of *force*. At about 5.5 to 6.5 months of age (e.g., Kotovsky & Baillargeon, 1994, 1998), infants come to realize that, in collision events, larger (heavier) objects displace stationary objects farther than do smaller (lighter) objects. One can see that such a rule (but not its reverse) would appear consistent with a simple notion of force as a unidirectional and incremental application of energy: all other things being equal, greater forces should yield greater effects. Infants' notion of force might also be implicated in their learning about support events. Recall our discussion, in Lesson 6, of how infants might have interpreted the teaching trials in our initial teaching experiment (Baillargeon et al., 2001). Implicit in this causal analysis is the notion that a heavier object (or portion of an object) can exert and resist a greater force than can a lighter object (or portion of an object). Other support variables might be interpreted by infants in similar terms. For example, consider the variable "middle or side contact," which is identified at about 8.5 months of age (Dan et al., 2000; Huettel & Needham, 2000; Wang & Baillargeon, 2001a), and was discussed in Lesson 2. Upon seeing that an object does not fall despite the fact that only its middle portion (e.g., the middle 30 percent of its bottom surface) is supported, a 6.5-month-old infant might assume that (1) the weight of the object is evenly distributed along its bottom surface and (2) when the middle portion of the object is supported, leaving two identical portions unsupported on either side, then each of these unsupported portions exerts an equal, and opposite, pull downward, thereby keeping the object balanced over the supported portion.

Pondering what causal analyses might underlie infants' acquisition of different variables, and what role innate notions of continuity, solidity, force, and so on, might play in these causal analyses, is a fascinating research exercise. But how can we determine whether our intuitions about infants' causal analyses are correct? As was discussed in Lesson 6, our current approach is to conduct teaching experiments in which infants (1) are exposed to appropriate outcome and condition data to learn a new variable, but (2) are given additional information that is either consistent or inconsistent with infants' hypothesized causal analysis of these outcome and condition data. We have already seen an example of such a teaching experiment, in Lesson 6: recall the experiment in which infants receive teaching trials on the support variable proportional distribution after holding asymmetrical boxes whose larger end is either heavier (consistent condition) or lighter (inconsistent condition) than their smaller end.

Parallel experiments can be designed for other variables. For example, consider once again the support variable middle or side contact discussed above. If our description of infants' causal analysis of this variable is correct, then it should be possible to teach infants that an object can remain stable (or balanced) with two *identical* unsupported portions extending on either side of a small supported portion; however, it should be very difficult to teach infants that an object can be stable with two *non-identical* unsupported portions extending on either side of a small supported portion. In the latter case, the object could in fact still be physically stable, but infants would no longer be able to produce a causal analysis in terms of two identical portions exerting equal and opposite downward pulls. A notion of center of mass is necessary to explain such cases of support, and children do not achieve such a notion until much later in development.

Conclusions

How do infants acquire their physical knowledge? The eight lessons discussed in this chapter point to the following picture. From birth, infants build specialized representations of physical events, which we termed physical representations. These physical representations initially include only limited spatial, temporal, and mechanical information. The interpretation of this information is guided by a few core principles, including continuity and solidity.

Over time, infants begin to form distinct event categories, such as occlusion, containment, and support events. For each category, infants identify a sequence of variables (condition–outcome rules) that enable them to predict outcomes more and more accurately over time. The identification of each variable depends on two main factors: (1) infants' exposure to appropriate outcome and condition data (exposure to contrastive outcomes not predicted by infants' current physical knowledge is thought to trigger the identification process; infants then seek out the conditions that map onto these outcomes); and (2) infants' ability to build a causal explanation for these condition–outcome data, based on their prior knowledge. This prior knowledge is presumed to include both the knowledge infants have already accumulated about the category, and the innate principles and concepts mentioned earlier.

The model just outlined shares many features in common with descriptions offered by other researchers. First, the notion that infants' naïve physics is a foundational or core domain, with its own innate concepts and specialized learning mechanism, is one with a long history in the field of conceptual development (e.g., Carey, 1985; Gelman, 1990; Keil, 1989; Spelke et al., 1992; Wellman & Gelman, 1992). Second, the idea that infants seek out causal rather than mere statistical regularities, and use their prior knowledge to build shallow explanations for these regularities, is highly consistent with the work of Keil, Leslie, Pauen, and their colleagues (e.g., Keil, 1991, 1995; Leslie, 1994, 1995; Pauen, 1999; Wilson & Keil, 2000). Finally, the notion that infants analyze the outcome and condition data to which they are exposed to abstract new rules about event categories is reminiscent of Mandler's proposal that infants analyze various data to form object categories (e.g., Mandler, 1992, 1998, 2000a; see also Quinn, ch. 4 this volume).

The model outlined here also differs from current models and descriptions in a number of important respects. First, it is not entirely consistent with the recent "theory theory" approach to knowledge acquisition (e.g., Gopnik & Wellman, 1994; Gopnik & Meltzoff, 1997). According to Gopnik and Wellman, for example, an important characteristic of scientists' as well as children's theories "is their coherence . . . changes in one part of the theory have consequences for other parts of the theory" (1994, pp. 260–261). As we saw in Lesson 3, however, the knowledge infants acquire about physical events during the first year of life is very piecemeal: infants' learning about height in occlusion events does not seem to have any consequence for their knowledge of containment or covering events (e.g., Hespos & Baillargeon, 2001a; McCall, 2001; Wang & Baillargeon, 2001b; Wang & Paterson, 2000; Wang et al., 2001); and infants' learning about transparency in occlusion events does not seem to have any distant effect on their knowledge of containment events (e.g., Luo, 2001; Luo & Baillargeon, 2001b, 2001d). In what sense

can an infant who must learn the same variables separately in different event categories be said to possess a theory-like understanding of these categories? To attribute to infants a theory – even an incorrect theory – of the physical world, one would want to see many more generalizations, and many fewer *décalages*, than our research has actually revealed. Infants are in the business of acquiring physical rules – they are not (or not yet) in the business of building a coherent physical theory (see also Wilson & Keil, 2000).

Another feature of the present model that differs from other descriptions concerns the way in which innate concepts are thought to contribute to infants' interpretations of physical events. In our model, innate concepts play an important but still limited role. Infants must discover for themselves, event category by event category and variable by variable, much that is relevant to continuity, solidity, and so on. True, to the extent that several of these variables – say, transparency in occlusion events, width in containment events, and height in covering events – are implied in the principles of continuity and solidity, one can say that infants are doing no more than discovering what they already know. But this process of discovery is an important and a protracted one, and to dismiss it is to ignore much of what happens in the first year of life. Infants are not philosophers who reflect at leisure upon abstract principles and infer from them new truths. Infants are concrete thinkers who draw novel inferences – who discover what they already know – only when challenged to do so by reality, in the form of unexplained variation in outcome.

Notes

1. Occlusion events are events in which an object becomes at least partly hidden behind a nearer object or occluder (e.g., as when a cup is lowered behind a teapot). Support events are events in which an object becomes supported by another object (e.g., as when a plate is placed on a table). Collision events are events in which an object hits another object (e.g., as when a toy car hits a shoe). Finally, containment events are events in which an object is placed inside a container (e.g., as when a ball is lowered inside a box). From an adult perspective, containment of course often involves occlusion. However, this occlusion is of a different form than that defined above: the contained object is hidden because it is lowered inside, not behind, the container. As we will see in Lesson 3, this distinction appears to be crucially important to infants.
2. From the present perspective, a variable is thus tantamount to a dimension; conditions correspond to values on the dimension, with each value (or discernable range of values) being associated with a distinct outcome.
3. Infants no doubt build several representations simultaneously, for different purposes. For example, another representation might focus on the features of the objects in the events, and be used for object recognition and categorization purposes – to ascertain whether these particular objects, or similar objects, have been encountered in the past (e.g., Needham & Modi, 2000; Quinn & Eimas, 1996).
4. It might be suggested that principles of continuity and solidity could be learned in the first few weeks or months of life, as infants observe the world around them. For example, when fixating stationary objects such as a chair, a cup on a table, or a phone on a wall, infants could notice that these objects do not disappear and reappear capriciously, but persist through time, in their same locations. Similarly, when watching moving objects such as a ball rolling across a floor, a parent walking across a room, or even their own hand fluttering back and forth,

infants could notice that these objects follow smooth, continuous paths – they do not abruptly disappear at one point in space and reappear at a different one. On the basis of such, very common, observations, infants could conclude that objects exist and move continuously in time and space.

One difficulty with this alternative account is that it is not consistent with the learning mechanism described in Lessons 4 to 6. We have seen that the acquisition of rules seems to be triggered by exposure to unpredicted variation in outcome – contrastive outcomes not predicted or explained by infants' current knowledge. But in the case of continuity, for example, what contrastive outcomes could be involved? Stationary objects persisting through time, in their same locations, moving objects following continuous paths: such events involve no contrastive outcomes. A learning mechanism that must be triggered by variation in outcome would be incapable of detecting the abstract regularities in these events. Hence, one is left with the following two possibilities: either these general principles, which are present in very young infants and affect their physical reasoning very broadly, are innate, as Spelke (1994; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995) has suggested; or else they are acquired by a learning mechanism very different from that responsible for the bulk of infants' acquisitions about occlusion, containment, support, and other physical events. At the present time, the first possibility seems to us more compelling.