Available online at www.sciencedirect.com

ScienceDirect

Cognitive Psychology 57 (2008) 262–291

Cognitive Psychology

www.elsevier.com/locate/cogpsych

The origins of causal perception: Evidence from postdictive processing in infancy ☆

George E. Newman *, Hoon Choi, Karen Wynn, Brian J. Scholl

Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA

Available online 13 June 2008

Abstract

The currency of our visual experience consists not only of visual features such as color and motion, but also seemingly higher-level features such as causality—as when we see two billiard balls collide, with one causing the other to move. One of the most important and controversial questions about causal perception involves its origin: do we learn to see causality, or does this ability derive in part from innately specified aspects of our cognitive architecture? Such questions are difficult to answer, but can be indirectly addressed via experiments with infants. Here we explore causal perception in 7-month-old infants, using a different approach from previous work. Recent work in adult visual cognition has demonstrated a postdictive aspect to causal perception: in certain situations, we can perceive a collision between two objects in an ambiguous display even after the moment of potential 'impact' has already passed. This illustrates one way in which our conscious perception of the world is not an instantaneous moment-by-moment construction, but rather is formed by integrating information over short temporal windows. Here we demonstrate analogous postdictive processing in infants' causal perception. This result demonstrates that even infants' visual systems process information in temporally extended chunks. Moreover, this work provides a new way of demonstrating causal perception in infants that differs from previous strategies, and is immune to some previous types of critiques.
Published by Elsevier Inc.

Keywords: Perception of causality; Event perception; Postdiction; Innateness; Michotte

☆ For helpful conversation and/or comments on earlier drafts, we thank Renée Baillargeon, Erik Cheries, Bruce Hood, Alan Leslie, Peter White, and an anonymous reviewer. We also thank Diana Damrosch, Lisa Halbsgut, Kiley Hamlin, Sydney Levine, Danielle Marquis, Koleen McCrink, Kathryn Melamed, and Ashley Norman for help testing the infants.

* Corresponding author. Fax: +1 203 432 7172.
E-mail addresses: george.newman@yale.edu (G.E. Newman), h.choi@yale.edu (H. Choi), karen.wynn@yale.edu (K. Wynn), brian.scholl@yale.edu (B.J. Scholl).

0010-0285/$ - see front matter Published by Elsevier Inc.
doi:10.1016/j.cogpsych.2008.02.003

1. Introduction

It is difficult to see how we could make sense of almost any aspect of the physical world without an appreciation of cause and effect. As such, this critical ability has received a great deal of study both in philosophy and psychology (see Sperber, Premack, & Premack, 1995). Philosophical theories of causality typically address the world directly—asking metaphysical questions about how (or whether) this property actually exists. Such theories may be deflationary, suggesting that causal inferences are not justified. For example, Hume famously argued that causal inferences are not supported even in simple mechanical interactions such as two billiard balls colliding: "Motion in the second Billiard ball is a quite distinct event from motion in the first; nor is there any thing in the one to suggest the smallest hint of the other" (Hume, 1748/1977, p. 18). Psychological theories of causality, in contrast, have explored the many ways in which the human mind traffics in causality despite such warnings. Much of this work explores higher-level cognitive processing, focusing on the dynamics of causal inferences (e.g., Sperber et al., 1995; White, 1995). Some of the most salient work on causality in psychology, however, has focused on perceptual processing.

1.1. The perception of causality

The study of causal perception was famously inaugurated by the Belgian psychologist Albert Michotte in his landmark book The Perception of Causality (1946/1963). Michotte had a great talent for noticing fairly mundane aspects of our visual experience, and then recognizing often deep implications that those observations had for the nature of perceptual processing (see Wagemans, Van Lier, & Scholl, 2006). In the case of causal perception, Michotte simply noticed that mechanical collisions are perceived in ways that transcend their spatiotemporal dynamics. When we see a collision between two billiard balls, for example, we do not simply see the cessation of one ball's motion followed by the onset of motion in the other: instead, we see one ball *cause* the other's motion (cf. Fig. 1a). Michotte further noticed (in over 100 studies reported in his book) that the perception of causality seemed to be determined by a highly constrained collection of visual cues, and was largely unaffected by higher-level beliefs or intentions. Accordingly, research following Michotte has largely been focused on the attempt to work out just what types of visual cues drive causal perception in simple mechanical interactions—for example, focusing on the relative and absolute speeds of the objects, various types of spatial and temporal gaps in the their trajectories, the smoothness and continuity of motion, differences in the durations and angles of each object's movement, cross-modal interactions, effects of perceptual grouping and attention, and several other factors (e.g., Boyle, 1960; Choi & Scholl, 2004, 2006a; Costall, 1991; Gordon, Day, & Stecher, 1990; Guski & Troje, 2003; Hubbard & Ruppel, 2002; Kruschke & Fragassi, 1996; Michotte & Thinés, 1963/1991; Natsoulas, 1961; Schlottmann & Anderson, 1993; Schlottmann, Ray, Mitchell, & Demetriou, 2006; Schlottmann & Shanks, 1992; Scholl & Nakayama, 2002, 2004; Weir, 1978; White, 2005a, 2005b, 2006; White & Milne, 1997, 1999; Yela, 1952). Other research has demonstrated more directly that causal perception is distinct from higher-level causal inference, dem-

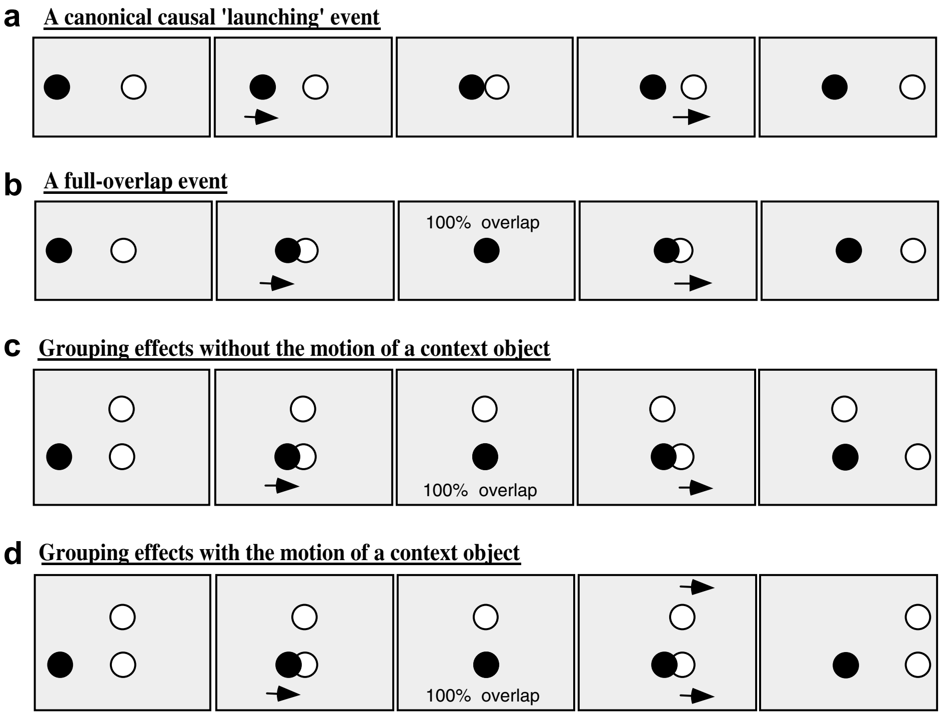

Fig. 1. Four standard visual events employed in studies of causal perception. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion. (a) A canonical 'causal launching' event, wherein one disc is seen to collide with a second disk, causing it to move. (b) A 'full-overlap' event, which observers can readily perceive as noncausal 'passing', wherein one moving object is seen to pass right over another stationary object (swapping colors). (c and d) Effects of perceptual grouping on causal perception: (c) when an additional context object remains stationary even after the second object begins to move, observers perceive the full-overlap event as noncausal passing. (d) When an additional context object remains grouped with the initially-stationary object via proximity and common motion, observers readily perceive the full-overlap event as causal launching. Animations of all conditions used in this study are available online at http://www.yale.edu/perception/infant-postdiction/.

onstrating for example that these two classes of interpretations can even pull in entirely opposite directions (e.g., Schlottmann & Shanks, 1992).

A general conclusion that can be drawn from this body of work is that the phenomenon of causal perception is aptly named: such impressions seem to reflect automatic visual processing, and are insulated to some degree from other aspects of cognition. For example, we can see—and can barely avoid seeing—causality in some situations even when we are certain it does not exist (as when a real billiard ball appears to cause a patch of light to move), but we will fail to perceive causality in many other situations when we know that it does exist. These observations have convinced several contemporary theorists that causal perception reflects a type of modular visual processing, similar to the processing of other visual features such as motion or form (e.g., Choi & Scholl, 2004; Leslie, 1986, 1988; Saxe & Carey, 2006; Scholl & Tremoulet, 2000)—though such proposals remain controversial (e.g., Cohen, Amsel, Redford, & Casasola, 1998; Schlottmann, 2000; Schlottmann et al., 2006; Young, Rogers, & Bechmann, 2005).

1.2. The origins of causal perception

The powerful and subtle nature of causal perception has lead directly to the question of where it comes from. All theorists seem to agree that such processing is driven by learning, but they differ about the scope of such learning. Some theorists prefer interpretations which isolate learning to a local ontogenetic timescale, suggesting that all humans develop the capacity and tendency to perceive causality as part of a domain-general process of perceptual development (e.g., Cohen et al., 1998). Other theorists, in contrast, suggest that powerful learning routines can also operate at longer phylogenetic timescales, resulting in the type of innate perceptual routines for perceiving causality that are also often assumed to be characteristic of other aspects of visual processing (e.g., Leslie, 1986; Scholl & Tremoulet, 2000).

Such debates are difficult to resolve empirically, of course, but one common strategy is to explore the nature of such processing developmentally. If even young infants are able to exhibit such competence, this at least severely limits the available time and scope of experience in giving rise to these abilities. Michotte maintained a nativist theory of the origin of causal perception for many reasons (see Saxe & Carey, 2006), but he did not conduct any developmental experiments: "It would clearly be very interesting if experiments such as those described in this book could be tried out on children of different ages... Unfortunately, plans for such research have not yet advanced beyond the project stage" (1946/1963, p. 255).

Though many developmental researchers (including Piaget) theorized about the developmental origin of the idea of cause and effect (e.g., Piaget, 1936/1963), perhaps the first infant studies of mechanical causal perception were conducted by Ball (1973).[1] In these studies, infants were habituated to an event that is naturally perceived by adults in terms of an occluded causal interaction. Infants were initially shown a stage containing a narrow central screen that partially obscured a stationary white block at its left edge. They then watched as a red block moved from the other side of the stage until it too was partially occluded by the screen at its right edge, at which point it stopped moving, and the white block began moving in the same direction. For adults, this display has a natural interpretation: a collision between the two blocks, with the moment of impact obscured by the screen. Infants of several ages (as young as 6 months) were habituated to this event, and their looking times were then measured to test events in which the motion of the two objects remained highly similar, but the screen was removed. Infants looked longer at test events during which the two objects did not contact each other, compared to when there was a visible connection (Ball, 1973). This effect can be interpreted in terms of infants' causal perception: the no-contact test event involved a switch from the habituated causal displays to a new noncausal display, whereas the test event with contact was just yet another instance of a causal interaction.

This initial experiment (later replicated by others; e.g., Kosugi, Ishida, & Fujita, 2003; Spelke, Phillips, & Woodward, 1995) provides an elegant way to study the inferences that infants will automatically draw on the basis of their causal percepts, though the basic design has been criticized in several ways (see Cohen et al., 1998, pp. 173–176). For present purposes, however, this study has a more fundamental problem: its experimental logic still depends on the assumption of causal perception in the first place. After all, it could be that infants simply make predictions based on the statistical properties of previously encountered

[1] There have also been several developmental studies of causal perception and inference in older children (e.g. Olum, 1956; Schlottmann, 1999; Schlottmann, Allen, Linderoth, & Heskith, 2002) but these studies do not typically have direct implications for the origins of causal perception in infancy, and are not discussed here.

'contact' events without necessarily perceiving that event in causal terms—e.g., "When A arrives adjacent to B, A typically stops, and B starts moving." This possibility also compromises the ability of several other studies of similar "causal" inferences to truly isolate causality per se. For example, Kotovsky and Baillargeon (2000) demonstrated that 7.5-month-old infants use knowledge about the possibility of spatial contact to inform their predictions about whether one object will cause another to move: when a barrier lies in the path between two "colliding" objects, infants will not expect the initially stationary object to move. Again, however, this inference could be statistically driven without any actual perception of causality. Similarly, the fact that infants can distinguish causal launching from noncausal events with spatial or temporal gaps (Leslie, 1982, 1984) implies that they are sensitive to these spatiotemporal differences, but does not actually require causal percepts: such differences could be especially salient based on their ability to predict ensuing motions, without actually looking causal. (For other examples of studies which explored the inferences drawn by infants on the basis of their causal percepts, see Saxe & Carey, 2006.)

This difficulty in isolating causality was eventually surmounted via the use of a novel 'reversal' logic (Leslie & Keeble, 1987). Six-month-old infants were first repeatedly shown a film of a red brick which moved until it was adjacent to a green brick, at which point the red brick stopped and the green brick began moving along the same trajectory. Once infants habituated to this event, they were then shown a test event in which the animation was simply reversed (implemented by actually playing the same film backwards through the projector). Two groups of infants were tested in the key experiment, differing only in the nature of the timing of the habituation event: infants in the 'Delayed Reaction' condition saw a pause between the motions of the two bricks, while infants in the 'Direct Launching' condition saw the second brick move immediately after the arrival of the first brick (as in Fig. 1a).

For adults, the introduction of a sizeable temporal gap between two motions typically destroys causal percepts (Michotte, 1946/1963). This factor also appeared to drive the results of Leslie and Keeble's (1987) study: infants in the Direct Launching condition dishabituated considerably more than infants in the Delayed Reaction condition. It is perhaps not surprising that dishabituation occurred for both groups, since the spatiotemporal dynamics were always novel (e.g., including motion in a new direction). But what led to greater dishabituation in the Direct Launching condition? Leslie and Keeble's hypothesis was that this event was now perceived with reversed causal roles: the red brick, previously the 'launcher', was now the 'launchee'. In contrast, no such different should occur in the Delayed Reaction condition, since by hypothesis this event wasn't perceived in terms of causal roles in the first place.

The fact that the relatively young infants in this study appeared to be to be sensitive to causal roles led Leslie and colleagues to conclude (1) that infants perceive causality, and (2) that causal perception may be realized via an innate and possibly modular process in the visual system (see also Leslie, 1986, 1988).[2] Other researchers have remained uncon-

[2] In fact, Leslie's theory goes on to distinguish between "the Michotte module", which is responsible for the perception of causality, and the mechanics module, "ToBy", which interprets movements and arrangements of objects relative to their potential force. Leslie (1994) therefore interprets the results of Leslie and Keeble (1987) as evidence that infants assign mechanical 'roles' to objects, and thus possess both kinds of modules. Though this distinction may have important implications for thinking about how causal perception might be instantiated in the mind, it is unclear how such a contrast directly maps onto the studies reported here. In any case, note that there are also several other studies of higher-level aspects of causal perception in infancy, e.g. exploring infants' ability to reason about likely mechanical interactions in collisions among objects of different sizes (e.g. Kotovsky & Baillargeon, 1994; Kotovsky & Baillargeon, 1998; Wang, Kaufman, & Baillargeon, 2003).

vinced, however, positing more 'constructivist' and domain-general developmental accounts (e.g., Cohen, 1998; Cohen & Amsel, 1998; Cohen & Oakes, 1993; Cohen et al., 1998; Oakes, 1994; Oakes & Cohen, 1990). These views have been fueled in part by other developmentally oriented studies. For example, one study found evidence for causal perception in 6-month-olds but not 4-month-olds (Cohen & Amsel, 1998); another study using more complex objects found evidence for causal perception in 10-month-olds but not 6-month-olds (Oakes & Cohen, 1990); and yet another study using objects which varied from trial to trial failed to find evidence for causal perception even in 10-month-olds (Cohen & Oakes, 1993). Moreover, these and other studies have demonstrated that causal perception is not an all-or-nothing achievement, and that its manifestation at different ages will be acutely sensitive to the nature and complexity of the objects used (Oakes, 1994). These developmental studies have been further modeled in an architecture that does not require domain-specific built-in notions of causality (Chaput & Cohen, 2001).

1.3. The status of the debate

The work described above has been used in support of two very different models of the origin of causal perception. On one hand, Leslie and colleagues argue that the notion of causation is rooted in an innate, domain-specific visual module—an "automatic starting engine" that operates "automatically and incorrigibly upon the spatiotemporal properties of events yet producing abstract descriptions of their causal structure" (Leslie, 1988, pp, 186–187, 194). In contrast, Cohen and colleagues draw conclusions that differ in almost every respect. For them, causal perception is demonstrably not innate: the failure of Oakes and Cohen (1990) to observe causal perception with real objects in 6-month-olds, for example, is an "embarrassment", "exactly the opposite of what a proponent of innate modularity would predict", and "there is no room in [this] formulation for causality perception to develop in stages" (Cohen et al., 1998, pp. 172, 183, 186). Nor is it automatic or modular: such results "provide little support for the notion that causality is directly or automatically perceived in events characterized by a set of necessary cues to causality" (Oakes, 1994, p. 878). And it is certainly not domain-specific: "infants come to perceive the causal nature of events as a result of their general perceptual development" (Oakes & Cohen, 1990, p. 193).

What to conclude from this previous work? It is difficult, of course, to conclusively prove the innate modularity view correct. However, in our view, the results and arguments of Cohen and colleagues do little to damage its plausibility for at least three reasons:

First, claims about specific ages of infants are largely irrelevant, since even innate modules may require both triggering and maturation. (To trot out a cliched example, puberty is thought to be driven by relatively specialized and innate processes, but is seldom seen in 6-month-old infants.) Thus a failure to observe causal perception in, say, 4-month-olds is not an "embarrassment" to this view—though the successful demonstration of causal perception in 6-month-olds does provide a useful upper bound on the amount of experience and instruction which is necessary for causal perception.

Second, claims about developmental progressions are also largely irrelevant (cf. Schlottmann, 2000). Cohen and colleagues take it to be an important discovery that "infant perception of causality does appear to develop" (Cohen & Oakes, 1993, p. 431), and this is taken to refute the innate and modular perspective. But Leslie's proposal for a 'causal module' is intended—like most modular proposals—to provide a sort of jumpstart to development, providing important abstract concepts that would otherwise be tremendously difficult to acquire (see Scholl & Leslie, 1999, 2001). It in no way denies a role for development; indeed, Leslie (1988) has explicitly noted that a modular origin for causal perception "is perfect for a mechanism whose job it is to help produce development" (p. 194).

Third, Cohen and colleagues have demonstrated in several ways that objects' features also influence causal perception, and they take this also to argue against a modular perspective. They point, in particular, to the fact that adults' causal perception does not depend much on the nature or surface features of the objects involved; what matters are only those objects' spatiotemporal dynamics: "the modular view assumes the perception of causality is automatic and independent of the particular objects used to portray the causal event. All the module seems to require is an input that has spatial and temporal contiguity..." (Cohen & Oakes, 1993, p. 431). They then assume that the operation of a causal module in infancy must similarly be blind to surface features, which (they argue) is not the case, since the complexity of the objects used will dramatically affect causal perception in both 6-month-olds (Oakes, 1994) and 10-month-olds (Cohen & Oakes, 1993).

In fact, the premise of this argument seems correct: a causal module, following Michotte, may in fact be largely blind to surface features. However, the demonstration that variance in surface features will influence the expression of causal perception does not show this. After all, the use of complex displays may simply overwhelm young infants' attention, and prevent them from focusing on the relevant features. This would surely be possible in adults too, given complex enough displays—but a failure to perceive causality in an event amidst Times Square would not disprove the operation of a causal module. Rather, it would just point out that perception of almost any event requires some degree of attention. (To drive this point home, notice that adults will fail to perceive entire salient objects directly in their field of view, given enough competing attentional demands—as in "inattentional blindness"; see Most, Scholl, Clifford, & Simons, 2005—but this does not disprove the existence of relatively hardwired mechanisms of object perception.)

Though we are unpersuaded by the arguments of Cohen and colleagues about the implausibility of the modular perspective, we are in full agreement about where to go next: "Whatever methods investigators choose to use to examine infant causal perception, they should go beyond asking if infants do or do not have causal perception and begin asking more detailed questions about the nature of that perception" (Cohen & Oakes, 1994, p. 431). In this paper, we follow this approach, drawing inspiration from recent adult studies that identify some peculiar temporal aspects of causal perception.

1.4. Grouping and postdiction in adults' causal perception

Nearly all of Michotte's studies of causal perception focused on the necessary spatiotemporal properties of the two objects actually involved in a putative collision. In contrast, more recent work has demonstrated that context also matters: whether an event involving two objects will be perceived in causal terms depends not only on the dynamics of those objects, but also on the dynamics of other independent objects and events in a display. In the initial demonstration of this grouping effect in adults, observers viewed a 'full-overlap' event that is similar to a canonical launching event (Fig. 1a) except that the two discs fully overlapped before the first disc (A) stopped and the second disc (B) began its motion (Fig. 1b). In isolation, this event can be readily perceived in terms of noncausal 'passing': one object remains stationary in the center of the screen, while another passes over it

(Scholl & Nakayama, 2002). This perception occurs despite salient feature differences between the objects. Though a red disc (A) may move to the center and stop, while a green disc (B) may move from the center after the arrival of the red disc, this is not what we see: rather, we see a single moving disc which changes color from red to green, and a single stationary disc which changes color from green to red.

A striking effect emerges, however, when a single additional object (C) is added to the display: the behavior of this completely independent object can control whether the full-

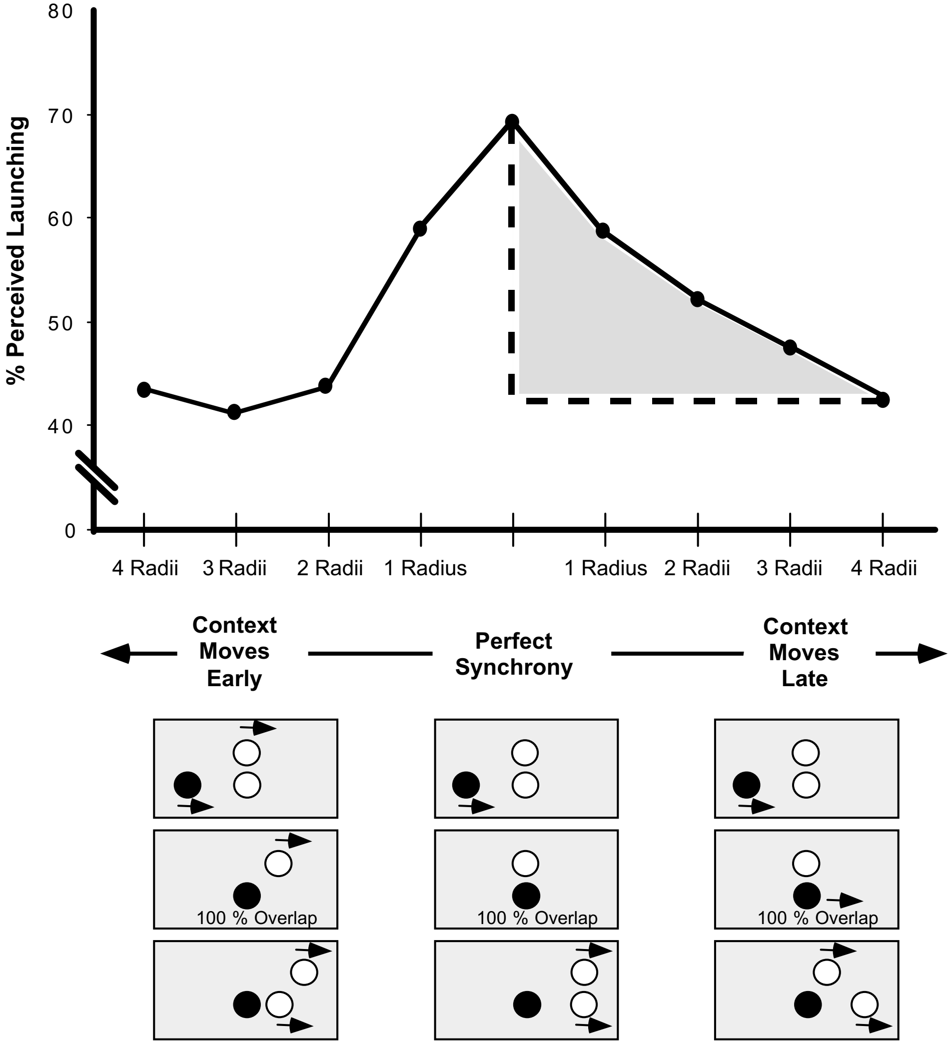

Fig. 2. The percentage of perceived causal launching for each temporal offset in Experiment 1 of the adult study of Choi and Scholl (2006b). The motion of the context object affected causal perception even when it was not temporally synchronized with the other motions. The dashed line indicates the result predicted by a view in which conscious awareness of the event is a moment-by-moment construction. The gray shaded area thus represents postdictive effects, in which the onset of the contextual motion yields a perceived collision even after the moment of full-overlap has already passed. (See the text for discussion.)

overlap event is seen in terms of causal launching or noncausal passing (Choi & Scholl, 2004). In all cases, C is initially vertically aligned with B, near the center of the display. When C stays in this stationary position throughout the animation (as in Fig. 1c), it has no effect, and observers can still readily perceive noncausal passing. However, when C remains aligned with B through its motion—i.e., remaining stationary until A and B fully overlap, and then moving along with B thereafter (as in Fig. 1d)—observers readily perceive the full-overlap event in terms of causal launching. This effect can be appreciated as a type of perceptual grouping based on proximity and common motion: these cues bind B and C into a single perceptual group (since otherwise it would be an unlikely coincidence that B and C moved at exactly the same time). Because of these grouping cues, observers do not perceive A and B to overlap. Rather, they perceive A only to partially overlap with the group formed by B and C—and because there is no ambiguity about the full-overlap, there is also no ambiguity about what is moving where, and therefore observers can readily perceive causal launching.

This grouping effect was later exploited in a study of the temporal dynamics of causal perception, by varying just when the context disc C moved relative to B (Choi & Scholl, 2006b). In fact, C does not need to move exactly when B does: the common motion, in other words, need not be perfectly "common". Even when C begins moving 200 ms before B (as in the left side of Fig. 2), observers are still more likely to perceive causal launching than if C remains stationary, or begins moving much earlier (360 ms) than B. This is in essence a measure of how much "commonality" among the motions is required to trigger perceptual grouping.

What was surprising about this study, however, was that this temporal window was symmetric about the point of full-overlap: causal launching was also perceived when C began moving 180 ms after B (as in the right side of Fig. 2). At first blush, this seems like a paradoxical result, as depicted in Fig. 3. The motions by which A and B become separated are fully visible, and during this time (before C beings its motion) the visual system must interpret this motion as belonging to one of the objects. (In other words, when the objects begin to separate, you must see the motion as belonging to either A or B. In this context, you cannot just see motion; you must see motion *of an object*.) The question, then, is which object is seen to move during this period, and the answer to this question is what effectively determines launching vs. passing. If the initial post-overlap motion is bound to A, you see noncausal passing (i.e., a continuation of A's previous motion), but if the motion is bound to B, you see causal launching (i.e., the sudden onset of a new object moving). We know that if C remains stationary observers are likely to see the movement of A (noncausal passing). This implies that in situations where C moves after B, observers should perceive noncausal passing in the window after A and B are no longer fully overlapped, but before C has moved. However, people continue to perceive causal launching even when C moves up to 180 ms after B (by which point A and B have become fully separated; Choi & Scholl, 2006b).[3]

[3] This particular duration is interesting to consider in the context of the suggestion that the duration of iconic processing might be critical for causal perception (White, 1988). Events that fall within the iconic time span may be seen as causally connected, while events that straddle this boundary are not. Originally, this idea was applied to Michotte's research on delays in launching events, which pointed to a similar critical delay of about 170 ms. This lends additional support to the notion that causal perception is served by integrating information over temporally extended chunks.

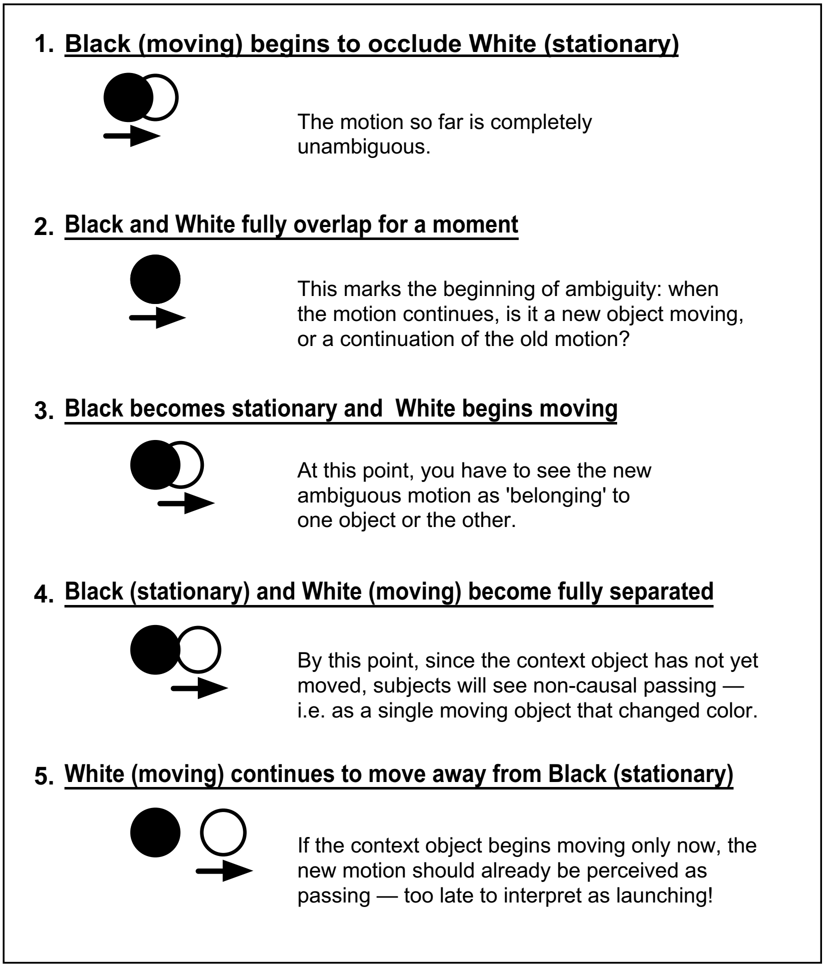

Fig. 3. A "moment-by-moment" analysis of postdictive effects on causal perception. Arrows indicate motion. (1) Motion is unambiguously bound to the moving Black object. (2) The Black object moves until it fully occludes the White object. (3) The motions by which the objects become separated are fully visible, at which point the motion can either be perceived as the onset of motion in a new object (causal launching), or as the continuation of motion with a color change (noncausal passing). (4) If the context object continues to remain stationary then observers are likely to perceive noncausal passing. (5) By this point the objects have separated and it should be too late to 're-perceive' this motion in terms of causal launching when the context object now begins to move. However, in fact people may continue to perceive causal launching within this window, at point after which the two objects have already separated.

In short, this effect demonstrates the postdictive nature of causal perception: a current stimulus can effectively 'rewrite' the immediate perceptual past in our conscious visual awareness. Such an effect illustrates that our conscious perception of the world is not an instantaneous moment-by-moment construction, but rather is formed by integrating information over short temporal windows (see also Dennett, 1991; Gao & Scholl, 2007; Eagleman & Sejnowski, 2003; Mitroff & Scholl, 2004). Here, a stimulus which should already be interpreted in terms of noncausal passing is effectively reinterpreted in terms of causal launching 'after the fact', despite the fact that the 'trigger' for this perception doesn't occur until it should be too late to have an effect.[4] Thus, the cues which drive causal perception in a putative collision are not limited to the objects in that collision, but can also encompass other objects and events which are nearby in both space and time.

1.5. The current experiments

Here we report the first tests in infancy of several phenomena borrowed from the adult visual cognition literature, including the contrast between causal launching and noncausal passing, the influence of grouping effects on causal perception, and the existence of postdictive processing.

The logic of the present studies is relatively straightforward. Seven-month-olds infants are habituated to a grouping display in which the motion of an isolated object (C) leads adults to perceive a full-overlap event between A and B in terms of causal launching, via its perfectly synchronized common motion with B.[5] We hypothesize that this manipulation will similarly influence infants' perception, such that they will effectively be habituated to an instance of causal launching. Infants are then tested with similar visual events, but in which C begins moving after A and B have separated. In Experiment 1 this temporal offset is considerable, which we predict may lead to a percept of noncausal passing (as in adults), and therefore to dishabituation based on the changed causal status of the display. In Experiment 2 this offset is more modest, which we predict will lead to percepts of causal launching via postdictive processing (as in adults)—and therefore to a failure to recover looking time despite the novel temporal offset, since the perceived causal status of the display will not have changed. In Experiments 3 and 4, we provide two demonstrations that infants are in fact sensitive to the shorter temporal offset itself, even though it may not lead to dishabituation in the full-overlap event. Finally, in Experiment 5, we provide a demonstration that the longer temporal offset does not lead to dishabituation by itself, without an associated change in the perceived causal status of an event. In sum, we predict that infants will dishabituate to these displays only when there is a change in causal status, and not due to lower-level variables.

These five studies are collectively designed to meet three interrelated goals: First, we seek to demonstrate causal perception in infants using a new experimental manipulation. Second, we seek to test one strong prediction of the modular view of the origins of causal perception. Recall that Michotte himself argued for what today we might interpret as a modular view of causal perception (see Scholl & Tremoulet, 2000; cf. Saxe & Carey, 2006), and that this helped to motivate Leslie's developmental hypothesis. If these modules are identical in infants and adults, then infants' causal perception should also be affected by the same types of features that drive causal perception in adults—including the effects

[4] Note that this is the same paradoxical logic that is involved in nearly any demonstration of apparent motion. When we see a disc move via apparent motion from location A to B, for example, we can in some cases see the motion traverse the intermediate positions. However, this seems impossible: the visual system cannot infer motion between A and B until B has already occurred—otherwise how would the visual system know which direction A is going to move? By the time that the flash at B has occurred, however, it should be too late to perceive motion leading up to B. Thus the perception of apparent motion itself indicates a type of postdictive processing.

[5] Animations of these and all displays used in our experiments are available online, at http://www.yale.edu/perception/infant-postdiction/.

of grouping and postdiction described above. Third, we seek to test whether postdictive processing operates in infancy for independent reasons: such processing has been identified in many areas of adult visual cognition, but to our knowledge has never been explored in infants.

2. Experiment 1: Grouping in causal perception

Infants in this first experiment were habituated to a full-overlap "grouping" event containing three green squares, perceived by adults in terms of causal launching: one square (A) moved until it was completely overlapped with a second square (B), at which point A stopped moving, and B and C both immediately began moving in the same direction (see Fig. 4a). After habituation, infants were then presented with two new test events involving novel objects. In a Synchronous Test event, infants were shown another animation in which B and C again moved perfectly synchronously following A's motion (as in the habituation event). The Large Offset Test event was identical, except that C now began moving 600 ms after the overlap between A and B. In adults, this duration of temporal offset exceeds the postdiction window: objects B and C do not group by common motion, and adults readily perceive noncausal passing. We predict that this pattern will also obtain in infants, and that they will dishabituate to this test event based on the changed causal status of the display (from launching during habituation, to passing during test). Furthermore, we predict that the degree of dishabituation in this event will exceed that in the Synchronous Test event, wherein the perceived causal status does not change.

2.1. Method

2.1.1. Participants

Twelve 7-month-old infants with a mean age of 6 months 28 days (range: 6 months 14 days to 7 months 12 days) participated. Two additional infants were tested but were excluded because of fussiness or disinterest.

2.1.2. Stimuli

Habituation trials involved two components: one visual event consisting of two objects (A and B), and one additional context object (C). Each object was a green square subtending 3.18°, drawn on a white background. Motion was always in the horizontal plane, since the perception of causality is weaker in other orientations (Michotte, 1946/1963).

A and B always moved in a 'full-overlap' event, as follows: A and B began horizontally aligned, with their centers 9.54° above the lower display border. A began near the left edge of the display, so that its most extreme edge was 2.12° from the display border. B was always initially drawn near the center of the display, such that the nearest edges of A and B were separated by 10.07°. At the beginning of each trial, A and B expanded and contracted 3 times: in a period of 100 ms, each square grew to a size of 4.24° before returning to its original size. The purpose of these motions was to draw infants' attention to A and B. After 2.2 s, A began moving at 17.65°/s toward B. When the two objects became fully overlapped (with A always drawn on top of B; cf. Scholl & Nakayama, 2004), A stopped moving and B instantly started moving at the same speed toward the other edge of the display (stopping when its most extreme edge was 2.12° from the display border).

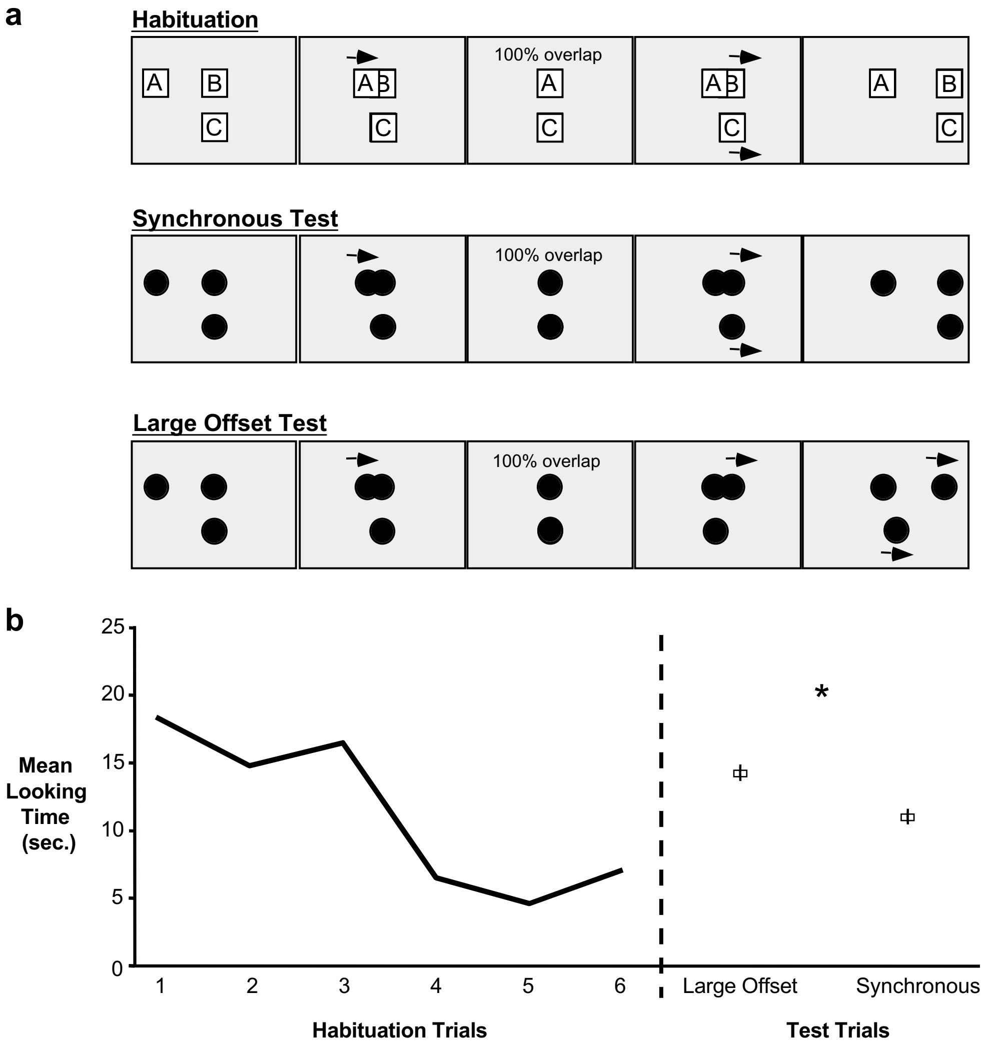

Fig. 4. (a) Depiction of the events presented to infants in Experiment 1. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion, and the letters did not appear in the actual displays. Infants were habituated to a full-overlap display in which an additional context object remains grouped with the initially-stationary object via proximity and common motion (perceived by adult observers as causal launching). After habituation criteria were met, infants were presented with two new test events featuring novel objects: a Synchronous Test event in which the context object continued to move synchronously with object B (i.e., with the same motion as during habituation) and a Large Offset Test event in which the context object moved 600 ms after B began to move (perceived by adult observers as noncausal passing). (b) Duration of looking-time across Habituation and Test events in Experiment 1. The asterisk indicates a statistically significant difference in the degree of dishabituation ($p < .05$).

C was initially positioned directly below B, so that their nearest edges were vertically separated by 3.18°. During habituation, C began moving at the same time—and in the same direction, with the same speed—as B. B and C then stopped together, at which point the display remained stationary for 100 ms before disappearing. This entire event lasted 5.24 s.

The Synchronous Test event was identical to the habituation trials, but the objects were now three red discs. The Large Offset Test event was identical, except that C began its motion 600 ms after B began moving. C then continued to move after B stopped until it was directly under B, at which point the display disappeared. This entire event lasted 5.88 s.

2.1.3. Procedure and design

Infants were seated in a car seat on a table, approximately 100 cm from a projected computer display subtending 39.5° by 28°. Mothers were present in the room during testing, but were positioned behind a curtain, facing the opposite direction from the computer monitor (so that they could not influence their child's attention). During each habituation trial, infants were repeatedly shown the habituation animation described above (see Fig. 4a). Each habituation trial ended when the infant looked away from the monitor for 2 consecutive seconds or when 34 s had elapsed. Habituation trials continued until infants met the habituation criterion, defined as 3 consecutive trials with summed looking time less than or equal to 50% of the sum of the looking times on the first 3 trials. Infants were presented with a minimum of 6 and a maximum of 14 habituation trials.

After infants reached the habituation criterion, they were presented with 6 test trials alternating between the Synchronous Test event and Large Offset Test event, with order (Synchronous first or Large Offset first) counterbalanced.

2.1.4. Data analysis

To be included in analyses, infants had to see at least two complete pairs of test events. A given test trial was considered to be completed if the infant watched for at least 3 s (the length of time it took for B to begin moving). Looking-time to each test event was measured by an observer who was hidden behind a curtain and was unaware of the infant's habituation group. Timing began at the beginning of the trial. A second experimenter, naive to the experimental design, reviewed video footage on 25% of trials and measured the infants' looking times to the test movies. In this and the following experiments, these times were found to correspond well to the on-line timing ($r = .97$ or better), and thus all data analyses were performed using results from the on-line timing. Infants completed an average of 9 habituation trials.

2.2. Results and discussion

Analyses confirmed that infants did habituate, as their average looking on the first three habituation trials was significantly greater than on the last three trials (16.50 s vs. 6.03 s, $p < .001$). Infants' looking times to the test events are depicted in Fig. 4b. Infants dishabituated more to Large Offset Test events than to Synchronous Test events, as confirmed by a mixed-design ANOVA with test event (Synchronous vs. Large Offset) as a within-subjects factor and presentation order as a between-subjects factor. This analysis revealed no effect of presentation order ($F < 1$), but a significant main effect of test-event type ($F(1,10) = 6.68$, $p = .027$): As predicted, infants looked reliably longer during Large Offset Test events (14.12 s) compared to Synchronous Test events (10.85 s). Analysis of the non-parametric data revealed that 10 out of 12 infants looked longer at Large Offset Test events than Synchronous Test events ($p = .019$, via a binomial test).

The results of this initial experiment are consistent with the interpretations that (1) infants, like adults, perceive both habituation events and Synchronous Test events in terms of causal launching, but perceive Large Offset Test events in terms of noncausal passing; and (2) infants dishabituated in part on the basis of the perceived causal status of the events.

3. Experiment 2: Postdiction in causal perception

Though the results of Experiment 1 were interpreted in terms of the events' categorical status as perceived launching vs. passing, there is another salient alternative which does not advert to causality at all: perhaps infants' looking times in Experiment 1 were only based on perceived temporal synchrony. In particular, infants may have looked longer at Large Offset Test events because of the offset, per se, along with its contrast to the no-offset habituation events. To properly test this alternative it is necessary to present infants with a full-overlap event involving a temporal offset between objects B and C that nevertheless yields the perception of causality. To do so in this experiment, we exploit the postdictive effect previously observed in adult subjects (Choi & Scholl, 2006b), where causal launching between A and B is perceived even when C starts moving up to 180 ms after B.

This experiment was thus identical to Experiment 1, except that the Large Offset Test condition was replaced with a Small Offset Test condition, where C began moving 120 ms after B (see Fig. 5a). This temporal offset still seems salient to adult observers (and in fact we show later that it is also salient to infants at this age), but in this experiment we predict that all events (both habituation and test) will be perceived in terms of causal launching, and therefore we do not expect any differential dishabituation.

3.1. Method

This experiment was identical to Experiment 1 except as noted here. Twelve new 7-month-old infants with a mean age of 6 months 29 days (range: 6 months 16 days to 7 months 12 days) participated. Two additional infants were tested but were excluded because of disinterest. The Small Offset Test event was identical to the Large Offset Test event from Experiment 1 except that C began moving 120 ms after B began moving. C then continued to move after B stopped until it was directly under B, at which point the display disappeared. This entire event lasted 5.36 s. Infants completed an average of 9 habituation trials.

3.2. Results and discussion

Analyses confirmed that infants did habituate, as their average looking on the first three habituation trials was significantly greater than on the last three trials (19.28 vs. 8.01 s, $p < .001$). The test event looking times from this experiment are depicted in Fig. 5b. In contrast to Experiment 1, infants did not differentially dishabituate to the two test event types, as confirmed by a mixed-design ANOVA with test event (Synchronous vs. Small Offset) as a within-subjects factor and presentation order as a between-subjects factor.

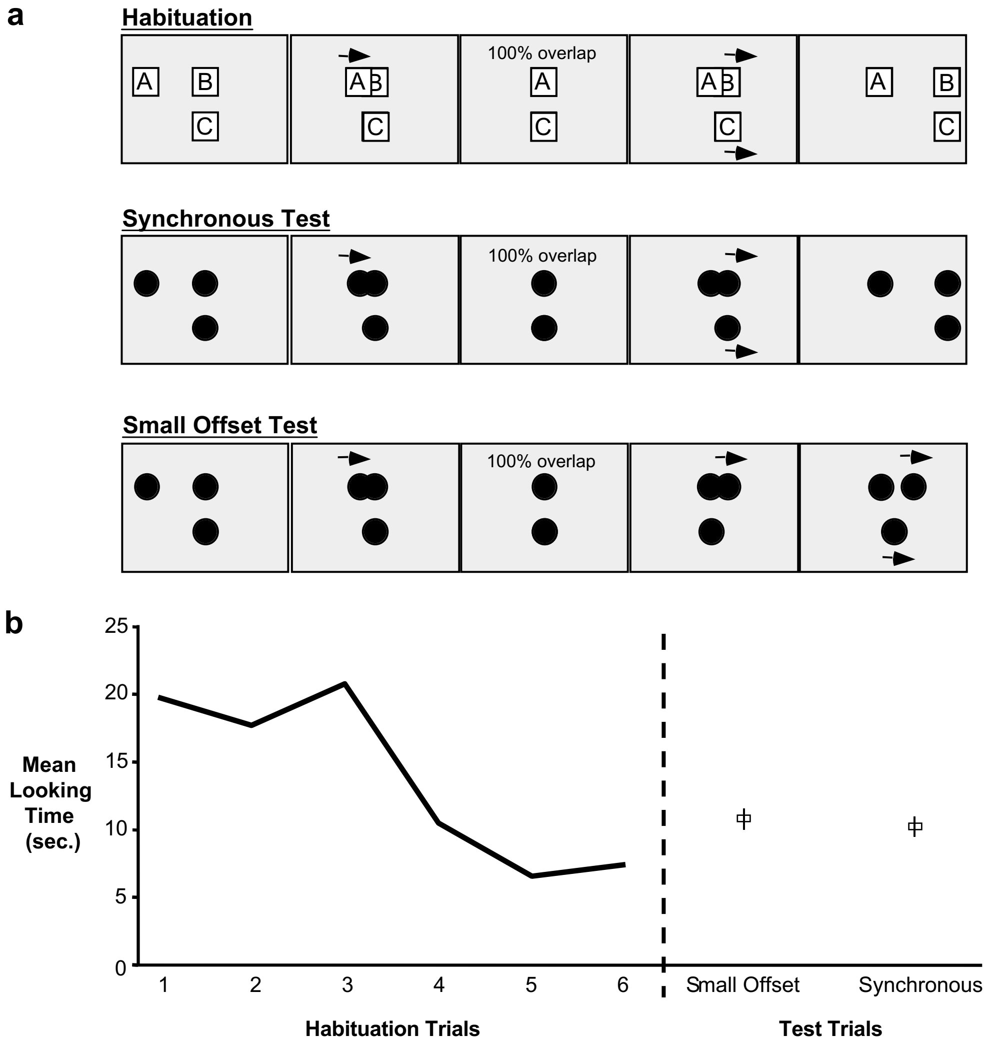

Fig. 5. (a) Depiction of the events presented to infants in Experiment 2. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion, and the letters did not appear in the actual displays. Infants were habituated to the same full-overlap display as in Experiment 1. After habituation criteria were met, infants were presented with two new test events featuring novel objects: a Synchronous Test event (identical to that in Experiment 1) and a Small Offset Test event in which the context object moved 120 ms after B began to move (perceived by adult observers as causal launching). (b) Duration of looking-time across Habituation and Test events in Experiment 2.

This analysis revealed no effect of either presentation order or test event type (both Fs < 1, both ps $> .5$): infants looked equally long at the Synchronous and Small Offset Test events (Ms $= 10.24$ and 10.54, respectively). Non-parametric data revealed that 6 of the 12 infants looked longer at the Small Offset Test event, while the other 6 infants looked longer at the Synchronous Test event ($p > .99$ via a binomial test).

We next compared test event looking times across Experiments 1 and 2, via a mixed-design ANOVA with experiment and presentation order as between-subjects factors and test event type (synchronous vs. offset) as a within-subjects factor. This analysis revealed a marginally significant interaction between experiment and test-event type ($F(1,22) = 4.15$, $p = .055$) The only other significant effect was a main effect of test event type ($F(1,20) = 5.94$, $p = .024$). This comparison indicates that the differential test looking times across these two experiments were driven by the nature of the temporal offset, since this was the only factor that differed between these two experiments.

The results of this experiment are consistent with an interpretation based on the perceived causal status of the events; namely, that infants (1) perceived the habituation events in both Experiments 1 and 2 in terms of causal launching; (2) perceived the Synchronous Test events in both experiments in terms of causal launching; (3) perceived the Large Offset Test event in Experiment 1 in terms of noncausal passing; but (4) perceived the Small Offset Test event in this experiment in terms of causal launching, via the operation of postdictive processing. In other words, if infants' looking time to these displays is driven by the perceived causal status of the events, then these results are exactly what would be expected if a causal module in infants employs the same rules of grouping and postdiction as in adults.

4. Experiment 3: The salience of small temporal offsets

Though the results of Experiments 1 and 2 were both interpreted in terms of the events' categorical status (i.e., perceived launching vs. passing), there is yet another alternative which does not advert to causality. This alternative, like that explored in Experiment 2, involves only a sensitivity to temporal offsets per se, but now does so with an added threshold assumption: perhaps infants in Experiment 1 were only looking differentially on the basis of the (600 ms) temporal offset, but infants in Experiment 2 did not do so because the shorter temporal offset (120 ms) was not salient enough to drive dishabituation. Indeed, we know relatively little about infants' motion sensitivity in this domain, and it is possible that infants did not even notice the shorter temporal offset from Experiment 2, in which case the Small Offset and Synchronous Test event would have seemed identical. To properly test this alternative it is necessary to assess more directly whether infants detect the shorter temporal offset from Experiment 2, and find it salient enough to drive dishabituation.

This experiment was thus identical to Experiment 2, except that the causal aspect of the displays was eliminated by removing object A from all animations (see Fig. 6a). In this case we predict that the smaller temporal offset will be sufficient on its own to drive dishabituation, when this difference is not overshadowed by the more salient causal interpretations that may arise in the full displays.

4.1. Method

This experiment was identical to Experiment 2 except as noted here. Twelve new 7-month-old infants with a mean age of 7 months 1 day (range: 6 months 21 days to 7 months 20 days) participated. Three additional infants were tested but were excluded because of fussiness or disinterest. Animations were identical to those used in Experiment 2, except that object A was removed from all displays during both habituation and test. Infants completed an average of 8 habituation trials.

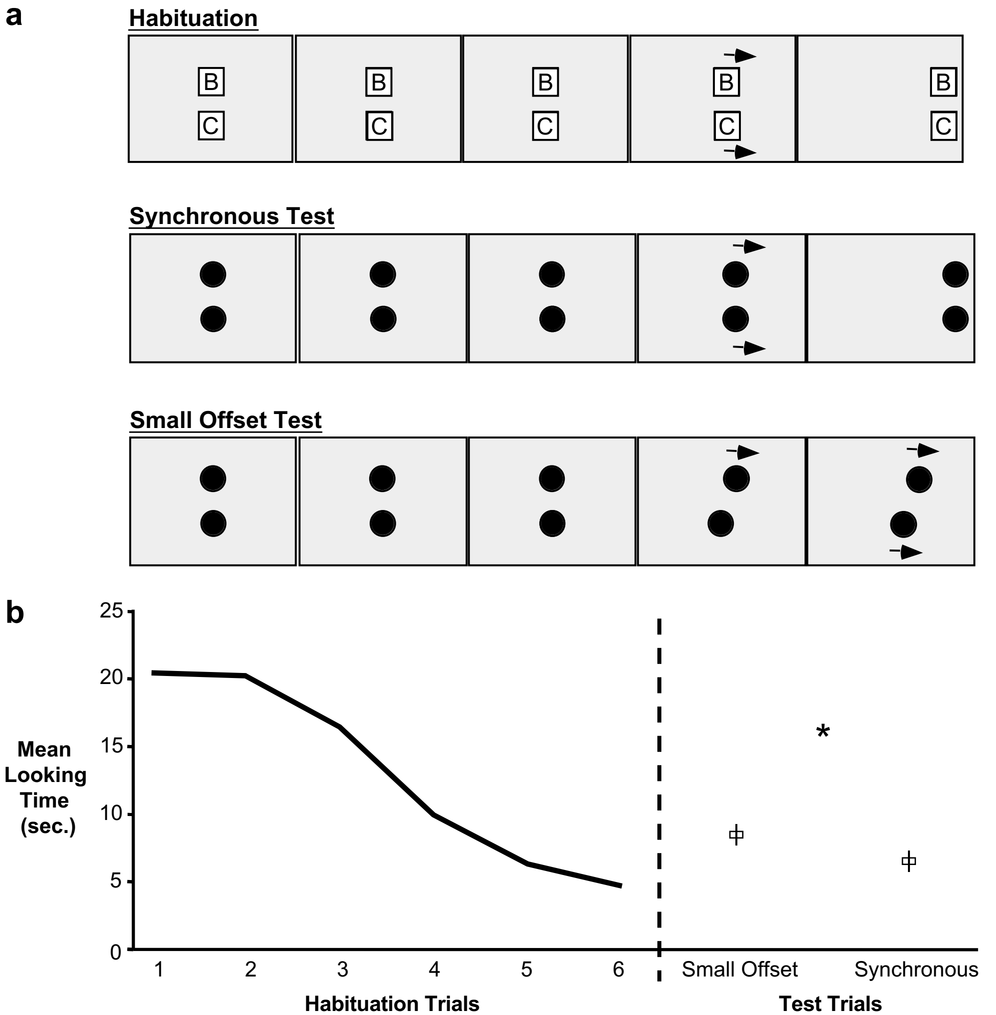

Fig. 6. (a) Depiction of the events presented to infants in Experiment 3. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion, and the letters did not appear in the actual displays. Infants were habituated to an event in which two objects were initially stationary and then began to move synchronously with one another. After habituation criteria were met, infants were presented with two new test events featuring novel objects: a Synchronous Test event in which the two objects continued to move synchronously (i.e., identical to the habituation) and a Small Offset Test event in which the lower object moved 120 ms after the top object began to move. (b) Duration of looking-time across Habituation and Test events in Experiment 3. The asterisk indicates a statistically significant difference in the degree of dishabituation ($p < .05$).

4.2. Results and discussion

Analyses confirmed that infants did habituate, as their average looking on the first three habituation trials was significantly greater than on the last three trials (18.67 vs.

6.85 s, $p < .001$). The test event looking times from this experiment are depicted in Fig. 6b. Inspection of this figure suggests that the relatively subtle difference of removing object A from the displays in this experiment (compared to the otherwise identical Experiment 2) made all the difference: now infants easily distinguished the two test event types, and dishabituated differentially based only on the short temporal offset. This impression was confirmed by submitting the data to a mixed-design ANOVA with test event (Synchronous vs. Small Offset) as a within-subjects factor and presentation order as a between-subjects factor. This analysis revealed no effect of presentation order ($F < 1$), but a significant main effect of test-event type ($F(1,10) = 10.27$, $p = .009$): as predicted, when salient causal interpretations were not present to drive infants' categorization and looking times, infants looked reliably longer during Small Offset Test events (8.44 s) compared to Synchronous Test events (6.45 s). Non-parametric data revealed that 10 out of 12 infants looked longer at Small Offset Test events than at Synchronous Test events ($p = .019$, via a binomial test).

We next compared test event looking times across Experiments 2 and 3, via a mixed-design ANOVA with experiment and presentation order as between-subjects factors and test event type (Synchronous vs. Small Offset) as a within-subjects factor. This analysis revealed a marginal interaction between experiment and test-event type ($F(1,20) = 3.13$, $p = .092$), and there was also a significant main effect of test event type ($F(1,20) = 5.67$, $p = .027$).

These results thus fail to support the deflationary interpretation of Experiments 1 and 2 wherein the differences (or lack thereof) are driven only by salient temporal offsets: in this experiment, the very same temporal offset used in Experiment 2 now had an effect, when more salient causal categorical interpretations were eliminated.

5. Experiment 4: The salience of small temporal offsets in more complex displays

The previous experiment demonstrated that infants dishabituate to changes in the small temporal offset per se, when the display is unambiguously noncausal—suggesting that the results of Experiment 2 reflect a postdictive effect on causal perception rather than a simple inability to perceive the small temporal delay. However, the display in Experiment 3 involved only two objects, whereas Experiment 2 involved three objects (compare Figs. 5 and 6a). It seems unlikely that this rather arbitrary difference between the displays was responsible for the difference in looking time patterns, but we wanted to ensure that this was not the case. Thus, this experiment replicated Experiment 3 while adding a third object back into the display, but in a way which did not support a causal interpretation.

5.1. Method

This experiment was identical to Experiment 2 except as noted here. Twelve new 7-month-old infants with a mean age of 7 months 2 days (range: 6 months 15 days to 7 months 30 days) participated. Three additional infants were tested but were excluded because of fussiness. The only difference in the habituation and test event animations compared to Experiment 2 was that object A traveled on a plane above B, so that their nearest edges were vertically separated by 3.18° (see Fig. 7a). Infants completed an average of 9 habituation trials.

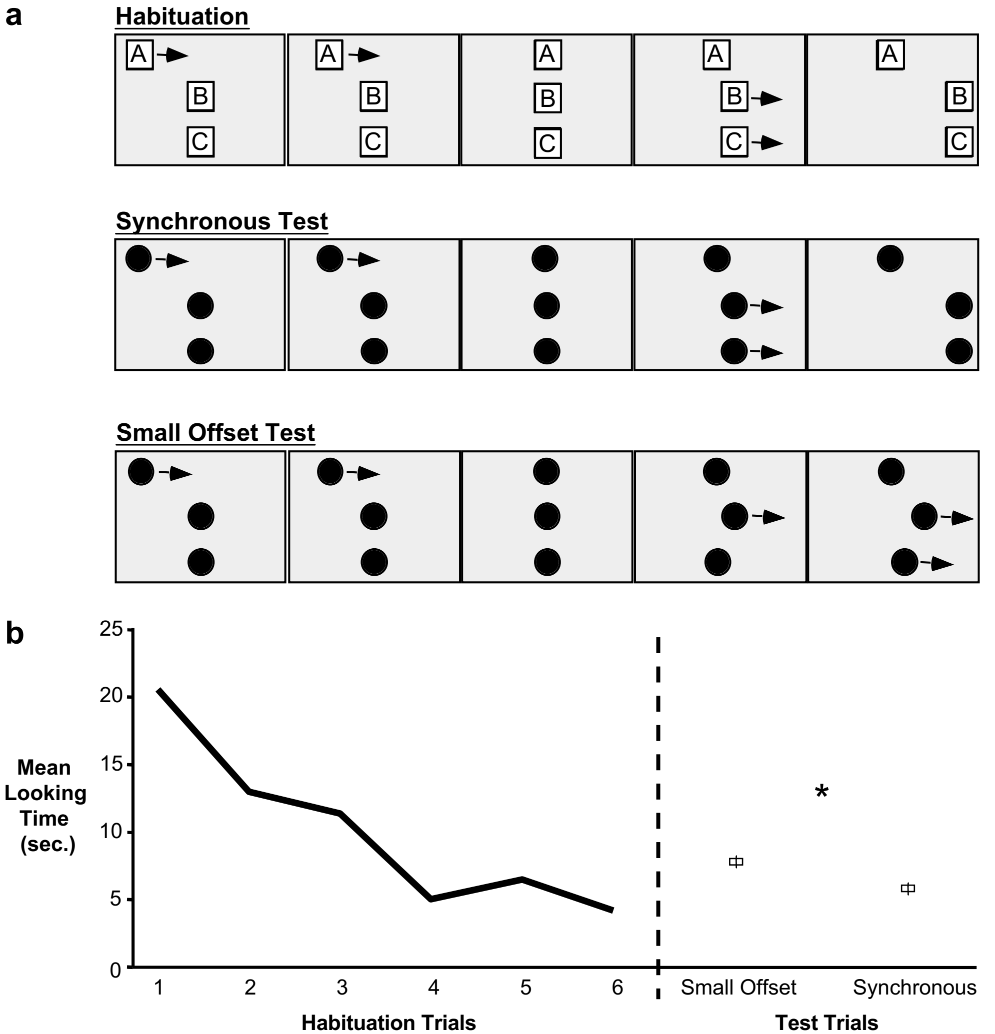

Fig. 7. (a) Depiction of the events presented to infants in Experiment 4. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion, and the letters did not appear in the actual displays. Infants were habituated to an event in which three objects were initially stationary and then began to move synchronously with one another. After habituation criteria were met, infants were presented with two new test events featuring novel objects: a Synchronous Test event in which the objects continued to move in a synchronous fashion (i.e., identical to the habituation) and a Small Offset Test event in which the bottom object began to move 120 ms after the middle object began moving. (b) Duration of looking-time across Habituation and Test events in Experiment 4. The asterisk indicates a statistically significant difference in the degree of dishabituation ($p < .05$).

5.2. Results and discussion

Analyses confirmed that infants did habituate, as their average looking on the first three habituation trials was significantly greater than on the last three trials (14.83 vs. 5.52 s,

$p < .001$). The test event looking times from this experiment are depicted in Fig. 7b. Consistent with the findings of Experiment 3, inspection of this figure suggests that even with the inclusion of object A (now moving in a different spatial plane) infants still distinguished the two test event types, and dishabituated differentially based only on the short temporal offset. This analysis revealed no effect of presentation order ($p > .14$), but a significant main effect of test-event type ($F(1,10) = 10.27$, $p = .009$): as predicted, when salient causal interpretations were not present to drive infants' categorization and looking times, infants looked reliably longer during Small Offset Test events (7.95 s) compared to Synchronous Test events (5.94 s). Non-parametric data revealed that 11 out of 12 infants looked longer at Small Offset Test events than at Synchronous Test events ($p < .01$, via a binomial test).

We next compared test event looking times across Experiments 2 and 4, via a repeated-measures ANOVA with experiment and presentation order as between-subjects factors and test event type (Synchronous vs. Small Offset) as a within-subjects factor. This analysis revealed a marginal interaction between experiment and test-event type ($F(1,20) = 3.42$, $p = .079$), and there was also a significant main effect of test event type ($F(1,20) = 6.17$, $p = .022$). Since these two experiments employed the same numbers of objects and motions yet yielded different results, this suggests that perceived causality was again the culprit, rather than temporal offsets per se.

The results of this experiment again support our interpretation of the results in terms of dishabituation driven by perceived causal status rather than sensitivity to temporal offsets. The alternative hypothesis (that differences in display complexity may have been responsible for the observed differences between Experiments 2 and 3) was not supported, since we find that even when display complexity is better equated with Experiment 2, infants remain sensitive to the Small Temporal Offset in a unambiguously noncausal display.

6. Experiment 5: The salience (or lack thereof) of large temporal offsets

The previous experiments argued against the possibility of interpreting the results on the basis perceived temporal offsets alone (i.e., without appeal to perceived causal status)—even with the post hoc assumption that the smaller temporal offsets were not salient enough by themselves to drive dishabituation. Can this more direct and deflationary interpretation still be salvaged? As is nearly always the case in such situations, it can, by again adding yet another post hoc assumption. In particular, suppose we retain the previous assumptions from Experiments 3 and 4, that infants will dishabituate on the basis of perceived temporal offsets, but we assume (1) that this effect is subject to a threshold cutoff, and (2) that this cutoff divides the two temporal offsets we used in these studies (120 vs. 600 ms). The challenge is then to explain why infants dishabituated on the basis of a 120 ms offset in Experiments 3 and 4 but not Experiment 2. Again, the logic here is straightforward: we can simply appeal to the visual character of the display, rather than to its categorical status. Suppose (1) that infants dishabituate only on the basis of perceived temporal offsets, and not due to perceived causality; (2) that they only find a certain range of temporal offsets salient—including 600 ms, but not including 120 ms; but (3) that this threshold itself moves based on the complexity of the display; and (4) that the behavior of object A makes all the difference here: when A interacts with another object in the display (unlike in Experiments 3 and 4), infants do not perceive the 120 ms offset or do not find it salient (as in Experiment 2). To properly test this alternative it is necessary to employ a temporal offset which we already know will drive dishabituation even in displays wherein A interacts with other objects—such as the 600 ms offset from Experiment 1—but to now show the impotence of this factor when perceived causality does not vary between test event types.

This experiment was thus identical to Experiment 1, except that the causal aspect of the displays was held constant in an especially direct way: in all test events in this experiment, object A stopped and object B began moving as soon as they became adjacent (rather than fully overlapped). This difference may be visually subtle (see Fig. 8a), but it essentially fixes adults' perception of causal launching such that no temporal offset can yield perceived passing. In other words, object C was initially added to the display to help disambiguate the full-overlap event, but the launching event used in this experiment is already unambiguous, rendering C's motion causally irrelevant. Critically, however, all three objects in these displays are still present and moving, and object A again interacts with another object (unlike Experiments 3 and 4). Thus, despite the fact that a short temporal offset (120 ms) is sufficient to drive dishabituation in a display that lacks any causal interpretation (in Experiments 3 and 4), for a display that is causal in nature and is impervious to grouping effects, we predict that even a large temporal offset (600 ms) will not drive dishabituation since a causal interpretation of the display should trump any lower-level differences (i.e., temporal asynchrony). However, if in fact temporal offsets alone are driving the effects (modulated by the magnitude of the offset and by the complexity or number of objects in the displays), the results of this experiment should mirror those of Experiment 1. Thus, this experiment contrasts a difference in spatiotemporal factors with no difference in causal status and therefore offers a direct test of the hypothesis that spatiotemporal factors, and not changes in perceived causality, drive infants' dishabituation in these displays.

6.1. Method

This experiment was identical to Experiment 1 except as noted here. Twelve new 7-month-old infants with a mean age of 6 months 28 days (range: 6 months 17 days to 7 months 12 days) participated. Two additional infants were tested but were excluded because of fussiness. The only difference in the test event animations compared to Experiment 1 was that object A always stopped and B began moving as soon as they became adjacent. Infants completed an average of 9 habituation trials.

6.2. Results and discussion

Analyses confirmed that infants did habituate, as their average looking on the first three habituation trials was significantly greater than on the last three trials (12.29 vs. 5.39 s, $p < .001$). It initially appeared that the amount of looking in the first trial of habituation was less than in previous experiments, but in fact a direct statistical comparison to Experiments 1 and 2 revealed that this was not the case (ps = .41 and .27, respectively).

The test event looking times from this experiment are depicted in Fig. 8b. In contrast to Experiment 1, infants did not differentially dishabituate to the two test event types. A mixed-design ANOVA with test event (Synchronous vs. Large Offset) as a within-subjects factor and presentation order as a between-subjects factor revealed no effect of either presentation order or test event type (both Fs < 1, both ps $> .2$): in this unambiguously causal

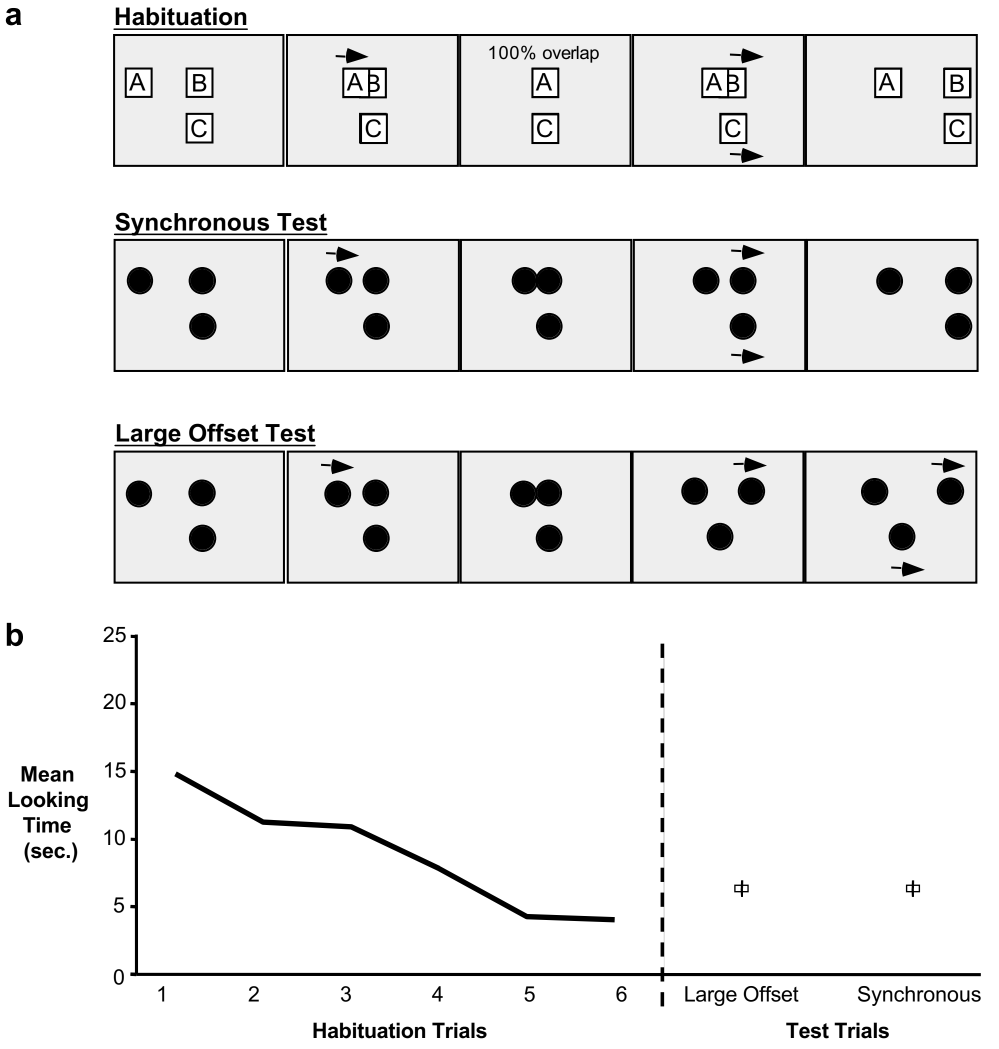

Fig. 8. (a) Depiction of the events presented to infants in Experiment 5. These events are presented as a sequence of static frames, with time increasing toward the right. Arrows indicate motion, and the letters did not appear in the actual displays. Infants were habituated to the same full-overlap display from Experiment 1. After habituation criteria were met, infants were presented with two new test events featuring novel objects: a Synchronous Test event (identical to that in Experiment 1), and a Large Offset Test event wherein A is seen to collide with B, causing it to move while the context object moved 600 ms after B began to move. (b) Duration of looking-time across Habituation and Test events in Experiment 2. Asterisks indicate statistically significant dishabituation ($p < .05$).

context infants looked equally long at the Synchronous and Large Offset Test events (Ms = 6.44 and 6.67, respectively). Non-parametric data revealed that 6 of the 12 infants looked longer the Large Offset Test event, while the other 6 infants looked longer at the Synchronous Test event ($p > .99$ via a binomial test).

We next compared test event looking times across Experiments 1 and 5, via a repeated-measures ANOVA with experiment and presentation order as between-subjects factors and test event type (Synchronous vs. Large Offset) as a within-subjects factor. This analysis revealed a significant interaction between experiment and test-event type ($F(1,20) = 5.30$, $p = .032$); there was also a significant main effect of test event type ($F(1,20) = 7.00$, $p = .015$) and a main effect of experiment ($F(1,20) = 5.16$, $p = .034$). Since these two experiments employed the same numbers of objects and motions yet yielded opposite results, this suggests that perceived causality drove infants' dishabituation, rather than temporal offsets per se.

The results of this experiment again support our interpretation of the results in terms of dishabituation driven by perceived causal status rather than temporal offsets. This alternative hypothesis cannot explain our results, even assuming a complex salience threshold for temporal offsets which divides the durations that we used in this study (ruled out in Experiments 3 and 4), and even assuming that this salience threshold is higher for interactive displays (ruled out here, since we already know from Experiment 1 that the 600 ms offset can drive dishabituation in such displays). Thus, the results of all of the experiments presented here can be better explained by appeal to the perceived causal status of the displays.

7. General discussion

The results of the five experiments presented here can be readily summarized by noting that infants appear to perceive various subtly different displays in terms of causality (or not) in the same way that adults do, and that these differences in perceived causality—rather than lower-level spatiotemporal factors—appear to drive dishabituation. In particular, our results cannot be explained by appeal to sensitivity to temporal offsets, since neither large nor small temporal offsets drive dishabituation when perceived causality does not vary. In the remainder of this paper, we suggest that these results have interesting implications for the nature of causal perception in infancy, the similarity between causal perception in infants and adults, the origins of perceived causality, and the ubiquity of postdictive processing.

7.1. But is it really causality?

Perhaps the dominant challenge in developmental work on the origins of causal perception has been to isolate perceived causality, per se, beyond correlated lower-level spatiotemporal factors. This has been most successfully done in Leslie's studies of infants' causal perception, using the 'reversal logic' described in Section 1.2. In the current project, we attempted to demonstrate causal perception in 7-month-old infants using a very different strategy. So how well did we isolate causal perception? In particular, would it be possible to craft an account that could explain our results without appeal to causal perception at all?

Certainly this is possible. Indeed, to do so one can pick up from where Experiment 5 left off, and employ the same strategy used throughout the paper to extend the alternative interpretation based on a sensitivity to temporal gaps: identify the difference between the most recent pair of contrasting experiments (e.g., Experiment 1 and 5), and then reinterpret that difference in terms of whatever lower-level factor was used to influence the perceived causal status of the events. Experiments 1 and 5, for example, differed in their perceived causality for adults, but they also differed in a lower-level way: objects A and B fully overlapped in Experiment 1 but not Experiment 5. So maybe that matters: maybe subjects are responding only to temporal offsets, but their sensitivity to these offsets is also impacted by whether two objects overlapped. You might think that overlapping is intrinsically more ambiguous or difficult to process, but in fact you would have to predict the opposite relationship: maybe the overlapping displays in Experiment 1 were easier to process because of their smaller total spatial extent, for example, which made the large temporal offset more salient—whereas with the non-overlapping displays in Experiment 5, the greater spatial extent was harder to process, and left too little resources available for processing the same temporal offset.

This type of explanation does fit the data: it explains all of the patterns of dishabituation in each of our five experiments without appeal to causal perception, per se. But note that in doing so it demands that we assume:

(1) that the ability of infants to detect temporal offsets (and/or to find them salient enough to drive dishabituation) is limited by their magnitude;
(2) that the critical threshold beyond which such offsets are or are not salient happens to lie in between the values of 120 ms and 600 ms tested in our studies;
(3) that this 'offset detection threshold' is modulated by the complexity of the display;
(4) that this degree of modulation by complexity happens to span the rather minimal difference between our Experiment 2 (which involved the extra object A to drive the grouping effect) and Experiment 3 (in which A was removed)—and also the even more minimal difference between Experiment 2 and Experiment 4 (in which A was present but did not interact with any other objects);
(5) that this threshold is further modulated by the degree to which two objects overlap in a display;
(6) that the direction of this 'overlap' effect is such to make temporal offsets easier to notice when objects overlap, compared to when they do not overlap;
(7) that the magnitude of modulation by this 'overlap' effect is strong enough to overpower the salience of the 600 ms temporal offset used in Experiments 1 and 5.

Using this set of 7 assumptions (or another equally complex set), one could explain the results from this paper without appeal to causal perception. But note also that: (a) you need every one of these assumptions; (b) none of them has independent supporting evidence, to our knowledge; and (c) several of them seem implausible. Applying the principle of Ockham's Razor, we conclude that it is more compelling to interpret these results in terms of the hypothesized differences in perceived causal status between these events—which, far from being arbitrary, all have independent evidence from very different types of experiments and dependent measures with adults. Certainly this view—that infants perceive causality in displays just as adults do (as predicted by the modular theory)—is a more convincing explanation of the results, if only because it has independent motivation, and requires considerably fewer post hoc assumptions.

We submit that this project provides a new type of converging evidence for causal perception in infancy that differs importantly from previous strategies. This puts the burden of proof back on opponents of this conclusion to support the required post hoc assumptions.

7.2. Innateness?

As noted in the introduction, much of the debate about the origins of causal perception has focused on the possibility of innate origins for such processing. The motivation for such views is largely theoretical: causality seems to be a difficult concept to acquire, and so having it embodied in an innate module might both (1) help to drive early identifications of this ubiquitous (and ubiquitously-important) property of the world, and (2) help to jumpstart the further development of causal cognition. Our results, of course, do not directly add any new constraints to this debate, beyond converging evidence that such abilities are in place by 7 months (which at least puts an upper bound on the degree of experience which is required to give rise to causal perception).

However, we think that the specificity of the effects observed in this project is perhaps relevant to this debate, via the following logic: the more subtle and nuanced the factors involved in such processing during infancy, the less likely that they were directly learned from a limited degree of environmental experience. Here we observed effects of grouping involving multiple objects and events, and postdictive effects involving sensitivity to subtle differences in the magnitude of temporal offsets between multiple objects' movements. The accumulation of such details to the underlying sensitivity during infancy, to our minds, makes it increasingly unlikely that sensitivity to these precise factors was directly learned by the individual infants in our experiments.

At a minimum, such details pose a greater challenge to those models that would account for causal perception without any inbuilt assumptions of innate processing principles driven by phylogenetic learning. This point can be made concrete: Chaput and Cohen (2001), for example, have proposed a learning model of the origin of causal perception which alleges to account for previous developmental results such as those of Leslie (1984) without any specific assumptions beyond "domain-general information processing principles". The power and plausibility of such a model, however, is only as great as the scope of effects it aims to account for. And here, again, this model (like most previous developmental work) is limited only to studies of spatial and temporal gaps. But we note in this context (1) that this model cannot even begin to account for any of our results, without building in an entirely new set of post hoc assumptions, and (2) that even with such assumptions it is far from clear how such a model might work. As such, the current project provides a case study of how to explore the specificity of principles of causal perception in infancy, which in turn can have important implications for the plausibility of models that attempt to account for this domain without any domain-specific assumptions.

7.3. Postdiction

Our focus throughout this paper has been on causal perception, and the postdictive effect that we exploited in these studies was simply a convenient way to explore the nature and specificity of such processing in infancy. However, it is worth noting the interest of demonstrating postdictive processing in infancy for independent reasons. The existence of postdiction in visual processing was initially highlighted only in very particular circumstances, mostly involving apparent motion (e.g., Ramachandran & Anstis, 1986). More recently, vision scientists have argued that such processing is a more general and ubiquitous feature of visual processing. These explorations of the breadth of postdictive processing have proceeded largely by adding to the catalog of phenomena where such processing seems to apply—e.g., including the flash-lag effect (Eagleman & Sejnowski, 2000), illusory line motion (Eagleman & Sejnowski, 2003), motion-induced blindness (Mitroff & Scholl, 2004), causal perception (Choi & Scholl, 2006b), and object reviewing (Gao & Scholl, 2007). But another way to explore the breadth of postdictive processing is to explore its developmental origin: are such effects gradually learned in adulthood, or do they characterize visual processing in young minds as well? To our knowledge, our results provide the first demonstration of postdictive processing in infancy, via the logic rehearsed in Section 1.4: in Experiment 2, an event (the onset of object C's motion) that occurred after objects A and B had fully separated was still able to influence which object the motion was bound to—though it should have been too 'late' to do so by that point. This in turn affected the perceived causal status of the otherwise-ambiguous event, leading infants to perceive causality 'after the fact'.

8. Conclusions: The nature of causal perception in infancy

In the English version of Michotte's original book on causal perception (1946/1963), he presented more than 100 experiments exploring dozens of subtle manipulations of causal launching and similar effects—and contemporary research has added to this catalog of 'rules' of causal perception along many dimensions (see Wagemans et al., 2006). However, very few of these manipulations have been represented in developmental work. Rather, nearly all infant studies of causal perception to date have been focused on contrasts with spatial and temporal gaps. In some ways this is understandable: these contrasts are among the most famous of Michotte's manipulations for good reason, since they collectively control for many extrinsic factors. However, while studies of launching vs. gaps can be tremendously useful in determining *that* infants perceive causality (as in Leslie, 1982, 1984; Leslie & Keeble, 1987), they can only do relatively little by way of explaining *how* causal perception works in infancy.[6]

Thus, a primary goal of this study was to explore entirely different aspects of causal perception in infancy, using manipulations and phenomena from the most recent published study of adults' causal perception that we were aware of when beginning this project (Choi & Scholl, 2006b). The results suggest that infants' causal perception, like adults', is influenced by both grouping and postdictive processing. Beyond implications for debates about innateness (as explored in Section 7.2), or the importance of postdiction itself (as explored in Section 7.3), these results are important for two more general reasons, both connected to the modular view of causal perception. First, it is notable that the relatively nuanced factors studied here—involving interactions beyond single events, and involving subtly dif-

[6] Moreover, despite their popularity, contrasts with spatial and temporal gaps may not always be ideal for more specific reasons. Some perceptual processes, in particular, may be sensitive to a combination of the continuity of motion in an event, and the number of objects in that event. This contrast has long been salient in work on causal perception, of course: in the contrast between causal launching and either spatial or temporal gaps, only launching involves two objects but a single continuous motion. This same issue has also been theoretically salient: indeed, Michotte (1946/1963) suggested that causal perception arose in the first place because of a conflict between two perceived objects but one perceived motion. It has recently been suggested, though, that the conjunction of these factors can explain some effects (e.g. involving representational momentum) even when causal perception is not involved (Choi & Scholl, 2006a)—and these authors further suggest that this conjunction could also explain some past infancy results without appeal to perceived causality. However, these potential problems do not apply to the current experiments.

ferent temporal offsets—have effects at all in infancy. Recall that one of Michotte's key arguments for the perceptual nature of "causal perception" was its relatively strict dependence on subtle display factors (e.g., Scholl & Tremoulet, 2000). This theme has been well explored in recent adult work on causal perception, but has never before been studied in infancy, to our knowledge. The present results suggest that infants' causal perception is also driven by nuanced rules that are highly sensitive to subtle differences in the visual input. We submit, using Michotte's original logic, that this discovery supports the interpretation of such processing in terms of true causal perception in infancy, and in more modern parlance, that this is consistent with the idea of a modular basis for such processing.

Of course, the results of these experiments go beyond demonstrating that infants' causal perception is driven by subtle bottom-up visual factors; it also suggests that these factors are the same ones that operate in adult visual cognition. This can be seen as an independent prediction of the modular view of causal perception: if a causal-perception module is in fact identical in infants and adults (which, with appropriate caveats regarding triggering and performance limitations, is what we would expect), then infants' causal perception should be affected by the same types of features that drive causal perception in adults—including the effects of grouping and postdiction studied here.

Most generally, however, perhaps the most important implication of these results lies not in their support for abstract themes of modularity, but in their demonstration of how work inspired by visual cognition studies in adults can be fruitfully employed to begin determining just how and when causal percepts are and are not generated in infancy.

References

Ball, W. A. (1973). *The perception if causality in the infant*. Paper presented at the biannual meeting of the Society for Research in Child Development, Philadelphia, PA.

Boyle, D. (1960). A contribution to the study of phenomenal causation. *Quarterly Journal of Experimental Psychology, 12*, 171–179.

Chaput, H. H., & Cohen, L. B. (2001). A model of infant causal perception and its development. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the twenty-third annual conference of the cognitive science society* (pp. 182–187). Mahwah: Lawrence Erlbaum Associates.

Choi, H., & Scholl, B. J. (2004). Effects of grouping and attention on the perception of causality. *Perception & Psychophysics, 66*, 926–942.

Choi, H., & Scholl, B. J. (2006a). Measuring causal perception: Links to representational momentum? *Acta Psychologica, 123*, 91–111.

Choi, H., & Scholl, B. J. (2006b). Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception. *Perception, 35*, 385–399.

Cohen, L. B. (1998). An information-processing approach to infant perception and cognition. In F. Simion & G. Butterworth (Eds.), *The development of sensory, motor, and cognitive capacities in early infancy* (pp. 277–300). East Sussex: Psychology Press.

Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of causality. *Infant Behavior and Development, 21*, 713–731.

Cohen, L. B., & Oakes, L. M. (1993). How infants perceive a simple causal event. *Developmental Psychology, 29*, 421–433.

Cohen, L. B., Amsel, G., Redford, M. A., & Casasola, M. (1998). The development of infant causal perception. In A. Slater (Ed.), *Perceptual development: Visual, auditory, and speech perception in infancy* (pp. 167–209). East Sussex, UK: Psychology Press Ltd.

Costall, A. (1991). Phenomenal causality. In G. Thines, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception* (pp. 51–64). Hove, England: Erlbaum.

Dennett, D. C. (1991). *Consciousness explained.* Little-Brown.

Eagleman, D. M., & Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science, 287*, 2036–2038.

Eagleman, D. M., & Sejnowski, T. J. (2003). The line-motion illusion can be reversed by motion signals after the line disappears. *Perception, 32*, 963–968.

Gao, T., & Scholl, B. (2007). Are objects required for object-files? [Abstract]. *Journal of Vision, 7*, 916a.

Gordon, I., Day, R., & Stecher, E. (1990). Perceived causality occurs with stroboscopic movement of one or both stimulus elements. *Perception, 19*, 17–20.

Guski, R., & Troje, N. (2003). Audiovisual phenomenal causality. *Perception & Psychophysics, 65*, 789–800.

Hubbard, T. L., & Ruppel, S. E. (2002). A possible role of naïve impetus in Michotte's "launching effect": Evidence from representational momentum. *Visual Cognition, 9*, 153–176.

Hume, D. (1977). *An enquiry concerning human understanding*. Indianapolis: Hackett.

Kotovsky, L., & Baillargeon, R. (1994). Calibration-based reasoning about collision events in 11-month-old infants. *Cognition, 51*, 107–129.

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

Kotovsky, L., & Baillargeon, R. (2000). Reasoning about collisions involving inert objects in 7.5-month-old infants. *Developmental Science, 3*, 344–359.

Kosugi, D., Ishida, H., & Fujita, K. (2003). 10-month-old infants' inference of invisible agent: Distinction in causality between object motion and human action. *Japanese Psychological Research, 45*, 15–24.

Kruschke, J. K., & Fragassi, M. M. (1996). The perception of causality: Feature binding in interacting objects. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 441–446). Hillsdale, NJ: Erlbaum.

Leslie, A. M. (1982). The perception of causality in infants. *Perception, 11*, 173–186.

Leslie, A. M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception, 13*, 287–305.

Leslie, A. M. (1986). Getting development off the ground: Modularity and the infant's perception of causality. In P. van Geert (Ed.), *Theory building in development* (pp. 405–437). Amsterdam: Elsevier, North-Holland.

Leslie, A. M. (1988). The necessity of illusion: Perception and thought in infancy. In L. Weiskrantz (Ed.), *Thought without language* (pp. 185–210). Oxford: Clarendon Press/Oxford University Press.

Leslie, A. M. (1994). ToMM, ToBy, and agency: Core architecture and domain specificity. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 119–148). New York: Cambridge University Press.

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition, 25*, 265–288.

Michotte, A. (1946/1963). *La perception de la causalite*. Louvain: Institut Superior de Philosophie, 1946. English translation of updated edition by T. Miles & E. Miles, *The perception of causality*, Basic Books, 1963.

Michotte, A., & Thinés, G. (1963/1991). La causalité perceptive. *Journal de psychologie normale et pathologique, 60*, 9–36, 1963. Translated and reprinted in Thinès, G., Costall, A., & Butterworth, G. (Eds.), *Michotte's experimental phenomenology of perception* (pp. 66–87). Hillsdale, NJ: Erlbaum, 1991.

Mitroff, S. R., & Scholl, B. J. (2004). Seeing the disappearance of unseen objects. *Perception, 33*, 1267–1273.

Most, S. B., Scholl, B. J., Clifford, E., & Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review, 112*, 217–242.

Natsoulas, T. (1961). Principles of momentum and kinetic energy in the perception of causality. *American Journal of Psychology, 74*, 394–402.

Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development, 5*, 193–207.

Oakes, L. M. (1994). The development of infants' use of continuity cues in their perception of causality. *Developmental Psychology, 30*, 869–879.

Olum, V. (1956). Developmental differences in the perception of causality. *American Journal of Psychology, 69*, 417–423.

Piaget, J. (1936/1963). *The origins of intelligence in children*. New York: W.W. Norton & Company, Inc.

Ramachandran, V. S., & Anstis, S. M. (1986). The perception of apparent motion. *Scientific American, 254*, 102–109.

Saxe, R., & Carey, S. (2006). The perception of causality in infancy. *Acta Psychologica, 123*, 144–165.

Schlottmann, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. *Developmental Psychology, 35*, 303–317.

Schlottmann, A. (2000). Is perception of causality modular? *Trends in Cognitive Science, 4*, 441–442.

Schlottmann, A., & Anderson, N. H. (1993). An information integration approach to phenomenal causality. *Memory & Cognition, 21*, 201–785.

Schlottmann, A., & Shanks, D. (1992). Evidence for a distance between judged and perceived causality. *Quarterly Journal of Experimental Psychology, 44A*, 321–342.

Schlottmann, A., Allen, D., Linderoth, C., & Hesket, S. (2002). Perceptual causality in children. *Child Development, 73*, 1656–1677.

Schlottmann, A., Ray, E., Mitchell, A., & Demetriou, N. (2006). Perceived social and physical causality in animated motions: Spontaneous reports and ratings. *Acta Psychologica, 123*, 112–143.

Scholl, B. J., & Leslie, A. M. (1999). Modularity, development, and 'Theory of Mind'. *Mind & Language, 14*, 131–153.

Scholl, B. J., & Leslie, A. M. (2001). Minds, modules, and meta-analysis. *Child Development, 72*, 696–701.

Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science, 13*, 493–498.

Scholl, B. J., & Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception, 33*, 455–469.

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences, 4*, 299–309.

Spelke, E. S., Phillips, A., & Woodward, A. L. (1995). Infants' knowledge of object motion and human action. In D. Sperber, D. Premack, & A. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 44–78). Oxford: Clarendon Press.

Sperber, D., Premack, D., & Premack, A. (1995). *Causal cognition: A multidisciplinary debate*. Oxford: Oxford University Press.

Wagemans, J., Van Lier, R., & Scholl, B. J. (2006). Introduction to Michotte's heritage in perception and cognition research. *Acta Psychologica, 123*, 1–19.

Wang, S., Kaufman, L., & Baillargeon, R. (2003). Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events. *Infant Behavior & Development, 26*, 529–567.

Weir, S. (1978). The perception of motion: Michotte revisited. *Perception, 7*, 247–260.

White, P. A. (1988). Causal processing: Origins and development. *Psychological Bulletin, 104*, 36–52.

White, P. A. (1995). *The understanding of causation and the production of action*. Hillside, NJ: Erlbaum.

White, P. A. (2005a). Visual causal impressions in the perception of several moving objects. *Visual Cognition, 12*, 395–404.

White, P. A. (2005b). Visual impressions of interactions between objects when the causal object does not move. *Perception, 34*, 491–500.

White, P. A. (2006). The role of activity in visual impressions of causality. *Acta Psychologica, 123*, 166–185.

White, P. A., & Milne, A. (1997). Phenomenal causality: Impressions of pulling in the visual perception of objects in motion. *American Journal of Psychology, 110*, 573–602.

White, P., & Milne, A. (1999). Impressions of enforced disintegration and bursting in the perception of collision events. *Journal of Experimental Psychology: General, 128*, 499–516.

Yela, M. (1952). Phenomenal causation at a distance. *Quarterly Journal of Experimental Psychology, 4*, 139–154.

Young, M. E., Rogers, E. T., & Bechmann, J. S. (2005). Predicting when, not just whether. *Memory & Cognition, 33*, 320–331.