

Processes of Change in Brain and Cognitive Development Attention and Performance XXI

Edited by

Yuko Munakata

Department of Psychology

University of Colorado at Boulder

Boulder, USA

and

Mark H. Johnson

Centre for Brain and Cognitive Development

School of Psychology

Birkbeck, University of London

London, UK

This book is based on the papers presented at the *Twenty-First International Symposium on Attention and Performance* held at Winter Park, Colorado, July 26–Aug 1, 2004.

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford ox2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© The International Association for the Study of Attention and Performance, 2005

The moral rights of the author have been asserted

Database right Oxford University Press (maker)

First published 2006

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,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

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

Library of Congress Cataloging in Publication Data
(Data available)

ISSN 1047-0387

Typeset by SPI Publisher Services, Pondicherry, India

Printed in Great Britain
on acid-free paper by
Biddles Ltd, King's Lynn

ISBN 0-19-856874-6 (Hbk) 978-0-19-856874-2 (Hbk)

10 9 8 7 6 5 4 3 2 1

Under what conditions do infants detect continuity violations?

Renée Baillargeon, Jie Li, Yuyan Luo,
and Su-hua Wang

Abstract

According to Spelke (1994), infants interpret physical events in accord with a core principle of *continuity*, which states that objects exist and move continuously in time and space. Here we adopt a stronger definition of the principle, which states that objects not only exist and move continuously in time and space, but also retain their physical properties as they do so. We then present a new account of infants' physical reasoning that specifies under what conditions infants succeed and fail in detecting violations of the principle. Finally, we describe new lines of research that test specific predictions from the account.

7.1 Introduction

For the greater part of the 20th century, researchers generally assumed that young infants understand very little about the physical world. Two related factors contributed to this assumption. First, the leading theoretical accounts of the time tended to portray young infants as limited sensorimotor processors incapable of representation or thought (e.g. Bruner 1964, 1968; Piaget 1952, 1954). Second, a dearth of methodological tools forced investigators to rely primarily on infants' manual actions to assess their physical knowledge. For example, Piaget (1954) observed that young infants typically do not search for objects hidden behind or beneath other objects, and concluded that they do not yet realize that objects continue to exist when hidden.

The situation today is markedly different. New methods have brought to light new findings which indicate that even very young infants possess expectations about physical events (e.g. Baillargeon 1987; Goubet and Clifton 1998; Gredebäck and von Hofsten 2004; Hespos and Baillargeon in press; Hofstader and Reznick 1996; Hood and

Willatts 1986; Kaufman *et al.* in press; Lécuyer and Durand 1998; Leslie 1984; Newcombe *et al.* 1999; Spelke and Kestenbaum 1986; Wilcox *et al.* 1996). As a result of these empirical advances, there is now widespread (though by no means universal) agreement that physical reasoning constitutes one of the fundamental domains of human cognition, and that core principles within the domain facilitate infants' reasoning and learning about events (e.g. Baillargeon 2002; Carey and Spelke 1994; Gelman 1990; Gopnik and Wellman 1994; Keil 1991; Leslie 1995; Wellman and Gelman 1992).

This new theoretical perspective has given rise to many new research questions. In particular, what specific core principles are infants endowed with? And how do these principles operate? In this chapter, we focus on the principle of *continuity*. The original definition of the principle, as proposed by Spelke and her colleagues (e.g. Carey and Spelke 1994; Spelke 1994; Spelke *et al.* 1992, 1995), was that objects exist and move continuously in time and space. For reasons that will become clear, here we adopt a stronger definition of the principle, which states that objects not only exist and move continuously in time and space, but also retain their physical properties (e.g. their size, shape, pattern, and color) as they do so.

How does the principle of continuity operate? Under what conditions do infants succeed in detecting continuity violations, and under what conditions do they fail? Over the past few years, we have been developing an account of infants' physical reasoning that attempts to answer this question (e.g. Baillargeon 2002, 2004; Luo and Baillargeon 2005 c; Wang *et al.* 2005). This chapter is organized into three main sections: in the first, we review recent findings on the development of infants' physical knowledge; in the second, we present our account of infants' physical reasoning; finally, in the third, we introduce new lines of research which test specific predictions from the account.

7.2 How do infants acquire their physical knowledge?

Research over the past 15 years has shed considerable light on the development of infants' knowledge about physical events (for recent reviews, see Baillargeon 2002, 2004). Much of this research has used the violation-of-expectation (VOE) method. In a typical experiment, infants see two test events: an *expected* event, which is consistent with the expectation examined in the experiment, and an *unexpected* event, which violates this expectation. With appropriate controls, evidence that infants look reliably longer at the unexpected than at the expected event is taken to indicate that infants: (1) possess the expectation under investigation; (2) detect the violation in the unexpected event; and (3) are 'surprised' by this violation. The term 'surprised' is used here simply as a short-hand descriptor, to denote a state of heightened interest or attention induced by an expectation violation.¹

¹ VOE reports that young infants possess rich cognitive abilities, such as the ability to represent hidden objects, have recently been criticized (e.g. Bogartz *et al.* 1997; Cashon and Cohen 2000; Haith and Benson 1998; Munakata *et al.* 1997; Rivera *et al.* 1999; Roder *et al.* 2000; Schilling 2000; Thelen and Smith 1994).

7.2.1 Event categories, vectors, and variables

Recent research suggests that infants ‘sort’ events into distinct categories. Many of these *event categories* capture relatively simple spatial relations between objects, such as ‘object behind nearer object, or occluder’ (occlusion events), ‘object inside container’ (containment events), and ‘object under cover’ (covering events) (e.g. Aguiar and Baillargeon 2003; Casasola *et al.* 2003; Hespos and Baillargeon 2001a; Luo and Baillargeon 2005b; McDonough *et al.* 2003; Munakata 1997; Wang *et al.* 2005; Wilcox and Chapa 2002; for a review, see Baillargeon and Wang 2002).

Each event category comprises one or more *vectors*, which correspond to separate problems that infants must solve in order to fully predict outcomes within the category. For example, in the case of occlusion events, infants must learn to predict whether an object will be hidden or visible when behind an occluder, and also how soon an object that moves behind an occluder will emerge from behind it; in the case of containment events, infants must learn to predict whether an object can be lowered inside a container, how much of an object inside a container will protrude above it, and whether the portion of an object that lies inside a container should be hidden or visible (e.g. Aguiar and Baillargeon 1998; Arterberry 1997; Gredebäck and von Hofsten 2004; Hespos and Baillargeon 2001b; Luo and Baillargeon 2005c; Sitskoorn and Smitsman 1995; Wang *et al.* 2004; Wilcox and Schweinle 2003).

For each vector in an event category, infants identify a sequence of *variables* that enables them to predict outcomes within the vector more and more accurately over time (e.g. Aguiar and Baillargeon 1999; Baillargeon and DeVos 1991; Hespos and Baillargeon 2001a; Luo and Baillargeon 2005a; Sitskoorn and Smitsman 1995; Wang *et al.* 2005; Wilcox 1999). Variables are akin to condition–outcome rules: for a set of contrastive outcomes, a variable specifies what condition produces each outcome. Each variable that is added along a vector revises and refines predictions from earlier variables. This process can be illustrated by a simple decision tree (for related ideas, see Mitchell 1997; Quinlan 1993; Siegler 1978). As an example, the decision tree in Fig. 7.1 depicts some of the variables infants identify as they learn when objects behind occluders should and should not be hidden.

At about 2.5 months of age, infants use only a simple *behind/not-behind* variable to predict when objects behind occluders should be hidden: they expect an object to be

Investigators have argued that VOE findings are often open to alternative, low-level interpretations, which cast doubt on claims that young infants possess rich cognitive abilities. However, converging evidence for VOE findings is steadily accumulating; at present, this evidence comes from action tasks (e.g. Goubet and Clifton 1998; Gredebäck and von Hofsten 2004; Hespos and Baillargeon *in press*; Hofstader and Reznick 1996; Hood and Willatts 1986), from habituation tasks (e.g. Casasola *et al.* 2003; McDonough *et al.* 2003), and from tasks tapping neural correlates (Kaufman *et al.* *in press*). In addition, experimental tests of specific alternative interpretations of VOE findings have not supported these interpretations (e.g. Luo and Baillargeon 2005c; Luo *et al.* 2003, 2005; Wang *et al.* 2004; for review and discussion, see Aslin 2000; Baillargeon 1999, 2000, 2004; Lécuyer 2001; Munakata 2000; Wang *et al.* 2004).

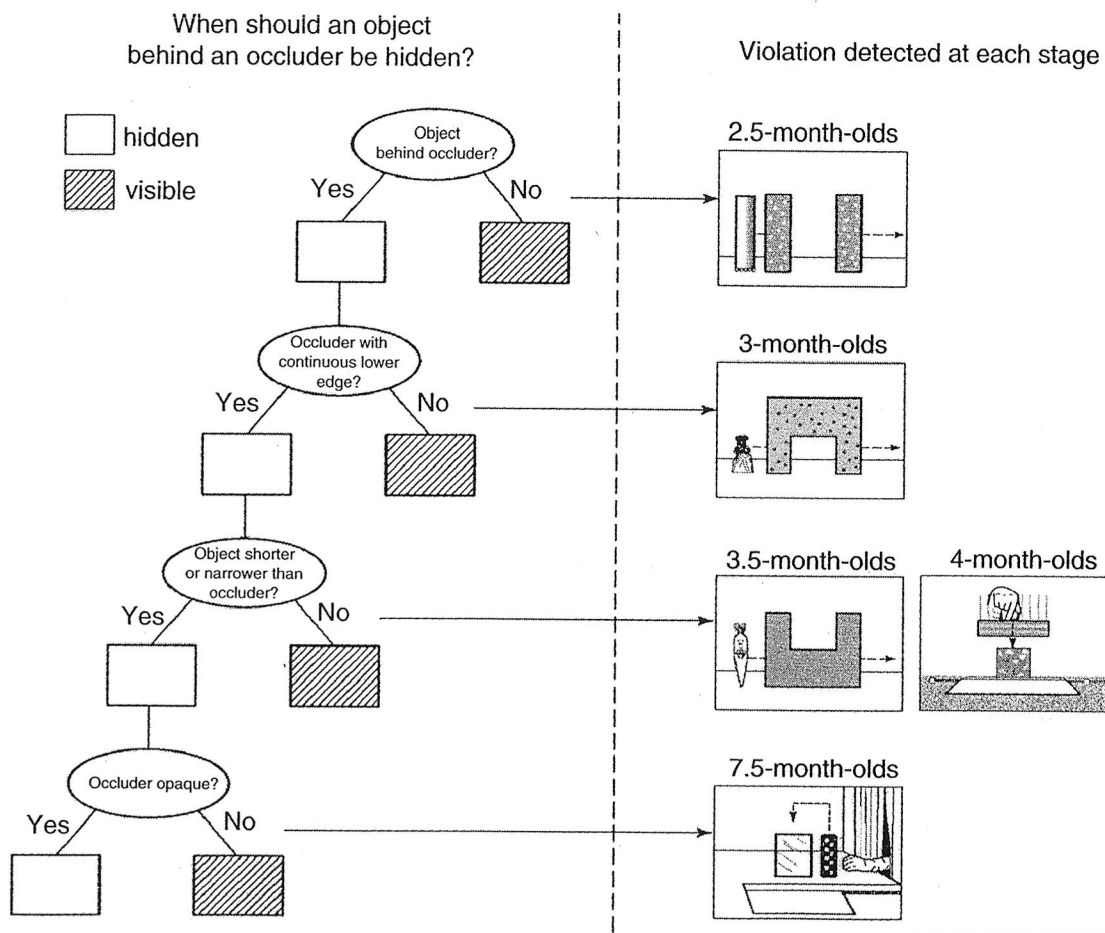


Fig. 7.1 Decision tree representing some of the variables infants identify to predict when objects behind occluders should be hidden or visible. The identification of each new variable enables infants to detect additional occlusion violations.

hidden when behind a closed occluder and to be visible when not (Aguiar and Baillargeon 1999; Lécuyer and Durand 1998; Luo and Baillargeon 2005c). Thus, when infants see an object move back and forth behind two screens placed a short distance apart, they expect the object to be hidden when behind each screen and to be visible when between them, because at that point the object does not lie behind any occluder. At about 3 months of age, infants identify a new occlusion variable, *lower-edge-discontinuity*: they now expect an object to be visible when behind a closed occluder whose lower edge is not continuous with the surface on which it rests, creating a gap between the occluder and surface (Aguiar and Baillargeon 2002; Luo and Baillargeon 2005c). Thus, infants expect an object to remain hidden when passing behind a screen shaped like a U, but not one shaped like an inverted-U. At about 3.5 to 4 months of age, infants identify *height* and *width* as occlusion variables: they now expect tall objects to remain partly visible when behind short occluders (Baillargeon and DeVos 1991), and wide objects to remain partly visible when behind narrow occluders (Wang *et al.* 2004; Wilcox 1999; Wilcox and Baillargeon 1998b).

Finally, at about 7.5 months of age, infants identify *transparency* as an occlusion variable: when an object is placed behind a transparent occluder, infants now expect the object to be visible through the front of the occluder, and are surprised if it is not (Luo and Baillargeon 2005a, 2005b).

7.2.2 Errors of omission and commission

We have just seen that, for each event category, infants identify variables – organized along vectors – which enable them to predict outcomes within the category more and more accurately over time. This description predicts that infants who have not yet identified a variable along a vector should err in two distinct ways in VOE tasks, when shown violation and non-violation events involving the variable. First, infants should respond to violation events consistent with their faulty knowledge as though they were expected. We refer to this first kind of error – viewing a violation event as expected – as an error of *omission*. Second, infants should respond to non-violation events inconsistent with their faulty knowledge as though they were unexpected. In other words, infants should respond to perfectly ordinary and commonplace events with increased attention, when these events happen to contradict their incomplete knowledge. We refer to this second kind of error – viewing a non-violation event as unexpected – as an error of *commission*.

To date, many investigations of infants' physical knowledge have revealed errors of omission: infants who have not yet identified a variable in a vector of an event category typically do not view violation events involving the variable as unexpected (e.g. Aguiar and Baillargeon 2002; Baillargeon and DeVos 1991; Hespos and Baillargeon 2001a; Luo and Baillargeon 2005c; Wang *et al.* 2005; Wilcox 1999). For example, 3-month-old infants, who have not yet identified height as an occlusion variable, view as expected a violation event in which a tall object remains fully hidden when passing behind a short occluder (Aguiar and Baillargeon 2002; Baillargeon and DeVos 1991; Luo and Baillargeon 2005c). At this age, infants have acquired only the variable lower-edge-discontinuity: as long as the lower edge of the occluder is continuous with the surface on which it rests, infants expect the object to remain hidden when behind the occluder, regardless of the heights of the object and occluder.

Recent experiments have also revealed errors of commission in infants' responses to occlusion events. In particular, there is now evidence that 2.5-month-olds, who have not yet identified the variable lower-edge-discontinuity, view as unexpected a non-violation event in which an object becomes visible when passing behind an inverted-U-shaped screen (Luo and Baillargeon 2005c). Similarly, 3-month-olds, who, as mentioned above, have not yet acquired the variable height, view as unexpected a non-violation event in which a tall object remains visible above a short occluder (Luo and Baillargeon 2005c). Finally, 7-month-olds, who have not yet identified transparency as an occlusion variable, view as unexpected a non-violation event in which

an object placed behind a transparent occluder remains visible through the occluder (Luo and Baillargeon 2005a).

7.2.3 Event-specific acquisitions

We have seen that, for each event category, infants identify variables, ordered along vectors, which allow them to better predict outcomes within the category. Recent experiments suggest that this acquisition process is event-specific: infants learn separately about each event category. This specificity manifests itself in at least two ways, described below.

First, variables identified in one event category remain tied to that category – they are not generalized to other categories, even when equally relevant. For example we saw above that infants identify the variable height at about 3.5 months in occlusion events (Baillargeon and DeVos 1991): they are now surprised when a tall object becomes fully hidden behind a short occluder. However, infants this age are not surprised when a tall object becomes fully hidden inside a short container, under a short cover, or inside a short tube. The variable height is not identified until about 7.5 months in containment events (Hespos and Baillargeon 2001a), until about 12 months in covering events (Wang *et al.* 2005), and until about 14 months in tube events (Wang *et al.* 2005). Similarly, we saw above that the variable transparency is identified at about 7.5 months in occlusion events (Luo and Baillargeon 2005a): infants are now surprised when an object placed behind a transparent occluder is not visible through the occluder. However, it is not until infants are about 9.5 months of age that they identify the same variable in containment events, and expect an object placed inside a transparent container to be visible through the container (Luo and Baillargeon 2005b). We use the Piagetian term *décalages* to refer to the lags in infants' identification of the same variable in different event categories.

Second, the same variable may be associated (at least initially) with different vectors in different event categories. This conclusion is suggested by different error patterns in infants' responses to occlusion and containment events. In occlusion events, as we just saw, height is identified at about 3.5 months and transparency at about 7.5 months. Both variables belong to a single vector having to do with when objects behind occluders should be hidden (see Fig. 7.1). Thus, 7-month-old infants, who have identified height but not transparency: (1) are surprised when a tall object becomes fully hidden behind a short occluder (a correct response; Baillargeon and Graber 1987; Hespos and Baillargeon 2005); and (2) are also surprised when an object placed behind a transparent occluder is visible through the occluder (an error of commission; Luo and Baillargeon 2005b). At this age, infants expect an object to be hidden when behind an occluder that is taller than the object – even if this occluder is in fact transparent.

In containment events, as we saw in the last section, height is identified at about 7.5 months and transparency at about 9.5 months. If these variables belonged to a

single vector specifying when objects inside containers should be hidden, then we should expect that at 8.5 months infants would produce responses similar to those described above for occlusion events. However, this is not the case. Although 8.5-month-old infants are surprised when a tall object becomes fully hidden inside a short container (a correct response; Hespos and Baillargeon 2001a, in press; Wang *et al.* 2005), they are not surprised when an object placed inside a transparent container is either visible or not visible through the container (an error of omission; Luo and Baillargeon 2005b). These results suggest that, in containment events, height and transparency belong to separate vectors: whereas height belongs to a vector specifying when an object inside a container should protrude above it, transparency belongs to a vector having to do with when an object inside a container should be hidden. Thus, when a short object is lowered inside a tall transparent container, 8.5-month-old infants bring to bear their knowledge of height to predict that no portion of the object will be visible above the container. However, they cannot make a prediction as to whether the portion of the object inside the container should be hidden or visible. Apparently, it is not until infants are about 9.5 months that they form a vector specifying when objects inside containers should be hidden.

This analysis leads to striking predictions concerning 7.5- and 8.5-month-old infants' responses to events involving transparent containers. When a tall object is lowered inside a short transparent container, infants should look reliably longer if the top of the object is not visible, as opposed to visible, above the container. However, as long as the top of the object protrudes above the container, infants should look about equally whether the bottom of the object is visible, or not visible, through the container. Experiments are planned to test these predictions.

7.2.4 Identifying variables

We have seen that, for each event category, infants identify variables which enable them to predict outcomes within the category more and more accurately over time. How do infants identify these variables? And why do they sometimes identify the same variable at different ages in different categories?

We have proposed that the process by which infants typically identify a new variable in an event category is one of explanation-based learning (EBL) and involves three main steps (Baillargeon 2002; Wang and Baillargeon 2005b; for a computational description of EBL in the machine learning literature, see DeJong 1993). First, infants must notice *contrastive outcomes* relevant to the variable (e.g. in the case of the variable height in covering events, infants must notice that when a cover is placed over an object, the object becomes sometimes fully and sometimes only partly hidden); because these contrastive outcomes are not predicted by infants' current physical knowledge, they serve to trigger learning. Second, infants must discover the *conditions* that map onto the outcomes (e.g. they must discover that an object becomes fully hidden when placed under a cover as tall as or taller than the object,

and becomes partly hidden otherwise). Third, infants must build an *explanation* for these condition–outcome data using their prior knowledge, which includes their core knowledge (e.g. because of their continuity principle, infants would readily understand that a tall object can extend to its full height inside a tall but not a short cover). Thus, according to the EBL account, only condition–outcome observations for which infants can build causal explanations are identified as new variables. These explanations are undoubtedly shallow (e.g. Keil 1995; Wilson and Keil 2000), but they still serve to integrate new variables with infants' prior causal knowledge.

The EBL account suggests at least two reasons why infants may identify a variable in one event category several weeks or months before they identify it in another event category. One reason has to do with the first step in the EBL process: because exposure to appropriate contrastive outcomes is necessary to trigger learning, it follows that variables will be learned later when exposure is less frequent. Thus, infants may identify height as a containment variable several months before they identify it as a covering variable (Hespos and Baillargeon 2001a; Wang *et al.* 2005) simply because, in everyday life, infants have more opportunities to notice that objects placed inside containers sometimes extend above them and sometimes not, than to notice that objects placed under covers sometimes extend beneath them and sometimes not.

A second reason why infants may identify a variable sooner in one event category than in another has to do with the second step in the EBL process. After noticing the contrastive outcomes for a variable, infants must discover the conditions that map onto these outcomes; this discovery may be more difficult in some categories than in others. To illustrate, consider the finding that infants identify height as an occlusion variable several months before they identify it as a containment variable (Baillargeon and DeVos 1991; Hespos and Baillargeon 2001a). Prior research (e.g. Baillargeon 1994, 1995) indicates that when infants begin to reason about a continuous variable in an event category, they can reason about the variable qualitatively but not quantitatively: they are not able at first to encode and reason about absolute amounts. In order to encode the heights of objects and occluders or containers qualitatively, infants must compare them as they stand *side by side*. It may be that infants have more opportunities to perform such qualitative comparisons with occlusion than with containment events. In the case of occlusion events, infants will often see objects move behind the side edges of occluders, making it easy to compare their heights as they stand next to each other (e.g. when a cereal box is pushed in front of a bowl). In the case of containment events, however, there may be relatively few instances in which objects are placed first next to and then inside containers; caretakers will more often lower objects directly into containers, giving infants no opportunity to compare their heights (e.g. Hespos and Baillargeon 2001a; Wang *et al.* 2004).

The preceding analysis predicts that infants who are exposed in the laboratory to appropriate outcome and condition data for a variable should identify it earlier than they would otherwise. To test this prediction, we recently attempted to 'teach'

9-month-old infants the variable height in covering events (Wang and Baillargeon 2005b); recall that this variable is typically not identified until about 12 months of age (Wang *et al.* 2005). Our results were positive and as such support both the EBL process and the speculation above that the *décalage* in infants' identification of the variable height in containment and covering events stems from the fact that infants are typically exposed to appropriate observations for this variable at different ages in the two categories.

7.3 An account of infants' physical reasoning

Armed with the findings presented in the last section, we now return to the question raised in the Introduction: How does infants' principle of continuity operate? Over the past few years, we have been developing an account of infants' physical reasoning that attempts to answer this question (e.g. Baillargeon 2002, 2004; Luo and Baillargeon 2005c; Wang *et al.* 2005).

7.3.1 Four assumptions

Our reasoning account rests on four assumptions. First, when watching a physical event, infants build a specialized physical representation of the event which is used to predict and interpret its outcome. Second, all of the information, but only the information, in infants' physical representation of an event becomes subject to a few core principles, including that of continuity (e.g. Leslie 1994; Spelke 1994; Wang and Baillargeon 2005b).

Third, in the first weeks of life, infants' physical representation of an event typically includes only basic spatial and temporal information about the event (Fig. 7.2a; e.g. Kestenbaum *et al.* 1987; Leslie 1994; Needham 2000; Slater 1995; Spelke 1982; Yonas and Granrud 1984). This *basic information* specifies primarily: (1) how many distinct objects are involved in the event (e.g. are there two objects present?); (2) what is the geometry, or distribution of open/closed surfaces, of each object (e.g. is one object open at the top to form a container, open at the bottom to form a cover, or open at both ends to form a tube?); and (3) what is the spatial arrangement of the objects and how does it change over time as the objects move or are moved (e.g. is one object being placed behind, inside, or under the other object?). The basic information thus captures essential aspects of the event, but leaves out most of its details: for example it includes no information about the relative sizes of the objects (e.g. is one object taller or wider than the other object?), or about their surface appearance (e.g. are the objects transparent or opaque?).

Fourth, as they form event categories and identify variables for each category, infants include more and more of this detailed information, or *variable information*, in their physical representations (Fig. 7.2b). When watching an event, infants first represent the basic information about the event, and use this information to

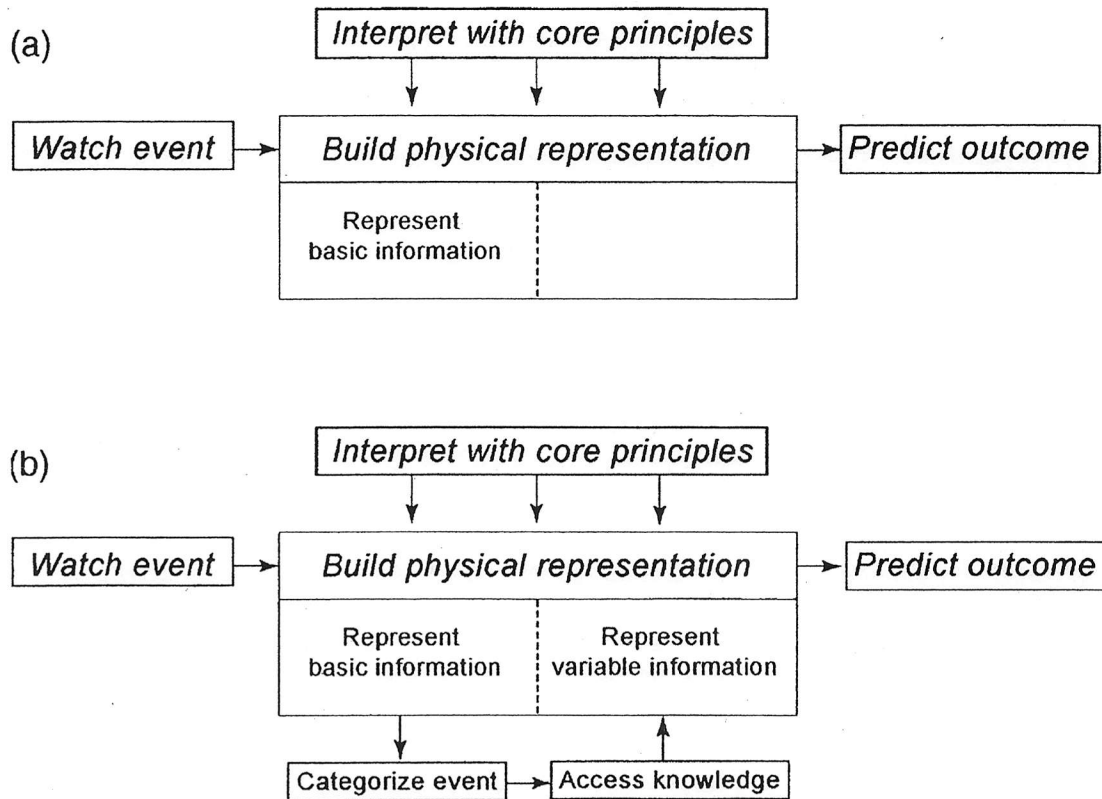


Fig. 7.2 Schematic presentation of the reasoning account for younger (a) and older (b) infants.

categorize it. Infants then access their knowledge of the event category selected. This knowledge specifies the variables that have been identified as relevant for predicting outcomes in the category, and hence that should be included in the physical representation of the event. Variables not yet identified are typically not included in the representation.

To illustrate our reasoning account, consider the finding that infants aged 7.5 months and older are surprised when a tall object is lowered inside a short container until it becomes hidden (Hespos and Baillargeon 2001a; Wang *et al.* 2005). The account suggests that, when watching this event, infants represent the basic information about the event and interpret this information in accord with their continuity principle ('object being lowered inside container'). Next, infants categorize the event as a containment event, and access their knowledge of this event category. Because at 7.5 months this knowledge includes the variable height, infants include information about the relative heights of the object and container in their physical representation of the event. This variable information then becomes subject to infants' continuity principle, making it possible for them to detect the continuity violation in the event: they recognize that the object is too tall to become hidden inside the short container. Infants younger than 7.5 months, who have not yet identified height as a containment variable, typically do not include height information in their physical

representations of containment events. As a result, this information is not available and hence cannot be interpreted in accord with infants' continuity principle. Infants thus fail to detect continuity violations involving tall objects and short containers (Hespos and Baillargeon 2001a).

7.3.2 Basic continuity violations

According to our reasoning account, young infants should succeed in detecting any continuity violation that involves only the basic information they typically include in their physical representations of events. We refer to such violations as *basic continuity violations*.

There are now several reports indicating that 2.5- to 3-month-old infants (the youngest tested successfully to date with the VOE method) can detect basic continuity violations in occlusion, containment, and covering events (e.g. Aguiar and Baillargeon 1999; Hespos and Baillargeon 2001b; Luo and Baillargeon 2005c; Spelke *et al.* 1992; Wang *et al.* 2005; Wilcox *et al.* 1996). For example there is evidence that these young infants are surprised: (1) when an object disappears behind one occluder and reappears from behind another occluder without appearing in the gap between them (Aguiar and Baillargeon 1999; Luo and Baillargeon 2005c; Wilcox *et al.* 1996); (2) when an object is lowered inside a container through its closed top (Hespos and Baillargeon 2001b); (3) when an object is lowered inside an open container, which is then slid forward and to the side to reveal the object standing in the container's initial position (Hespos and Baillargeon 2001b); (4) when a cover is lowered over an object, slid to the side, and lifted to reveal no object (Wang *et al.* 2005); and (5) when a cover is lowered over an object, slid behind the left half of a screen taller than the object, lifted above the screen, moved to the right, lowered behind the right half of the screen, slid past the screen, and finally lifted to reveal the object (Wang *et al.* 2005).

To succeed in detecting the continuity violations in these events, infants need not represent any variable information, only basic information: in each event, they must specify how many objects are present, which of their surfaces are open/closed, how the objects are spatially arranged, and how this arrangement changes over time. This basic information, as it is represented, becomes subject to the continuity principle. When the event evolves in a manner inconsistent with the principle, it is tagged as a violation, causing infants to respond with increased attention.

To illustrate, consider the finding that infants are surprised when a cover is lowered over an object, slid to the side, and then lifted to reveal no object (Wang *et al.* 2005). We would argue that infants represent the following basic information: (1) an object open at the bottom, or cover, is held over a closed object; (2) the cover is lowered over the object [the continuity principle would specify at this point that the object continues to exist, in its same location, under the cover]; (3) the cover is slid to the side [the continuity principle would specify at this point that the object cannot pass through the sides of the cover and hence must be displaced with the cover to its new

location]; and (4) the cover is lifted to reveal no object [the continuity principle would signal at this point that a violation has occurred: the object should have been revealed when the cover was lifted].

7.3.3 Variable continuity violations

According to our reasoning account, infants should fail to detect any continuity violation that involves a variable they have not yet identified as relevant to an event category, and hence do not yet include in their physical representations of events from the category. We refer to such violations as *variable continuity violations*.

When infants identify a variable earlier in one event category than in another, striking discrepancies can arise in their responses to perceptually similar events from the two categories: infants will detect a continuity violation involving the variable in one category, but not in the other. There are now several reports of such discrepancies (e.g. Hespos and Baillargeon 2001a; Luo and Baillargeon 2005b; Wang *et al.* 2005). For example, 4.5-month-old infants watched an experimenter lower a tall cylindrical object either behind (occlusion condition) or inside (containment condition) a container until only the knob at the top of the object remained visible (Hespos and Baillargeon 2001a). In one test event, the container was as tall as the cylindrical portion of the object (tall event); in the other test event, the container was only half as tall (short event), so that it should have been impossible for the cylindrical portion of the object to become fully hidden inside the container. The infants in the occlusion condition looked reliably longer at the short than at the tall event, but those in the containment condition tended to look equally at the two events. The infants thus detected the variable continuity violation in the occlusion but not the containment condition.

According to our reasoning account, the infants in the occlusion condition represented the basic information about each test event, categorized it as an occlusion event, and then accessed their knowledge of this event category. Because at 4.5 months this knowledge comprises the variable height (recall that this variable is identified at about 3.5 months; Baillargeon and DeVos 1991), the infants included information about the relative heights of the object and container in their physical representation of the event. This information became subject to the continuity principle, and the short event was marked as a violation event. The infants in the containment condition went through a similar reasoning process; however, because at 4.5 months infants' knowledge of containment events does not yet comprise the variable height (recall that this variable is not identified until about 7.5 months; Hespos and Baillargeon 2001a), the infants did not include information about the relative heights of the object and container in their physical representation of each test event. As a result, this (missing) information could not be interpreted in accord with the continuity principle, and the infants failed to detect the variable continuity violation in the short event.

In the experiments just described, the infants responded differently when an object was lowered behind or inside a container. These events were perceptually similar, but not identical. In a recent experiment, 9-month-old infants responded differently to perceptually *identical* events, when prior information led them to categorize the events as containment or as tube events (Wang *et al.* 2005). The infants watched two test events in which an experimenter lowered a tall object inside a container (containment condition) or a tube (tube condition) until it became fully hidden. In one event, the container or tube was slightly taller than the object (tall event); in the other event, the container or tube was only half as tall (short event), so that it should have been impossible for the object to become fully hidden. Prior to the test session, in an orientation procedure, the experimenter showed the infants each container (containment condition) or tube (tube condition) one at a time, calling attention to its top and bottom. When standing upright on the apparatus floor during the test events, the containers and tubes were indistinguishable. The infants in the containment and tube condition thus saw perceptually identical test events; only the information provided in the orientation procedure could lead them to believe that they were watching events involving containers or tubes.

The infants in the containment condition looked reliably longer at the short than at the tall event, but those in the tube condition tended to look equally at the two events. Thus, the infants detected the variable continuity violation in the containment but not the tube condition. According to our reasoning account, the infants in the containment condition categorized each test event as a containment event (based in part on information remembered from the orientation procedure), and then accessed their knowledge of this event category. Because at 9 months the variable height is known to be relevant for predicting outcomes in containment events (recall that this variable is identified at about 7.5 months; Hespos and Baillargeon 2001a), the infants included information about the relative heights of the object and container in their physical representation of each test event. This information became subject to the continuity principle, and the short event was tagged as violating the principle. The infants in the tube condition underwent a similar reasoning process; however, because height is not identified as a tube variable until about 14 months (Wang *et al.* 2005), the infants included no information about the relative heights of the object and tube in their physical representation of each test event. As a result, this (missing) information could not become subject to the continuity principle, and the infants failed to detect the violation in the short event.

The infants in this last experiment thus detected the variable continuity violation they were shown when they believed that they were facing containers, but not tubes. Such a finding provides strong evidence for our reasoning account and more specifically for the claim that: (1) infants detect a variable continuity violation in an event when they include information about the variable in their physical representation of the event; and (2) infants include this information when they have identified the variable as relevant for predicting outcomes in the event's category.

7.4 Tests of our reasoning account

The reasoning account presented in the last section suggests several interesting predictions. Two such predictions are examined here; both focus on infants' ability to detect variable continuity violations.

7.4.1 Change blindness

According to our reasoning account, infants who have not yet identified a variable as relevant to an event category typically do not include information about this variable in their physical representations of events from the category. If infants do not include information about a variable when representing an event, then they should be unable to detect surreptitious changes involving the variable: in other words, they should be *blind* to such changes (for related findings in the adult perception literature, see Rensink 2002; Rensink *et al.* 1997; Simons 1996, 2000).

7.4.1.1 Occlusion events

There have been many experiments over the past 30 years examining infants' ability to detect a change in an object that is briefly hidden (e.g. Bower 1974; Bower *et al.* 1971; Goldberg 1976; Gratch 1982; Meicler and Gratch 1980; Muller and Aslin 1978; Newcombe *et al.* 1999; Simon *et al.* 1995; Wilcox 1999; Wilcox and Baillargeon 1998a, 1998b). In the case of VOE tasks involving occlusion events, two factors appear to determine whether infants will respond with increased attention to an event in which one object disappears behind an occluder and a different object reappears (after an appropriate interval) from behind it. The first factor has to do with the width of the occluder relative to that of the objects. As we saw earlier, by 4 months of age, infants have identified width as an occlusion variable: they realize that a narrow occluder cannot hide a wide object, nor can it hide two narrow objects whose combined width is greater than that of the occluder (e.g. Wang *et al.* 2004; Wilcox 1999; Wilcox and Baillargeon 1998b). Infants who detect an object change in an occlusion event typically respond with increased attention only if the screen is too narrow to hide the two objects at the same time. If the screen is wide enough to hide both objects, infants do not respond with increased attention, because they can readily make sense of the event: they infer that two different objects are present behind the screen (e.g. Wilcox 1999; Wilcox and Baillargeon 1998a, 1998b).

The second factor has to do with the variable information infants include about the objects. In an extensive series of experiments using narrow-screen events, Wilcox (1999) found that infants detect differences in size and shape at about 4 months, differences in pattern at about 7.5 months, and differences in color at about 11.5 months. To restate this last result in terms of change blindness, infants younger than 11.5 months are blind to color changes in narrow-screen events. When a green ball disappears behind a narrow screen and a red ball reappears from behind it, infants do not realize that two different balls are present. Because they have not yet identified

color as an occlusion variable, they do not include color information in their representation of the event. As a result, they assume that the event involves a single ball which becomes briefly occluded as it moves back and forth behind the screen.

7.4.1.2 Other events

We have recently begun to examine change blindness in event categories other than occlusion. Some experiments have focused on height changes in covering events (Wang and Baillargeon in press). Ongoing experiments (Li and Baillargeon 2005a) are focusing on height changes in containment and tube events. For example, in one experiment, 8-month-old infants first see a familiarization event in which an experimenter's gloved hand rotates a container (containment condition) or a tube (tube condition) forward and backward to show its top and bottom; the tube is identical to the container with its bottom removed. Next, all of the infants see the same change and no-change test event. At the start of the change event, the container/tube stands on the apparatus floor (because the container and tube are indistinguishable when standing upright, the infants can only assume that they are facing a container or tube based on the information provided in the familiarization event). Next to the container is a tall cylindrical object with a knob attached to its top; the container/tube is as tall as the cylindrical portion of the object. The experimenter's hand grasps the knob at the top of the object, lifts the object, and lowers it inside the container/tube until only the knob and very top of the object remain visible above the rim. The hand then gently twists the object back and forth for a few seconds. Finally, the hand lifts the object and returns it to its original position on the apparatus floor. When removed from the container/tube, the object is much shorter: its cylindrical portion is only half as tall as previously. The no-change event is identical to the change event, except that the short object is used throughout the event.

Because height is identified at about 7.5 months in containment events (Hespos and Baillargeon 2001a, 2005), we predict that the infants in the containment condition will include information about the relative heights of the object and container in their physical representation of each test event, and hence will detect the difference in the height of the object in the change event. As a result, the infants will look reliably longer at the change than at the no-change event. Conversely, because height is not identified until about 14 months in tube events (Wang *et al.* 2005), we predict that the infants in the tube condition will include no height information in their physical representations of the test events, will fail to detect the change in the height of the object, and hence will look about equally at the two test events.

Preliminary results support these predictions: although the infants in the two conditions view exactly the same test events, only the infants in the containment condition look reliably longer at the change than at the no-change event. Thus, at 8 months, infants appear to detect a change in the height of an object if this change takes place when the object is lowered inside a container – but not inside a tube.

7.4.1.3 A caveat

Before leaving this section, we offer an important caveat. When we say that infants do not include information about a variable in their physical representation of an event, we do not mean to claim that they do not represent this information *at all*. Whether they do or not is an empirical question we are currently investigating. Our working hypothesis is that when infants watch a physical event, different computational systems form different representations simultaneously, for distinct purposes. In particular, infants' *object-recognition* system represents detailed information about the objects in the event, for recognition and categorization purposes. At the same time, infants' *physical-reasoning* system forms a physical representation of the event, to predict and interpret its outcome. Thus, when infants determine that they must include information about a variable in their physical representation of an event, and this information is no longer perceptually available, they access their object-recognition system to retrieve the necessary information. On this view, the 8-month-old infants in the experiment just described (Li and Baillargeon 2005a) *all* encoded information about the relative heights of the object and container or tube in their object-recognition system. However, *only* the infants in the container condition retrieved this height information and included it in their physical representations of the test events.

These speculations suggest that even infants who fail to include variable information in their representation of an event could nevertheless have this information available in their object-recognition system. Tasks designed to tap this system directly should thus reveal this knowledge. This means, for example, that infants who have not yet identified the variable height in containment, covering, or tube events should nevertheless detect a change in the relative heights of an object and container, cover, or tube, when given a task that taps their object-recognition rather than their physical-reasoning system. Experiments are under way to test this prediction.

7.4.2 Inducing infants to detect variable continuity violations

We have argued that infants fail to detect a variable continuity violation in an event when: (1) they have not yet identified the variable as relevant to the event's category; (2) they do not include information about the variable in their physical representation of the event; and (3) the missing information cannot be interpreted in accord with the infants' continuity principle, so the event cannot be tagged as violating the principle. This analysis predicts that if infants could be *induced*, through various contextual manipulations, to include information about the variable in their physical representation of the event, then this information should become subject to the continuity principle, and infants should be able to detect the violation in the event. According to our reasoning account, the missing variable information, once represented, should be immediately interpretable by the continuity principle.

7.4.2.1 Priming effects

Recent evidence indicates that infants can be primed to include information about a variable they have not yet identified in their physical representations of events. Wilcox and Chapa (2004) built on the finding, described earlier, that infants younger than 11.5 months do not detect color changes in occlusion events: they are not surprised when a green ball and a red ball appear successively from behind a screen that is too narrow to hide them both (Wilcox 1999). Wilcox and Chapa set out to prime 7.5-month-old infants to attend to the color information in their narrow-screen event. Prior to the test trials, the infants received three pairs of priming trials. In each pair, the infants saw a pound event, in which a green cup was used to pound a peg, and a pour event, in which a red cup was used to pour salt. Green and red cups of different sizes and shapes were used in the three priming pairs. Next, the infants saw a test event in which a green and a red ball appeared successively from behind a narrow or a wide screen. The infants who saw the narrow-screen event looked reliably longer than those who saw the wide-screen event, suggesting that the priming trials had induced the infants to include color information in their physical representation of each test event. This priming effect was eliminated when the infants received only two priming pairs, or when the same cups were used across all three pairs.

Additional priming experiments built on the finding, reported earlier, that infants younger than 7.5 months do not detect pattern changes in occlusion events: they are not surprised when a dotted and a striped green ball appear successively from behind a screen that is too narrow to hide them both (Wilcox 1999). Using similar priming trials involving dotted and striped green cups, Wilcox and Chapa (2004) found that 5.5- and even 4.5-month-old infants could be primed to include pattern information in their physical representations of narrow- and wide-screen events.

These priming results provide strong support for the expanded definition of the continuity principle proposed in the Introduction – objects not only exist continuously in time and space, but also retain their physical properties as they do so. Young infants typically do not include color or pattern information in their representations of occlusion events. However, if primed to do so, infants immediately expect objects to retain these properties throughout the events. Thus, infants who are primed to represent an object as green and dotted expect it to remain green and dotted when passing behind an occluder. Note that this expectation is not a low-level response to a perceptual change: following priming, infants respond with increased attention when a green ball disappears behind an occluder and a red ball reappears from behind it *only* if the occluder is too narrow to hide the two balls at the same time. In other words, the priming experience does not merely heighten infants' sensitivity to color and pattern changes: rather, it leads them to include color and pattern information in their physical representations of events. This information, once included, becomes subject to infants' continuity principle, and events in which objects appear to spontaneously change color or pattern are tagged as violating the principle.

7.4.2.2 Carry-over effects

We have been developing a very different approach for inducing infants to detect variable continuity violations. The point of departure for this approach was the following question: What happens when infants see the same objects in two successive events from different event categories? Do they represent each event separately? Or do they carry over whatever variable information they included in their representation of the first event to their representation of the second event? The second alternative seemed to us more efficient and hence more plausible (e.g. Aguiar and Baillargeon 2003).

We reasoned that if infants carry over variable information from one event representation to the next, then infants who see an event in which a variable has been identified, followed by an event in which this same variable has not yet been identified, should show a *positive* carry-over effect: the variable information included in the first event representation should be carried over to the second event representation, allowing infants to detect continuity violations involving the variable earlier than they would otherwise. Exposure to a single initial event would thus be sufficient to induce infants to detect a variable continuity violation in a subsequent event: as long as infants spontaneously include the appropriate variable information in their representation of the first event, this information should be available to them when reasoning about the second event (e.g. Wang and Baillargeon 2005b).

At the same time, we realized that the converse should also be true: if variable information is carried over from one event representation to the next, then infants who see an event in which a variable has not yet been identified, followed by an event in which this same variable has been identified, should show a *negative* carry-over effect: the information about the variable should be absent from the second event representation, causing infants to fail to detect continuity violations they would otherwise have been able to detect.

Do infants show negative as well as positive carry-over effects when they see the same objects in two successive events from different categories? An ongoing experiment (Li and Baillargeon 2005b) addresses this question. This experiment examines 8.5-month-old infants' ability to detect a surreptitious change in the height of an object in an event sequence comprising an occlusion and a covering event.

The infants are assigned to an occlusion–covering or a covering–occlusion condition. The infants in the occlusion–covering condition receive a change or a no-change test trial. At the beginning of the change trial, a short cylindrical object stands next to a tall rectangular cover with a knob attached to its top; the object is half as tall as the rectangular portion of the cover. To start, an experimenter's gloved hand grasps the knob at the top of the cover, rotates the cover forward to show its hollow interior, and then replaces the cover next to the object (pretrial). Next, the hand slides the cover in front of the object, fully hiding it, and then returns it to its original position on the apparatus floor (occlusion event). Finally, the hand lowers the cover over the object, again fully hiding it, and then returns it to the apparatus floor (covering event). When

the cover is removed from over the object in the covering event, the object is now as tall as the rectangular portion of the cover. In the no-change trial, the tall object is used throughout the trial. The infants in the covering-occlusion condition receive similar change and no-change trials, except that the occlusion and covering events are performed in the reverse order: the cover is placed first over and then in front of the object. The surreptitious change in the height of the object in the change trial thus takes place in the occlusion rather than in the covering event.

Because the variable height is identified at about 3.5 months in occlusion events (Baillargeon and DeVos 1991), but only at about 12 months in covering events (McCall 2001; Wang *et al.* 2005), we expect that the infants in the occlusion–covering condition will show a positive carry-over effect. When watching the occlusion event, the infants will categorize the event, access their knowledge of occlusion events, and include information about the relative heights of the cover and object in their physical representation of the event. When the infants next see the covering event, this height information will be carried over into this new representation; the information will then be interpreted in terms of the continuity principle, allowing the infants to detect the violation in the change event.

In contrast, the infants in the covering–occlusion condition should show a negative carryover effect. When watching the covering event, the infants will include no height information in their representation of the event. As a result, no height information will be carried over when the infants next represent the occlusion event. The infants will thus fail to detect the continuity violation in the change event.

Preliminary results support our predictions: in the occlusion–covering condition, the infants who see the change trial look reliably longer than those who see the no-change trial; in the covering–occlusion condition, the infants look about equally during the two trials.

These results are interesting for several reasons. First, they provide strong support for the notion that infants detect variable continuity violations when they include information about the relevant variables in their physical representations of the events. Second, they provide additional support for the expanded principle of continuity introduced in this chapter. Wilcox and Chapa (2004) found that young infants who were primed to encode the color or pattern of an object in an occlusion event were surprised when the object changed color or pattern behind the occluder. In a similar vein, the infants in the occlusion–covering condition who were induced to include height information in their representation of the covering event were surprised when the object changed height under the cover (see also Wang and Baillargeon 2005b). Together, these results make clear that infants who are led by contextual manipulations to represent the color, pattern, or height of an object in an event immediately expect these properties to remain stable – they do not need to learn that green balls cannot turn into red balls, dotted balls into striped balls, or short cylinders into tall ones (cf. Scholl and Leslie 1999).

Third, the present results suggest that when infants see objects involved in a sequence of two events from different event categories, they carry over whatever variable information they included in their physical representation of the first event to that of the second event. In some cases, this carry-over can *induce* infants to detect a variable continuity violation they would otherwise have failed to detect: the infants in the occlusion–covering condition detected at 8.5 months a violation that is typically not detected until about 12 months (Wang and Baillargeon 2005b; Wang *et al.* 2005). In other cases, the carry-over of variable information can *prevent* infants from detecting a variable continuity violation they would otherwise have been able to detect: the infants in the covering–occlusion condition failed to detect at 8.5 months a violation that is typically detected at 3.5 months (Baillargeon and DeVos 1991).

What mechanism might underlie these positive and negative carry-over effects between physical representations? One possibility is suggested by a recent model of object-based attention in infants (e.g. Kaldy and Leslie 2003; Leslie *et al.* 1998; Scholl and Leslie 1999; for related models of visual attention in adults, see Kahneman *et al.* 1992; Pylyshyn 1989, 1994). According to this model, when infants attend to an event involving a few objects, they assign an index to each object. These indexes serve as pointers that help keep track of the objects as the event unfolds (each index ‘sticks to’ its object as it moves). Typically, indexes are assigned based on spatiotemporal information and contain no featural information; however, such information can be added through a binding process. As the objects engage in first one and then another event, the same indexes continue to be used as long as the infants keep attending to the objects. Finally, in any event, the maximum number of indexes that can be assigned concurrently is three or four (four is the limit in adults).

This model suggests a simple explanation for the carry-over effects reported here. First, consider the infants in the occlusion–covering condition. When representing the occlusion event, the infants assigned an index to the cover and to the object, and bound height information to these indexes; when the infants next saw the covering event, which involved the same cover and object, they continued to use the same indexes, so that the height information bound to these indexes became, fortuitously, available to them. Next, consider the infants in the covering–occlusion condition. When representing the covering event, the infants again assigned an index to the cover and object – but bound no height information to these indexes. Because the occlusion event involved the same cover and object, the infants continued to use the same indexes and thus failed to detect the height continuity violation in the change event.

Blending together our reasoning account and Leslie’s model of object-based attention (e.g. Kaldy and Leslie 2003; Leslie *et al.* 1998; Scholl and Leslie 1999) may thus provide useful insights into infants’ representations of single as well as multiple events. In particular, such a hybrid account may help explain what variable information is bound to indexes in any one event representation, what variable information is carried over from one event representation to the next, and more generally what

variable continuity violations infants succeed or fail to detect in the context of single and multiple events.

7.5 Concluding remarks

The research reviewed in this chapter indicates that infants detect continuity violations when they include in their physical representations the basic and variable information necessary to detect these violations. Violations that involve only basic information are typically detected at an early age, because even very young infants generally include adequate basic information in their physical representations (for exceptions, see Baillargeon 1987; Baillargeon and DeVos 1991). In contrast, violations that involve variable information are typically detected at later ages, because infants who have not yet identified a variable as relevant to an event category typically do not include information about this variable when representing events from the category. An infant who does not represent the heights of a tall object and short container cannot be surprised when the object becomes fully hidden inside the container.

Infants who have not yet identified a variable as relevant to an event category can nevertheless be induced to include information about this variable in their physical representations of events from the category, through appropriate contextual manipulations. For example infants can be induced to include information about the color or pattern of an object in an occlusion event (Wilcox and Chapa 2004), or about the height of an object in a covering event (Li and Baillargeon 2005b; Wang and Baillargeon 2005b). This variable information, once included in the physical representation, becomes subject to the continuity principle (objects exist and move continuously through time and space, retaining their physical properties as they do so), allowing infants to detect any surreptitious change or other continuity violation involving the variable. The world of infants is thus not a fairy-tale one: objects that are represented (either spontaneously or as a result of contextual manipulations) as small, green, and frog-like, are expected *not* to spontaneously turn into objects that are large, blond, and prince-like.

Acknowledgment

The preparation of this chapter was supported by a grant from the National Institute of Child Health and Human Development to the first author (HD-21104).

References

- Aguiar A and Baillargeon R (1998). 8.5-month-old infants' reasoning about containment events. *Child Development*, 69, 636–653.
- Aguiar A and Baillargeon R (1999). 2.5-month-old infants' reasoning about when objects should and should not be occluded. *Cognitive Psychology*, 39, 116–157.

- Aguiar A and Baillargeon R (2002). Developments in young infants' reasoning about occluded objects. *Cognitive Psychology*, 45, 267–336.
- Aguiar A and Baillargeon R (2003). Perseverative responding in a violation-of-expectation task in 6.5-month-old infants. *Cognition*, 88, 277–316.
- Arterberry ME (1997). Perception of object properties over time. In: C Rovee-Collier and LP Lipsitt, eds. *Advances in infancy research*, Vol. 11, pp. 219–268. Greenwich, CT: Ablex.
- Aslin RN (2000). Why take the cog out of infant cognition? *Infancy*, 1, 463–470.
- Baillargeon R (1987). Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology*, 23, 655–664.
- Baillargeon R (1994). How do infants learn about the physical world? *Current Directions in Psychological Science*, 3, 133–140.
- Baillargeon R (1995). A model of physical reasoning in infancy. In: C Rovee-Collier and LP Lipsitt, eds. *Advances in infant research*, Vol. 9, pp. 305–371. Norwood, NJ: Ablex.
- Baillargeon R (1999). Young infants' expectations about hidden objects: A reply to three challenges (article with peer commentaries and response). *Developmental Science*, 2, 115–163.
- Baillargeon R (2000). Reply to Bogartz, Shinskey, and Schilling; Schilling; and Cashon and Cohen. *Infancy*, 1, 447–462.
- Baillargeon R (2002). The acquisition of physical knowledge in infancy: A summary in eight lessons. In: U Goswami, ed. *Handbook of childhood cognitive development*, pp. 47–83. Oxford: Blackwell.
- Baillargeon R (2004). Infants' reasoning about hidden objects: Evidence for event-general and event-specific expectations (article with peer commentaries and response). *Developmental Science*, 7, 391–424.
- Baillargeon R and DeVos J (1991). Object permanence in young infants: Further evidence. *Child Development*, 62, 1227–1246.
- Baillargeon R and Graber M (1987). Where's the rabbit? 5.5-month-old infants' representation of the height of a hidden object. *Cognitive Development*, 2, 375–392.
- Baillargeon R and Wang S (2002). Event categorization in infancy. *Trends in Cognitive Sciences*, 6, 85–93.
- Bogartz RS, Shinskey JL and Speaker CJ (1997). Interpreting infant looking: The event set x event set design. *Developmental Psychology*, 33, 408–422.
- Bower TGR (1974). *Development in infancy*. San Francisco: WH Freeman.
- Bower TGR, Broughton JM and Moore MK (1971). Development of the object concept as manifested in the tracking behavior of infants between seven and twenty weeks of age. *Journal of Experimental Child Psychology*, 11, 182–193.
- Bruner JS (1964). The course of cognitive growth. *American Psychologist*, 19, 1–15.
- Bruner JS (1968). *Processes of cognitive growth: Infancy*. Worcester, MA: Clark University Press and Barre Press.
- Carey S and Spelke ES (1994). Domain-specific knowledge and conceptual change. In: LA Hirschfeld and SA Gelman, eds. *Mapping the mind: Domain specificity in cognition and culture*, pp. 169–200. New York: Cambridge University Press.
- Casasola M, Cohen L and Chiarello E (2003). Six-month-old infants' categorization of containment spatial relations. *Child Development*, 74, 679–693.
- Cashon CH and Cohen LB (2000). Eight-month-old infants' perceptions of possible and impossible events. *Infancy*, 1, 429–446.
- DeJong GF (1993). *Investigating explanation-based learning*. Boston, MA: Kluwer Academic Press.

- Gelman R (1990). First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cognitive Science*, 14, 79–106.
- Goldberg S (1976). Visual tracking and existence constancy in 5-month-old infants. *Journal of Experimental Child Psychology*, 22, 478–491.
- Gopnik A and Wellman HM (1994). The theory theory. In: LA Hirschfeld and SA Gelman, eds. *Mapping the mind: Domain specificity in cognition and culture*, pp. 257–293. New York: Cambridge University Press.
- Goubet N and Clifton RK (1998). Object and event representation in 6.5-month-old infants. *Developmental Psychology*, 34, 63–76.
- Gratch G (1982). Responses to hidden persons and things by 5-, 9-, and 16-month-old infants in a visual tracking situation. *Developmental Psychology*, 18, 232–237.
- Gredebäck G and von Hofsten C (2004). Infants' evolving representations of object motion during occlusion: A longitudinal study of 6- to 12-month-old infants. *Infancy*, 6, 165–184.
- Haith MM and Benson JB (1998). Infant cognition. In: W Damon, series ed. and D Kuhn and R Siegler, vol. eds. *Handbook of child psychology*, Vol. 2, pp. 199–254. New York: Wiley.
- Hespos SJ and Baillargeon R (2001a). Infants' knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science*, 12, 140–147.
- Hespos SJ and Baillargeon R (2001b). Knowledge about containment events in very young infants. *Cognition*, 78, 204–245.
- Hespos SJ and Baillargeon R (in press). Décalage in infants' knowledge about occlusion and containment events: Converging evidence from action tasks. *Cognition*.
- Hofstadter M and Reznick JS (1996). Response modality affects human infant delayed-response performance. *Child Development*, 67, 646–658.
- Hood B and Willatts P (1986). Reaching in the dark to an object's remembered position: Evidence of object permanence in 5-month-old infants. *British Journal of Developmental Psychology*, 4, 57–65.
- Kahneman D, Treisman A and Gibbs BJ (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 174–219.
- Kaldy Z and Leslie AM (2003). Identification of objects in 9-month-old infants: Integrating 'what' and 'where' information. *Developmental Science*, 6, 360–373.
- Kaufman J, Csibra G and Johnson MH (in press). Oscillatory activity in the infant brain reflects object maintenance. *Proceeding of the National Academy of Science of the United States of America*.
- Keil FC (1991). The emergence of theoretical beliefs as constraints on concepts. In: S Carey and R Gelman, eds. *The epigenesis of mind: Essays on biology and cognition*, pp. 237–256. Hillsdale, NJ: Erlbaum.
- Keil FC (1995). The growth of causal understandings of natural kinds. In: D Sperber, D Premack and AJ Premack, eds. *Causal cognition: A multidisciplinary debate*, pp. 234–62. Oxford: Clarendon Press.
- Kestenbaum R, Termine N and Spelke ES (1987). Perception of objects and object boundaries by 3-month-old infants. *British Journal of Developmental Psychology*, 5, 367–383.
- Lécuyer R (2001). Rien n'est jamais acquis. De la permanence de l'objet... de polemiques. *Enfance*, 55, 35–65.
- Lécuyer R and Durand K (1998). Bi-dimensional representations of the third dimension and their perception by infants. *Perception*, 27, 465–472.
- Leslie AM (1984). Infant perception of a manual pick-up event. *British Journal of Developmental Psychology*, 2, 19–32.

- Leslie AM (1994). ToMM, ToBY, and agency: Core architecture and domain specificity. In: LA Hirschfeld and SA Gelman, eds. *Mapping the mind: Domain specificity in cognition and culture*, pp. 119–148. New York: Cambridge University Press.
- Leslie AM (1995). A theory of agency. In: D Sperber, D Premack and AJ Premack, eds. *Causal cognition: A multidisciplinary debate*, pp. 121–149. Oxford: Clarendon Press.
- Leslie AM, Xu F, Tremoulet PD and Scholl BJ (1998). Indexing and the object concept: Developing ‘what’ and ‘where’ system. *Trends in Cognitive Sciences*, 2, 10–18.
- Li J and Baillargeon R (2005a). Change detection in infancy. Event-category effects. Manuscript in preparation.
- Li J and Baillargeon R (2005b). Mapping object representations across successive events in infancy: Positive and negative carry-over effects. Manuscript in preparation.
- Luo Y and Baillargeon R (2005a). Development of infants’ reasoning about transparent occluders. Manuscript in preparation.
- Luo Y and Baillargeon R (2005b). Infants’ reasoning about transparent occluders and containers. Manuscript in preparation.
- Luo Y and Baillargeon R (2005c). When the ordinary seems unexpected: Evidence for incremental physical knowledge in young infants. *Cognition*, 95, 297–328.
- Luo Y, Baillargeon R, Brueckner L and Munakata Y (2003). Reasoning about a hidden object after a delay: Evidence for robust representations in 5-month-old infants. *Cognition*, 88, B23–32.
- Luo Y, Baillargeon R and Lécuyer R (2005). Young infants’ reasoning about height in occlusion events. Manuscript in preparation
- McCall D (2001). *Perseveration and infants’ sensitivity to cues for containment*. Paper presented at the biennial meeting of the Society for Research in Child Development, Minneapolis, MN, April, 2001.
- McDonough L, Choi S and Mandler JM (2003). Understanding spatial relations: Flexible infants, lexical adults. *Cognitive Psychology*, 46, 229–259.
- Meidler M and Gratch G (1980). Do 5-month-olds show object conception in Piaget’s sense? *Infant Behavior and Development*, 3, 265–282.
- Mitchell TM (1997). *Machine learning*. New York: McGraw-Hill.
- Muller AA and Aslin RN (1978). Visual tracking as an index of the object concept. *Infant Behavior and Development*, 1, 309–319.
- Munakata Y (1997). Perseverative reaching in infancy: The roles of hidden toys and motor history in the AB task. *Infant Behavior and Development*, 20, 405–416.
- Munakata Y (2000). Challenges to the violation-of-expectation paradigm: Throwing the conceptual baby out with the perceptual processing bathwater? *Infancy*, 1, 471–490.
- Munakata Y, McClelland JL, Johnson MH and Siegler R (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104, 686–713.
- Needham A (2000). Improvements in object exploration skills may facilitate the development of object segregation in early infancy. *Journal of Cognition and Development*, 1, 131–156.
- Newcombe N, Huttenlocher J and Learmonth A (1999). Infants’ coding of location in continuous space. *Infant Behavior and Development*, 22, 483–510.
- Piaget J (1952). *The origins of intelligence in children*. New York: International Universities Press.
- Piaget J (1954). *The construction of reality in the child*. New York: Basic Books.

- Pylyshyn ZW (1989). The role of location indexes in spatial perception: A sketch of the FINST spatial index model. *Cognition*, 32, 65–97.
- Pylyshyn ZW (1994). Some primitive mechanisms of spatial attention. *Cognition*, 50, 363–384.
- Quinlan JR (1993). *C4.5: Programs for machine learning*. San Mateo, CA: Morgan Kaufmann.
- Rensink RA (2002). Change detection. *Annual Review of Psychology*, 53, 245–277.
- Rensink RA, O'Regan JK and Clark JJ (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–373.
- Rivera SM, Wakeley A and Langer J (1999). The drawbridge phenomenon: Representational reasoning or perceptual preference? *Developmental Psychology*, 35, 427–435.
- Roder BJ, Bushnell EW and Sasseville AM (2000). Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy*, 1, 491–507.
- Schilling TH (2000). Infants' looking at possible and impossible screen rotations: The role of familiarization. *Infancy*, 1, 389–402.
- Scholl BJ and Leslie AM (1999). Explaining the infants' object concept: Beyond the perception/cognition dichotomy. In: E Lepore and Z Pylyshyn, eds. *What is cognitive science?*, pp. 26–73. Oxford: Blackwell.
- Siegler RS (1978). The origins of scientific reasoning. In: RS Siegler, ed. *Children's thinking: What develops?*, pp. 109–149. Hillsdale, NJ: Erlbaum.
- Simon T, Hespos SJ and Rochat P (1995). Do infants understand simple arithmetic? A replication of Wynn (1992). *Cognitive Development*, 10, 253–269.
- Simons DJ (1996). In sight, out of mind: When object representations fail. *Psychological Science*, 7, 301–305.
- Simons DJ (2000). Current approaches to change blindness. *Visual Cognition*, 7, 1–15.
- Sitskoorn SM and Smitsman AW (1995). Infants' perception of dynamic relations between objects: Passing through or support? *Developmental Psychology*, 31, 437–447.
- Slater A (1995). Visual perception and memory at birth. In: C Rovee-Collier and LP Lipsitt, eds. *Advances in infancy research*, Vol. 9, pp. 107–162. Norwood, NJ: Ablex.
- Spelke ES (1982). Perceptual knowledge of objects in infancy. In: J Mehler, E Walker and M Garrett, eds. *Perspectives on mental representation*, pp. 409–430. Hillsdale, NJ: Erlbaum.
- Spelke ES (1994). Initial knowledge: Six suggestions. *Cognition*, 50, 431–445.
- Spelke ES, Breinlinger K, Macomber J and Jacobson K (1992). Origins of knowledge. *Psychological Review*, 99, 605–632.
- Spelke ES and Kestenbaum R (1986). Les origines du concept d'objet. *Psychologie Française*, 31, 67–72.
- Spelke ES, Phillips A and Woodward AL (1995). Infants' knowledge of object motion and human action. In: D Sperber, D Premack and AJ Premack, eds. *Causal cognition: A multidisciplinary debate*, pp. 44–78. Oxford: Clarendon Press.
- Thelen E and Smith LB (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Wang S and Baillargeon R (in press). Infants' physical knowledge affects their change detection. *Development Science*.
- Wang S and Baillargeon R (2005a). Can infants' physical knowledge be trained? Manuscript in preparation.
- Wang S and Baillargeon R (2005b). Inducing infants to detect a physical violation in a single trial. *Psychological Science*, 16, 542–549.

- Wang S, Baillargeon R and Brueckner L (2004). Young infants' reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition*, 93, 167–198.
- Wang S, Baillargeon R and Paterson S (2005). Detecting continuity violations in infancy: A new account and new evidence from covering and tube events. *Cognition*, 95, 129–173.
- Wellman HM and Gelman SA (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, 43, 337–375.
- Wilcox T (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition*, 72, 125–66.
- Wilcox T and Baillargeon R (1998a). Object individuation in infancy: The use of featural information in reasoning about occlusion events. *Cognitive Psychology*, 17, 97–155.
- Wilcox T and Baillargeon R (1998b). Object individuation in young infants: Further evidence with an event-monitoring task. *Developmental Science*, 1, 127–142.
- Wilcox T and Chapa C (2002). Infants' reasoning about opaque and transparent occluders in an object individuation task. *Cognition*, 85, B1–10.
- Wilcox T and Chapa C (2004). Priming infants to attend to color and pattern information in an individuation task. *Cognition*, 90, 265–302.
- Wilcox T, Nadel L and Rosser R (1996). Location memory in healthy preterm and fullterm infants. *Infant Behavior and Development*, 19, 309–323.
- Wilcox T and Schweinle A (2003). Infants' use of speed information to individuate objects in occlusion events. *Infant Behavior and Development*, 26, 253–282.
- Wilson RA and Keil FC (2000). The shadows and shallows of explanation. In: FC Keil and RA Wilson, eds. *Explanation and cognition*, pp. 87–114. Cambridge, MA: MIT Press.
- Yonas A and Granrud CE (1984). The development of sensitivity to kinetic, binocular, and pictorial depth information in human infants. In: D Engle, D Lee and M Jeannerod, eds. *Brain mechanisms and spatial vision*, pp. 113–145. Dordrecht: Martinus Nijhoff.