Early-Developing Causal Perception is Sensitive to Multiple Physical Constraints

Jonathan F. Kominsky (jkominsky@g.harvard.edu) &
Susan Carey (scarey@wjh.harvard.edu)
Harvard University Department of Psychology, William James Hall # 1154, 33 Kirkland St.
Cambridge, MA 02138 USA

Abstract

If an object A moves until it is adjacent with a stationary object B, at which point object A stops and object B begins moving, adults and infants 6 months of age and older perceive that A caused B to move. These "launching" events correspond to real-world collisions, which are governed by Newtonian mechanics. Previous work showed that infants were sensitive to Newtonian constraints on relative speed. Here, we show that infant causal perception is sensitive to other physical constraints on collision events as well. Infants habituated to a launching event will dishabituate to an event in which object B moves at a 90° angle relative to object A, but not to a rotated version of the launching event. This selective dishabituation was not found for non-causal events. The results suggest that early-developing causal perception is sensitive to the many physical principles of real-world collision events.

Keywords: Causal perception, naïve physics, cognitive development, infant

Introduction

Consider an event involving two objects. One object, call it A, moves toward a second, stationary object B, until they are adjacent, at which point A immediately stops moving and B immediately begins moving in the same direction and at the same speed. This "launching" event (rendered schematically in Fig. 1a) is viscerally and irresistibly perceived as A causing B to move (Michotte, 1963). (Animated events can be found at http://www.jfkominsky.com/demos.html, and readers are strongly encouraged to view it in order to experience the phenomenology for themselves.)

This launching event is the prototypical example of causal perception, the automatic detection of cause and effect relationships within certain types of events (Hubbard, 2013; Michotte, 1963; Scholl & Tremoulet, 2000). Causal perception is fundamentally different from causal learning or reasoning. While there have been arguments that such events merely lead us to make a cognitive inference of causality based on extensive experience (Hume, 1902; Piaget & Garcia, 1977; Rips, 2011), recent work has provided strong evidence that causal perception is truly a product of "low-level" automatic perceptual processing. For example, causal perception can influence the perception of apparent motion (Kim, Feldman, & Singh, 2013), causal events enter awareness faster than non-causal events (Moors, Wagemans, & de-Wit, 2017), and causal perception is subject to the uniquely perceptual phenomenon of retinotopically-specific visual adaptation (Karaminis et al., 2015; Rolfs, Dambacher, & Cavanagh, 2013).

Furthermore, causal perception is early-developing, robustly present in human infants by six months of age (Cohen & Amsel, 1998; Leslie & Keeble, 1987; Newman, Choi, Wynn, & Scholl, 2008). The typical finding is that infants who are habituated to a launching event in which A launches B will dishabituate to an event in which the causal roles are reversed, i.e. B launches A (Leslie & Keeble, 1987).

Controlled rearing experiments have even suggested that causal perception may even be innate in some species. Chicks raised from birth in an environment where their only visual input was a launching event will later spend more time with object A than object B (Mascalzoni, Regolin, & Vallortigara, 2010). Such selective imprinting does not occur when object A's motion onset is hidden (and therefore may or may not be self-generated), suggesting that chicks are responding to the *spontaneous* motion onset as an indicator of which object could be a caregiver. It is notable that object B has a motion onset in the causal event as well, but chicks appear to be sensitive to the fact that the motion onset is due to the impact with A. Put differently, chicks in this experiment seem to have concluded that object B, as a causal patient, is inert.

Human infants also infer that causal patients in launching events are inert (Luo, Kaufman, & Baillargeon, 2009; Saxe & Carey, 2006). This inference of inertness is not trivial, especially considering that it seems unique to launching events, and does not occur in contact events in which B's movement is delayed relative to the moment of contact. Consider the subtlety of this distinction: In both cases, infants are observing B's motion onset following A's movement, but only in the very specific spatiotemporal parameters of launching does this onset fail to indicate that B is capable of self-propelled motion. Implicitly, this suggests sensitivity to something about the physics of the real-world collisions approximated by launching events: B's motion in a launching event is completely accounted for by the collision with A.

This analogy to real-world collisions is obvious to the point of being trivial (how could causal perception exist if it did not apply to some natural events?), but the implications are far-reaching. In particular, in order for causal perception to serve as a *reliable* cue to whether B's motion onset indicates that object B is self-propelled or inert, it would have to be sensitive to the physical limits on B's motion.

Intuitive physics and triggering events

The human mind is simultaneously extremely adept and extremely poor at understanding Newtonian mechanics. We are prone to making egregious errors in explicit prediction tasks (Hecht & Bertamini, 2000; McCloskey, 1983; Liu & MacIsaac, 2005), but perception is so well-calibrated to the physical parameters of the natural world that our implicit predictions (e.g., where we put our hands to catch a falling object) are remarkably accurate and robust. In fact, such predictions are only hopelessly thrown off when we alter the most fundamental parameters of the physical environment (e.g., gravity; McIntyre, Zago, Berthoz, & Lacquaniti, 2001).

Recent work has proposed that we may even possess a mental "physics engine" (Ullman, Spelke, Battaglia, & Tenenbaum, 2017), and our judgments can be modeled successfully as accurate physical principles applied to noisy sensory input (Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2012). However, even if such results do not indicate a sophisticated mental model of physics, it is undeniable that perception is sensitive to certain physical principles, like spatiotemporal continuity and solidity, from very early in life (Carey, 2009; Spelke, Breinlinger, Macomber, & Jacobson, 1992).

In the context of causal perception, the relevant physical constraints are of course those that apply to collisions. For example, in a perfectly elastic collision involving two objects, with no outside forces acting on the system, Newtonian mechanics imposes a "speed limit" on the relative speed of B following the collision. In simple terms, and provable mathematically, B can never move at more than double the speed of A, without some extra force acting on the system. Michotte described events that violate this Newtonian speed limit as "triggering" events, in which B's motion is *autonomous*, but still initiated by contact with A (Michotte, 1963; Natsoulas, 1961).

Recent work has found that these triggering events are still perceived as causal, but are categorically different from launching events (Kominsky et al., 2017). These causal event categories are distinguished in automatic perceptual processing in adults: triggering events in an array of launching events produce an oddball advantage, but no such advantage is found for non-causal events with the same speed ratios.

Critically, this categorical distinction was also found in 7-9-month-old infants, using a dishabituation paradigm. Within a month of the earliest age at which causal perception is reliably found, infants are not only sensitive to whether an event is causal or not, they are sensitive to the kind of causal relationship, as informed by physical constraints. However, this experiment only examined physical constraints on speed ratio.

Kominsky et al. (2017) concluded that the categorical distinction between launching and triggering was not due to a precise representation of this Newtonian constraint on real-world collision events. Rather, the boundary seems to be defined by a detectable increase in B's speed relative to A's, and the threshold for detectability is in the vicinity of the Newtonian limit. Thus, this categorical distinction could mimic real-world physical constraints only to the degree that there is a sensitivity to speed ratio information in causal events, not sensitivity to the physics of collisions in general.

Alternatively, the categorical distinction between launching and triggering could reflect a broader sensitivity to the physical constraints on collision events, limited by the inherent noise in sensory input. Under this view, we would expect that causal perception should distinguish different categories of causal events based on whether object B violates physical constraints on collision events. The threshold for detecting that violation might be imprecise, but as long as the violation is detected, it might demarcate a categorical boundary in perception.

The contrast between these two hypotheses has particular relevance for our theories of infant causal perception. In some domains, infants are narrowly sensitive to specific types of information for specific types of events at different points in development, e.g. the use of height and width information in containment versus occlusion events (Baillargeon & Wang, 2002; Strickland & Scholl, 2015). Thus, in the case of causal perception, one might expect that infants could be sensitive to some physical constraints and not others.

Alternately, causal perception may intrinsically incorporate more general physical principles. Unlike containment and occlusion, causality in motion events is in some sense completely defined by the specific physical constraints imposed by Newtonian mechanics. Contact and immediate reaction are both intrinsic aspects of Newtonian physics, but so is the speed restriction, and every other physical constraint on collisions. It may be impossible for the visual system to define causal events without incorporating more general physical principles. Under this view, we should expect that at any age we can demonstrate causal perception, we should also be able to demonstrate a sensitivity to *multiple* physical constraints (though not necessarily all).

The current experiment

Here, we investigate whether infant causal perception is sensitive to physical constraints on the *angle* of B's motion relative to A.

The Newtonian speed limit constraint can be easily proven mathematically, but the angle constraint is best understood intuitively. Imagine the momentum vector of object A as it impacts object B, pointing along the direction of A's motion. Even in an off-center collision, some component of this vector must be preserved in B's resulting momentum vector. In other words, if A was moving forwards, even in the most glancing collision B must end up moving forward to some degree. Therefore, the angle limit on B's motion relative to A is a straightforward 90°. If B moves at an angle $\geq 90°$ relative to A, by definition B's momentum vector contains none of A's.

Of course, we once again run into the issue of how precise we can expect the visual system to be. Comparing 89.9° to 90° movement in an off-center collision between two spheres would quite likely fail to identify a categorical boundary, especially in infants. The perceptual system is unlikely to represent the boundary so precisely, given the inherent noise in perceptual processing. Therefore, to make the violation completely unambiguous, we created displays involving fully on-center contact between two square objects, and relative angles of 0° or 90°.

In the current experiment, we follow the design of Kominsky et al. (2017)'s Experiment 3, and use a classic habituation/dishabituation design (Colombo & Mitchell, 2009). Infants in the causal condition are habituated to a launching event, and are then shown either a test event which violates this angle constraint (Fig. 1b, top), or “rotated” launching event in which *the whole launching event* is rotated 90°, not just the movement of one of the objects (Fig. 1b, bottom).

The prediction is straightforward: If infants are sensitive to the angle constraint, then they should look longer to the violation test event than the control test event. Notably, there is a clear and contrary alternative hypothesis based on low-level features: the rotated event is actually more different from the habituation event in terms of the motion characteristics of each object, the area of space occupied by the whole event, etc., even though the relationship between the objects is unchanged. Therefore, if infants are not sensitive to the causal relationship between A and B, only to the differences in low-level features or the behavior of each individual object, they should look longer at the rotated event than the violation event.

Furthermore, as in Kominsky et al. (2017), this selective dishabituation should only hold for causal events. For non-causal delay test events (identical except for a 500ms delay before the start of B’s movement), infants should either look longer at the rotated event (based on the low-level cues) or look equally at both (if they are only sensitive to the causal relationship between A and B, which is nonexistent in both non-causal test events).

Methods

Participants Sixty-four infants (30 female, 34 male) age 7.5 months to 9.5 months participated in the experiment. An additional 19 infants (8 female, 11 male) were recruited but excluded from the final analysis for fussing out (2), moving off-camera during the experiment (5), parental interference (3), experimenter error (2), or an above-threshold discrepancy during offline re-coding (7, see below). One additional (female) participant was replaced in the final sample due to having a test trial looking time >3 standard deviations from the average for her condition (a predetermined exclusion criterion).

Apparatus Stimuli were controlled using PyHab (Kominsky, 2017). PyHab uses the PsychoPy stimulus presentation libraries as its base (Peirce, 2007). It is designed to provide all of the functionality of other looking time coding software like XHab or JHab (recording infants’ looking time by holding down a key whenever they are looking at the display and releasing the key when they are not), but in addition PyHab automatically controls the timing and content of stimulus presentation according to the experimenter’s live coding of looking times and a pre-set experimental design.

Stimuli were presented on a 25” wide by 15.5” high Apple Cinema Display operating at 1280x800 pixel resolution and 60 frames per second. The edges of the display were hidden behind a black foamcore frame, with black fabric around the frame running floor to ceiling and about a foot on either side. Beige curtains obscured the rest of the room from the infants’ view. Infants sat on their parent’s lap about 56” from the display screen. A hidden camera located directly under the center of the display monitor recorded infants’ looking behavior and displayed a live feed to the experimenter. Light was provided by four overhead dimmable compact fluorescent track lights set at approximately 10% brightness.

Stimuli and procedure After providing informed consent, parents were instructed to sit in a chair facing the display screen with their infant in their lap, and asked to close their eyes and avoid interacting socially with their infant for the duration of the experiment (they were shown the stimuli afterward). They were also asked to try to prevent their infant from standing up on their lap, in order to keep the infant’s face in view of the camera.

This experiment used a 2 (causal vs. delay) x 2 (rotated vs. angle violation test) between-subjects design. Condition assignment was randomized, and looking-time coding was always done with the experimenter blind to condition.

In all four conditions, the basic parameters of a trial were the same: At the start of the trial, the experimenter pressed a key to play an attention-getter, consisting of a rapidly looming and spinning yellow rectangle and a rapid rising series of notes, taking exactly 1.1 seconds. Immediately following this, two squares appeared on the screen, one red and one green, each one 80 pixels on a side with a 240 pixel gap between them. One square was always adjacent to the center of the screen, and the other to its left or right (counterbalanced between subjects). The two squares appeared static on the screen for a minimum of 200ms. The experimenter could play the attention-getter again if the infant failed to look at the screen initially, and the trial did not start until the infant looked at the screen.

Each trial started after the attention-getter when the infant initially looked at the screen, and lasted until the infant looked away for 2 *consecutive* seconds or 60 seconds had passed. The stimuli are depicted in Fig. 1.

First, infants saw up to 14 habituation trials consisting of either launching events (causal condition) or contact events with a 500ms delay before object B began moving (delay condition). Each object’s movement took one second and

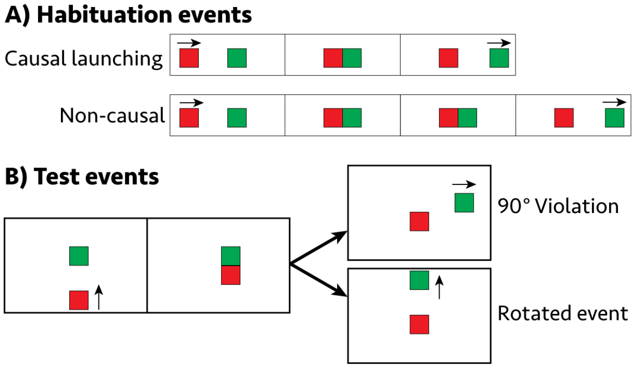

Figure 1. Stimuli used in the experiment.

covered 240 pixels. Following object B’s movement, both squares vanished for 333ms, and then the animation repeated from the beginning.

The habituation criterion was calculated as the sum of the infants’ looking time over the first three trials, divided by two. The experiment moved on to the test trial when the sum of the infant’s looking time across three (subsequent) consecutive trials was less than the criterion. The experimenter was not informed when the criterion was met. (Five infants never reached criterion but saw the test trial after 14 habituation trials, and their data were included in the analysis.)

Infants then saw a single test trial. In the test trial, the arrangement of the squares was different. Object B was now in the center of the screen, and object A 240 pixels directly below it. In both test conditions, object A moved up toward object B until they were adjacent, at which point object B began moving either immediately (in the causal condition) or after a 500ms delay (in the delay condition). In the rotated condition, object B moved up as well, preserving the linearity of the habituation trials. In the angle violation condition, object B moved identically to how it had moved in the habituation trials, i.e. to the left or right. Thus, in the causal-angle violation condition, this event violated the angle constraint on collision events.

Results

Coding and reliability In addition to the live-coding, each participants’ experimental session was re-coded from video by a separate coder, who was also blind to condition. We used predetermined exclusion criteria to identify significant discrepancies and remove participants for whom the initial coding (and therefore stimulus presentation) appeared to be in error. The criteria were as follows:

If there was a disagreement on the total looking time of a trial of greater than 10% of the greater looking time, and the discrepant trial either changed when the habituation criterion would have been met or was the test trial itself, the participants’ data were excluded (this occurred for seven participants’ data). If the margin by which that threshold was exceeded was less than 250ms, a third independent coder re-coded the video, and if they were in agreement with the live coding the data were left in (this occurred for six participants’ data). If the second and third coder were in agreement, and the discrepancy was specifically that the test trial looking time was longer in the live coding (i.e., the stimuli were still displayed for an appropriate period of time), then the live coding was replaced with the secondary coding and the infant’s data were included in the analysis (this occurred for two participants’ data).

Ultimately the exclusions by condition were as follows: 4 from the causal angle violation condition, 9 from the causal rotated condition, 4 from the delay angle violation condition, and 2 from the delay rotated condition.

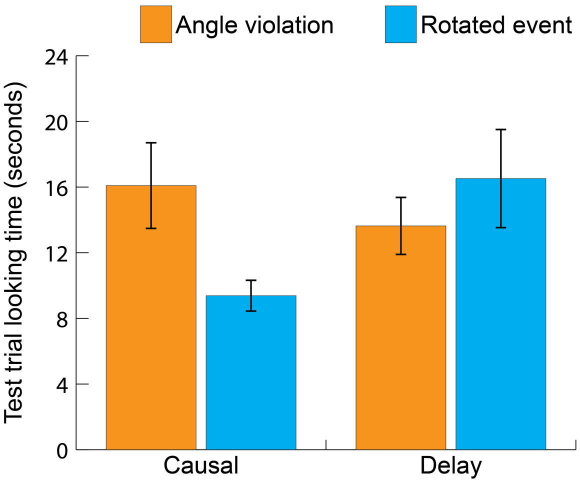

Figure 2. Results. Error bars represent ± 1 SEM.

Analyses Average test trial looking time by condition can be found in Fig. 2. Immediately, one can see a clear difference in looking times in the causal condition such that infants looked longer at the angle violation test event, and no such selective dishabituation in the delay condition.

We confirmed these impressions with the following analyses: A 2 x 2 ANOVA showed no main effect of causal versus delay, $F(1, 60) = 1.07$, $p = .3$, and no main effect of test event, $F(1, 60) = .72$, $p = .4$, but a significant interaction, $F(1, 60) = 4.50$, $p = .038$, $\eta_p^2 = .07$. Planned comparisons of the effect of test event in each causal condition found that infants in the causal condition looked significantly longer at the angle violation test event ($M = 16.09$s, $SD = 10.63$) than the rotated test event ($M = 9.38$s, $SD = 3.95$), $t(31) = 2.37$, $p = .029$, $d = .84$. In the delay condition, there was no significant difference between the angle violation ($M = 13.63$s, $SD = 7.14$) and rotated test events ($M = 16.52$s, $SD = 12.15$), $t(31) = .82$, $p = .4$.

Discussion

By 7-9 months of age, human infants draw categorical boundaries between causal events that reflect the Newtonian constraints on real-world collisions. In previous work we have demonstrated sensitivity for constraints on relative speed (Kominsky et al., 2017), and here we demonstrate a nearly identical pattern for relative angle.

This result is a striking contrast with other infant work on event perception, which has often found that being sensitive to one kind of physical information about an event does not entail being sensitive to other kinds physical information about an event. For example, 6.5-month-old infants are sensitive to width information in containment events, but do not show sensitivity to height information until 7.5 months, or understand transparent containers until 10 months (Baillargeon & Wang, 2002). This event type distinction impacts perceptual processing of height and width into adulthood (Strickland & Scholl, 2015).

Of course, causal perception is often found at 6 months of age (Cohen & Amsel, 1998; Leslie & Keeble, 1987), and the current experiment started at 7.5 months. It is possible that when causal perception first emerges, infants are sensitive to angle information or speed information, but not both. Future work with narrower and earlier age ranges will clarify whether causal perception *intrinsically* includes these Newtonian constraints, i.e. if there is any point in development when infants perceive launching as causal but do not distinguish it from triggering.

Why do these physical constraints matter?

Infants and newborn chicks treat causal patients in launching events as inert objects (Luo et al., 2009; Mascalzoni et al., 2010), which is appropriate when the parameters of those events fall within the Newtonian constraints on collision events. What about events that defy those constraints?

We refer to these events as physical violations, but in truth they are only violations of what is physically possible *based on the force of the collision alone* (plus gravity and friction). Objects with an internal source of motive force are not subject to such restrictions. This, of course, is why these events are often described as “triggering”: real-world events of this type involve one object triggering the autonomous motion of a different object, or the addition of some other force beyond the force imparted by the collision with object A (such as a chemical reaction). In the natural environment, we might observe triggering events when the causal patient is animate. If you touch an animal that was previously unaware of your presence, you will often find they move much faster and in a different direction than your hand.

It is easy enough to see the evolutionary conjecture that results from this logic. Especially given its role in chick imprinting behavior, causal perception may have come about specifically to aid in the identification of agents and animate entities, and filter out false-positive motion onsets by inert objects in causal events. The importance of identifying such entities quickly, efficiently, and accurately is obvious. In the natural environment, a self-propelled entity is going to be prey, predator, or social partner. In the case of an infant organism, it is especially critical to identify a caregiver who is animate and agentic to provide adequate care. Motion cues are not the only cues to agency by any means (Muentener & Carey, 2010; Saxe, Tenenbaum, & Carey, 2005; Setoh, Wu, Baillargeon, & Gelman, 2013; Träuble & Pauen, 2011), but they are good cues when static object features are ambiguous.

Such an account would predict that physical violations which require the presence of additional forces should all be detected equally easily: failing to recognize self-propelled motion in the natural environment because a particular physical constraint does not yet define a categorical boundary could be deeply costly, as a self-propelled entity is most likely prey, predator, or social partner, all survival-relevant categories. The current results are perfectly in line with this expectation, but there are many critical questions left to answer. For one, there are physical constraints that have not been studied. However, the most important question is what inferences infants draw about triggered objects, and how those inferences compare to the ones they make about launched objects, such as whether they are expected to be fully self-propelled (Luo et al., 2009), or have insides (Setoh et al., 2013). Work is ongoing to address these questions.

However, there are alternative explanations for our results. For example, these 90° angle violation events may not be seen as causal at all. It is well-established that when infants are habituated to causal events, they dishabituate to non-causal events (e.g., Cohen & Amsel, 1998). Indeed, work with adults using subjective ratings scales has suggested that the greater the angle difference, the less causal the event appears (White, 2012). Such measures do not necessarily capture perceptual processing, but this result raises a valid concern. Future work is planned to explore this possibility by determining whether habituating to a causal angle violation event leads to dishabituation on a non-causal angle violation event, i.e. if infants consider the addition of a temporal delay to be a meaningful change.

Conclusion

Infants show remarkably sophisticated sensitivity to the physical constraints on the natural world, not just on the behavior of individual objects, but also on interactions between objects. Here, we provide initial evidence that infants make a categorical distinction between different types of causal events on the basis of angle constraints. In doing so, we raise the possibility, yet to be tested, that general principles of Newtonian mechanics may be intrinsically incorporated into causal perception.

References

Baillargeon, R., & Wang, S.-h. (2002). Event categorization in infancy. *Trends in Cognitive Sciences*, *6*(2), 85-93.

Carey, S. (2009). *The origin of concepts.* Oxford ; New York: Oxford University Press.

Cohen, L. B., & Amsel, G. (1998). Precursors to infants’ perception of the causality of a simple event. *Infant behavior and development*, *21*(4), 713-731.

Colombo, J., & Mitchell, D. W. (2009). Infant visual habituation. *Neurobiology of learning and memory*, *92*(2), 225-234.

Gerstenberg, T., Goodman, N., Lagnado, D. A., & Tenenbaum, J. B. (2012). *Noisy Newtons: Unifying process and dependency accounts of causal attribution*. Proceedings of the 34th annual meeting of the Cognitive Science Society.

Hecht, H., & Bertamini, M. (2000). Understanding projectile acceleration. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(2), 730-746.

Hubbard, T. L. (2013). Phenomenal Causality I: Varieties and Variables. *Axiomathes*, *23*(1), 1-42.

Hume, D. (1902). *An enquiry concerning human understanding*. Retrieved from Project Gutenberg

Karaminis, T., Turi, M., Neil, L., Badcock, N. A., Burr, D., & Pellicano, E. (2015). Atypicalities in perceptual adaptation in autism do not extend to perceptual causality. *PLoS One*, *10*(3), e0120439.

Kim, S.-H., Feldman, J., & Singh, M. (2013). Perceived causality can alter the perceived trajectory of apparent motion. *Psychological Science*, *24*(4), 575-582.

Kominsky, J. F. (2017) PyHab: Looking time and stimulus presentation script for PsychoPy. [Computer software]. Retrieved from https://github.com/jfkominsky/PyHab

Kominsky, J. F., Strickland, B., Wertz, A. E., Elsner, C., Wynn, K., & Keil, F. C. (2017). Categories and Constraints in Causal Perception. *Psychological Science*, *28*(11), 1649-1662.

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*(3), 265-288.

Liu, X., & MacIsaac, D. (2005). An investigation of factors affecting the degree of naive impetus theory application. *Journal of Science Education and Technology*, *14*(1), 101-116.

Luo, Y., Kaufman, L., & Baillargeon, R. (2009). Young infants' reasoning about physical events involving inert and self-propelled objects. *Cognitive Psychology*, *58*(4), 441-486.

Mascalzoni, E., Regolin, L., & Vallortigara, G. (2010). Innate sensitivity for self-propelled causal agency in newly hatched chicks. *Proceedings of the National Academy of Sciences, 107*(9), 4483-4485.

McCloskey, M. (1983). Naive theories of motion. *Mental models*, 299-324.

McIntyre, J., Zago, M., Berthoz, A., & Lacquaniti, F. (2001). Does the brain model Newton's laws? *Nature Neuroscience*, *4*(7), 693-694.

Michotte, A. (1963). *The Perception of Causality*. New York, NY: Basic Books. Original printing 1946.

Moors, P., Wagemans, J., & de-Wit, L. (2017). Causal events enter awareness faster than non-causal events. *PeerJ*, *5*, e2932.

Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology*, *61*(2), 63-86.

Natsoulas, T. (1961). Principles of Momentum and Kinetic Energy in the Perception of Causality. *The American Journal of Psychology*, *74*(3), 394-402.

Newman, G. E., Choi, H., Wynn, K., & Scholl, B. J. (2008). The origins of causal perception: evidence from postdictive processing in infancy. *Cognitive Psychology*, *57*(3), 262-291.

Peirce, J. W. (2007). PsychoPy--Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1-2), 8-13.

Piaget, J., & Garcia, R. (1977). *Understanding causality* (D. Miles & M. Miles, Trans.). New York: W.W. Norton.

Rips, L. J. (2011). Causation From Perception. *Perspectives on Psychological Science*, *6*(1), 77-97.

Rolfs, M., Dambacher, M., & Cavanagh, P. (2013). Visual Adaptation of the Perception of Causality. *Current Biology*, *23*(3), 250-254.

Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret Agents: Inferences About Hidden Causes by 10-and 12-Month-Old Infants. *Psychological Science*, *16*(12), 995-1001.

Saxe, R., & Carey, S. (2006). The perception of causality in infancy. *Acta Psychologica*, *123*(1-2), 144-165.

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends Cognitive Science*, *4*(8), 299-309.

Setoh, P., Wu, D., Baillargeon, R., & Gelman, R. (2013). Young infants have biological expectations about animals. *Proceedings of the National Academy of Sciences*, *110*(40), 15937-15942.

Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological review*, *99*(4), 605-632.

Strickland, B., & Scholl, B. J. (2015). Visual perception involves event type representations: The case of containment vs. occlusion. *Journal of Experimental Psychology: General, 144*(3), 570-580.

Träuble, B., & Pauen, S. (2011). Infants' reasoning about ambiguous motion events: The role of spatiotemporal and dispositional status information. *Cognitive Development*, *26*(1), 1-15.

Ullman, T. D., Spelke, E., Battaglia, P., & Tenenbaum, J. B. (2017). Mind Games: Game Engines as an Architecture for Intuitive Physics. *Trends in Cognitive Science*.

White, P. A. (2012). Visual impressions of causality: Effects of manipulating the direction of the target object's motion in a collision event. *Visual Cognition*, *20*(2), 121-142.