



Published in final edited form as:

Psychol Rev. 2019 March ; 126(2): 196–225. doi:10.1037/rev0000136.

Catastrophic Individuation Failures in Infancy: A New Model and Predictions

Maayan Stavans^a, Yi Lin^b, Di Wu^c, and Renée Baillargeon^b

^aPsychology Department, Bar-Ilan University, Ramat-Gan 5290002, Israel

^bPsychology Department, University of Illinois, Urbana-Champaign, IL 61820, USA

^cPsychology Department, Cedarville University, Cedarville, OH 45314, USA

Abstract

Comparison of infant findings from the physical-reasoning and object-individuation literatures reveals a contradictory picture. On the one hand, physical-reasoning results indicate that young infants can use featural information to guide their actions on objects and to detect interaction violations (when objects interact in ways that are not physically possible) as well as change violations (when objects spontaneously undergo featural changes that are not physically possible). On the other hand, object-individuation results indicate that young infants typically cannot use featural information to detect individuation violations (when the number of objects revealed at the end of an event is less than the number of objects introduced during the event). In this article, we attempt to reconcile these two bodies of research. In a new model of early individuation, we propose that two systems help infants individuate objects in physical events, the object-file and physical-reasoning systems; under certain conditions, disagreements between the systems result in catastrophic individuation failures, leading infants to hold no expectation at all about how many objects are present. We report experiments with 9–11-month-old infants ($N = 216$) that tested predictions from the model. After two objects emerged in alternation from behind a screen, infants detected no violation when the screen was lowered to reveal no object. Similarly, after two objects emerged in alternation from inside a box, which was then shaken, infants detected no violation when the box remained silent, as though empty. We end with new directions, suggested by our model, for research on early object representations.

Keywords

infant cognition; object individuation; object-file system; physical-reasoning system

Corresponding author: Maayan Stavans, Psychology Department, Bar-Ilan University, Ramat-Gan 5290002, Israel. maayan.stavans@biu.ac.il.

Disclosures The authors declare that they have no conflicts of interests to disclose.

Footnotes

1. Introduction

Research over the past few decades has begun to shed light on how infants represent and reason about the physical world (for reviews, see Baillargeon et al., 2012; Baillargeon, Li, Gertner, & Wu, 2011; Baillargeon, Li, Ng, & Yuan, 2009a; Baillargeon, Wu, Yuan, Li, & Luo, 2009b). These investigations have produced four main findings. First, infants' reasoning about physical events is guided by a skeletal framework of core principles and concepts (Carey, 2011; Gelman, 1990; Keil, 1995; Leslie, 1995; Luo, Kaufman, & Baillargeon, 2009; Mascalzoni, Regolin, Vallortigara, & Simion, 2013; Spelke, Breinlinger, Macomber, & Jacobson, 1992; Ullman, Spelke, Battaglia, & Tenenbaum, 2017). One of these principles is that of *persistence*: All other things being equal, objects persist as they are, with all of their properties (Baillargeon, 2008; Baillargeon et al., 2009a). The persistence principle has many corollaries, which dictate that an object cannot spontaneously disappear (continuity), break apart (cohesion), pass through another object (solidity), fuse with another object (boundedness), or change into a different object (unchangeableness) (Baillargeon et al., 2009a; Spelke et al., 1992; Spelke, Phillips, & Woodward, 1995; Stahl & Feigenson, 2015).

Second, infants' representations of physical events are initially very sparse and become gradually richer with experience (Aguiar & Baillargeon, 1999; Baillargeon, 1991; Kotovsky & Baillargeon, 1998; Luo & Baillargeon, 2005). As infants observe and produce events, they form distinct event categories, such as occlusion, containment, support, and collision events (Casasola, 2008; Hespos & Baillargeon, 2001a; Mou & Luo, 2017; Wang, Baillargeon, & Paterson, 2005). For each category, infants acquire physical rules that identify causally relevant features for predicting outcomes; information about these identified features (e.g., properties of objects and their arrangements) is then included when representing events from the category, resulting in increasingly detailed event representations over time (Aguiar & Baillargeon, 2002; Baillargeon, Needham, & DeVos, 1992; Wang, Zhang, & Baillargeon, 2016; Wilcox, 1999).

Third, one of the processes by which infants acquire their physical rules is explanation-based learning (Baillargeon & DeJong, 2017; DeJong, 2014; Wang & Baillargeon, 2008; Wang & Kohne, 2007). When infants encounter contrastive outcomes they cannot predict, they search for a feature that distinguishes these outcomes. If they find one, they attempt to construct a plausible explanation for it, using their physical knowledge (i.e., skeletal framework and acquired rules). This explanation is then generalized into a candidate rule and evaluated against a few additional empirical observations. In some cases, the new rule serves to establish a new event category; in other cases, it serves to revise an existing rule. In the case of occlusion, for example, infants' initial rule is akin to "an object is hidden when behind an occluder" (Aguiar & Baillargeon, 1999; Luo & Baillargeon, 2005). With experience, this rule is extensively elaborated as infants identify causally relevant features for predicting, for example, whether an object will be fully or only partly hidden when behind an occluder (e.g., is the object larger than the occluder?) and whether the object that reappears from behind an occluder is the same one that disappeared behind it or a different one (e.g., is the object that reappeared the same size, shape, pattern, or color as the one that disappeared?)

(Baillargeon & DeVos, 1991; Wang, Baillargeon, & Brueckner, 2004; Wilcox, 1999; Wilcox & Baillargeon, 1998).

Finally, featural information, once included in an event representation and interpreted by infants' physical knowledge, is used both to guide actions and to detect violations (Hauf, Paulus, & Baillargeon, 2012; Wang, 2011; Wang & Kohne, 2007; Wilcox & Chapa, 2004). To illustrate, consider the feature size, which is identified as causally relevant to occlusion events by about 4 months of age (Baillargeon & DeVos, 1991; Wang et al., 2004). Infants who have acquired this feature use it to act prospectively: When searching for a large object, for example, they are more likely to reach for a large as opposed to a small occluder (Hespos & Baillargeon, 2006). Infants also detect several types of persistence violations involving this feature, including *interaction* violations (i.e., objects interact in ways that are not physically possible given their respective properties) and *change* violations (i.e., objects spontaneously undergo changes that are not physically possible). Thus, infants detect a violation if a large object becomes fully hidden behind a small occluder (interaction violation; Hespos & Baillargeon, 2001a), or if a small object is much larger after being briefly occluded (change violation; Wang & Baillargeon, 2006). Other occlusion features, such as shape (identified at about 4 months), pattern (7 months), and color (11 months) have yielded similar findings; for example, 11-month-olds detect a violation if an object surreptitiously changes color when occluded, but younger infants do not (Kaldy & Leslie, 2003; Wilcox, 1999; for neuroimaging evidence showing the same developmental sequence, see Wilcox & Biondi, 2015; Wilcox, Stubbs, Hirshkowitz, & Boas, 2012).

1.1. Early Failures at Individuation Tasks

Strikingly, however, infants age 12 months and younger fail to detect another type of persistence violation, *individuation* violations (i.e., the number of objects revealed at the end of an event is less than the number of objects introduced during the event, as though one or more of the objects spontaneously disappeared). Xu and Carey (1996) were the first to discover this baffling failure in infants' physical reasoning. Negative results have been obtained with at least three types of individuation tasks (see Fig. 1; for ease of comparison across tasks and ages, the objects to individuate are always two different objects from the same basic-level category, namely, two balls that differ in size, pattern, and color).

Standard within-subject task.—In the within-subject version of the standard task (Fig. 1A; Xu & Carey, 1996; Xu, Carey, & Quint, 2004), infants first see one ball emerge on one side of a large screen and then return behind it; next, the other ball emerges on the opposite side of the screen and again returns behind it. After several repetitions of this occlusion event, the screen is removed to reveal either both balls (*two-objects* event) or only one of the balls (*one-object* event). Because infants typically show a baseline preference for two-objects over one-object displays (as displays with two different objects often take longer to encode and are more interesting visually), infants who correctly individuate the two balls are expected to look about equally at the two test events: They should show enhanced attention to the one-object event due to the individuation violation in this event; they should show enhanced attention to the two-objects event due to their baseline preference for two-objects displays; and these opposing responses should cancel each other, resulting in approximately

equal looking times at the two test events. In contrast to this pattern, however, infants age 12 months and younger typically show only a baseline preference for the two-objects event, suggesting that they fail to detect the violation in the one-object event and hence that they do not clearly expect to see two objects when the screen is removed (e.g., Bonatti, Frót, Zangl, & Mehler, 2002; Futo, Téglás, Csibra, & Gergely, 2010; Krajgaard, 2000; Leslie, Xu, Tremoulet, & Scholl, 1998; Rivera & Zawaydeh, 2007; Surian & Caldi, 2010; Xu & Carey, 1996; Xu et al., 2004).

Standard between-subjects task.—The between-subjects version of the standard task was designed to circumvent the interpretive difficulties associated with infants' baseline preference for two-objects displays (Fig. 1B; Wilcox & Baillargeon, 1998; Wilcox & Chapa, 2002). To start, infants watch either a *different-objects* event, in which the two different balls emerge in alternation from behind the screen, or a *same-object* event, in which the same ball emerges on either side of the screen. Following this occlusion event, all infants see the same final display:

The screen is lowered to reveal only one ball. If infants who see the different-objects event correctly individuate the two balls, they should look significantly longer at the final display than infants who see the same-object event. However, infants age 12 months and younger tend to look equally at the final display across events, suggesting that infants who see the different-objects event do not clearly expect two objects to be present when the screen is removed (Stavans & Baillargeon, 2018; Wilcox, 2007; Wilcox & Baillargeon, 1998; Wilcox & Chapa, 2002).

Search task.—Yet another individuation task involves a containment event followed by a manual-search event (Fig. 1C; Van de Walle, Carey, & Prevor, 2000). Infants face a large box that has a spandex-filled opening with a narrow slit, so that the box's contents are not visible. At the start of each trial, infants see either a different-objects or a same-object event. In the *different-objects* event, an experimenter removes one of the balls from the box, returns it to the box, and then repeats these actions with the other ball; the *same-object* event is identical except that the same ball is removed from and returned to the box. Following this containment event, the box is moved within infants' reach, and they are encouraged to search.¹ If infants correctly determine how many balls are involved in each event, they should search significantly more, after the first ball is retrieved, when shown the different-objects as opposed to the same-object event. However, infants age 12 months and younger tend to search equally across events, suggesting that they do not clearly expect to find two objects inside the box after seeing the different-objects event (Van de Walle, et al., 2000).

¹Readers might wonder why we describe this sequence as a containment event followed by a search event, rather than as an ongoing containment event, given that the box remains in the scene. In the containment event, infants passively view objects being removed from and returned to the box; in the search event, however, infants plan and perform search actions, and it seems plausible that they would represent these actions as a new event or physical interaction with its own causally relevant features (e.g., the distance of their hand to the box, the orientation of their hand relative to the box's narrow opening, and so on). In Experiment 3, we showed infants a novel sequence of two events involving a box: Objects were removed from and returned to a box, which was then shaken. Results indicated that infants represented the shaking as a separate event, in which they expected the objects in the box to noisily collide with each other and with the box's interior walls.

In sum, infants' failures at detecting individuation violations in standard and search tasks are highly puzzling given their successes at detecting interaction and change violations involving similar occlusion or containment features (Aguiar & Baillargeon, 1998; Kaldy & Leslie, 2003; Kibbe & Feigenson, 2016; Wang et al., 2004; Wang & Onishi, 2017; Wilcox, 1999). When two balls that differ in size, pattern, and color emerge in alternation from behind a screen or from inside a box, infants age 4 months and older represent at least some of the balls' distinguishing features and view these features as persistent. Why, then, do even 12-month-olds give no indication that they expect both balls to be present when the screen is removed or when they search the box?

1.2. Early Successes at Individuation Tasks

Infants' perplexing failures at standard and search individuation tasks gave rise to a large body of empirical work as researchers attempted to make sense of these failures. In the course of these investigations, researchers identified at least four separate factors that allow infants age 12 months and younger to succeed at individuation tasks (see Fig. 2; for ease of comparison across tasks and ages, we demonstrate all four factors using the different-objects event from a standard between-subjects task, and where possible we continue to use the two-balls example).

Different-locations task.—Young infants succeed at an individuation task involving two different balls if at some point during the occlusion event (standard task) or the containment event (search task), the two balls briefly emerge into view simultaneously (Fig. 2A; Van de Walle, et al., 2000; Xu & Carey, 1996; Xu et al., 2004; Zosh & Feigenson, 2015). Because the balls occupy visibly different locations, infants establish separate representations for them and detect a violation if one of them spontaneously disappears.

Different-categories task.: Young infants succeed at an individuation task if they are tested with two objects they assign to different categories, as opposed to two objects they assign to the same category. In other words, infants succeed when they encode the objects as *categorically* distinct, but they fail when they encode the objects as merely *featurally* distinct. Critically, the nature of the object categories that infants spontaneously encode changes over development, with significant impact on their performance. Prior to their first birthday, most infants do not spontaneously encode an isolated object's *basic-level* category, such as ball, toy duck, block, or cup (Pauen, 2002; Xu & Carey, 1996). However, young infants do encode more abstract or ontological categorical information about the object, such as whether it is human-like or non-human and whether it is animate or inanimate (Bonatti, et al., 2002; Bonatti, Frot, & Mehler, 2005; Setoh, Wu, Baillargeon, & Gelman, 2013; Surian & Caldi, 2010). Thus, 9–10-month-olds succeed when tested with two objects from different ontological categories, such as a human-like and a non-human object (e.g., a doll and a ball; Fig. 2B; Bonatti et al., 2002, 2005), but they fail when tested with two objects from the same ontological category that differ only in their basic-level categories and/or their individual features (e.g., a toy duck and a ball, or two different balls; Bonatti et al., 2002; Surian & Caldi, 2010; Xu & Carey, 1996; Xu et al., 2004).

By their first birthday, infants begin to spontaneously encode objects' basic-level categories (Cacchione, Schaub, & Rakoczy, 2013; Mandler & McDonough, 1998; Pauen, 2002; Xu & Carey, 1996). Thus, infants now succeed when tested with two objects from different basic-level categories (e.g., a toy duck and a ball), but still fail when tested with two objects from the same basic-level category that differ only in their individual features (e.g., two different balls) (Cacchione et al., 2013; Schaub, Bertin, & Cacchione, 2013; Van de Walle et al., 2000; Wilcox & Baillargeon, 1998; Xu & Carey, 1996; Xu et al., 2004). There is also evidence that infants in the second year of life spontaneously assign novel animate entities with different (visually accessible) insides to different categories (Newman, Herrmann, Wynn, & Keil, 2008; Welder & Graham, 2006). Building on these findings, Taborda-Osorio and Cheries (2018) recently tested 13-month-olds in a search task using two transparent animate entities that differed only in the color of their insides or only in the color of their outsides. As expected, infants succeeded with the different insides (they encoded the entities as categorically distinct), but they failed with the different outsides (they encoded the entities as merely featurally distinct).

Although infants younger than 12 months typically do not encode objects' basic-level categories, they can be induced to do so via experimental manipulations, with positive effects on their individuation performance (Futó et al., 2010; Stavans & Baillargeon, 2018; Xu, 2002). For example, in a standard within-subject task using a language-based manipulation (Xu, 2002), 9-month-olds heard a distinct label (i.e., "Look, [baby's name], a duck!" or "Look, [baby's name], a ball!") as each object emerged into view during the occlusion event. Following this manipulation, infants detected a violation when the screen was lowered to reveal only one of the objects. The same positive result was obtained with novel objects and labels ("a fendle", "a toma"), but was eliminated if the two objects received the same label ("a toy"). In a standard between-subjects task using a function-based manipulation (Stavans & Baillargeon, 2018), 4-month-olds first saw a functional demonstration for each of two tools, a masher and a pair of tongs: In separate trials, the masher was used to compress sponges, and the tongs were used to pick them up. In the test trial, infants saw the two tools in alternation (different-objects event) or the same tool (same-object event) emerge on either side of a screen. When the screen was finally lowered to reveal one tool, infants who had seen the different-objects event looked significantly longer at this final display, suggesting that they expected to see both tools. The same positive result was obtained with two other tools, a marker and a knife, but was eliminated if the tools were used in similar but non-functional demonstrations.

Thus, contrastive categorical information, whether encoded spontaneously or as a result of experimental manipulations, enables young infants to succeed at individuation tasks.

One-event task.—Young infants succeed at a standard task, even when tested with two objects they assign to the same category and view as merely featurally distinct, if instead of being presented with an occlusion event followed by a no-occlusion event, they are presented with a continuing occlusion event, because a transparent panel stands behind the screen (Fig. 2C). Infants thus see an ongoing occlusion event involving first an opaque and then a transparent occluder. To date, the only published evidence that young infants succeed at one-event tasks has come from a brief report by Wilcox and Chapa (2002). In one

experiment, 9-month-olds saw a block and a ball in alternation (different-objects event) or the same ball (same-object event) emerge on either side of a screen. When the center portion of the screen was lowered, it left a thin frame surrounding a transparent panel, through which the ball could be seen. Infants looked significantly longer at this final display if they had seen the different-objects event, suggesting that they expected to see both the block and the ball behind the panel. This effect was eliminated, however, if the transparent panel was absent, leaving only an empty frame.

These results suggest that *during* an occlusion event, young infants can establish separate representations for two objects even when encoded as merely featurally distinct, and they detect a violation if one of the objects spontaneously disappears. Apparently, it is only when the event comes to an end that infants have difficulty forming an expectation about how many objects should be present.

Remainder task.—In the standard and search tasks discussed earlier, when the first event involved two different objects, it always ended with both objects out of view (i.e., behind the screen in the case of standard tasks or inside the box in the case of search tasks; see Fig. 1). Young infants succeed at these tasks, even when tested with two objects they view as merely featurally distinct, if only one of the objects remains hidden when the first event ends (Fig 2D; Wilcox & Baillargeon, 1998; Wilcox & Schweinle, 2002; McCurry, Wilcox, & Woods, 2009).

To illustrate, 5–9-month-olds first saw either a block (different-objects event) or a ball (same-object event) move along a platform and disappear behind a screen. Next, the ball emerged on the other side of the screen and paused in full view. Finally, the screen was lowered to reveal no block—only the ball was visible to the right of the screen. Infants looked significantly longer at this final display if they had seen the different-objects event, suggesting that they expected to see the block when the screen was lowered (Wilcox & Baillargeon, 1998; Wilcox & Schweinle, 2002). In another experiment in this series (McCurry et al., 2009), the screen was made of cloth fringe, and the occlusion event was followed by a manual-search event, instead of a no-occlusion event: After the ball emerged from behind the screen and paused in full view, the platform was moved within infants' reach, and duration and frequency of reaching were measured. On both measures, infants reached significantly more for the screen if they had seen the different-objects event (suggesting that they believed an object was hidden behind the screen), but they reached significantly more for the visible ball if they had seen the same-object event (suggesting that they believed it was the only object present).

Similar results were obtained with 10-month-olds in a modified remainder search task (Xu & Baker, 2005). In each trial, infants first saw a same-object containment event: The same object (e.g., a toy duck) was removed from and returned to a box several times. Next, infants were allowed to search the box and retrieved either that object (e.g., the duck; same-object trial) or a different object (e.g., a ball; different-objects trial). Duration of search after this first retrieval was measured. Infants searched the box significantly longer on different-objects trials, suggesting that they realized the retrieved object was not the one they had seen previously, which must therefore still be in the box.

Thus, young infants can succeed at an individuation task involving two objects they view as merely featurally distinct as long as only one of the objects is hidden when the first event ends.

1.3. Prior Explanations for Individuation Findings

The evidence reviewed in the two preceding sections provides a bewildering array of negative and positive findings, as infants fail to detect individuation violations in some tasks but succeed in other, superficially similar tasks. Several different explanations have been proposed for these findings, though none so far has been successful at accounting for all the findings we have discussed. One suggestion has been that infants succeed at individuation tasks when they can assign the two hidden objects to distinct kinds (i.e., categories that allow rich inductive inferences; Carey, 2011; Gelman, 2003; Xu, 2007). Because infants understand that objects cannot change kinds, then upon noticing that the second object to emerge into view belongs to a different kind than the first object, they infer that two objects are present, and they expect both to be revealed (Carey, 2011; Xu, 2007; Xu & Carey, 1996; Xu et al., 2004). A related suggestion has been that infants succeed when they can assign the two hidden objects to distinct categories because this reduces processing-load demands: Representing the objects as members of separate categories provides “symbols that can be directly placed in short-term memory” (Xu & Carey, 2000, p. 294) or “makes it possible to store them in short-term memory without having to preserve their idiosyncratic features” (Futó et al., 2010, p. 1). One difficulty with these suggestions is that they cannot explain why infants detect individuation violations in one-event and remainder tasks based on featural differences alone. Another suggestion has been that infants succeed at individuation tasks that require monitoring object representations within a single event, but they fail at more demanding tasks that require mapping object representations from one event to the next (Needham & Baillargeon, 2000; Wilcox & Chapa, 2002; Wilcox & Baillargeon, 1998; Wilcox, Schweinle, & Chapa, 2003). One difficulty with this suggestion is that it does not explain why infants succeed at different-categories and remainder tasks that require mapping object representations across events.

In this article, we propose a new model of early individuation that accounts for all of the negative and positive individuation findings we have reviewed. Before outlining our model, we briefly discuss three types of information adults use to individuate objects; as will become clear, at the heart of our model is the notion that infants use these same types of information but cannot always resolve conflicts between them, leading to individuation failures. In the following section, we describe the main assumptions of our model and explain how they help make sense of infants’ responses in different individuation tasks. Next, we briefly review evidence from other infant-cognition tasks that provides support for our assumptions. Finally, we report three experiments that tested and supported key predictions from our model.

2. Object Individuation in Adults

Imagine that adults are shown a simple individuation scene: An object disappears on one side of a large screen, and then an object appears on the opposite side of the screen. To

decide whether the objects on either side of the screen are the same object or two separate objects, adults would use at least three types of information.

Spatiotemporal information.

Two objects viewed at different times tend to be perceived as numerically identical if a continuous spatiotemporal trace can be established between successive views; if no such trace can be established, then the objects are perceived as numerically distinct (Flombaum, Scholl, & Santos, 2009; Pylyshyn, 1989, 2007; Rips, Blok, & Newman, 2006). Thus, in our simple scene, spatiotemporal information consistent with a continuous trace would suggest a single object, whereas spatial discontinuities (e.g., the object did not appear in a gap at the center of the screen) or temporal discontinuities (e.g., the object reappeared too soon, given its current speed, to have travelled from one side of the screen to the other) would suggest two objects.

Object-representation information.

Even if the spatiotemporal information available in our simple scene was consistent with a continuous trace, adults could still posit two objects if they noticed that the object that emerged from behind the screen was different from the one that had disappeared behind it. When viewing an object, adults build a temporary representation of the object that includes both categorical (i.e., more general or abstract) and featural (i.e., more specific or detailed) information (Archambault, O'Donnell, & Schyns, 1999; Gordon & Irwin, 1996, 2000; Kahneman, Treisman, & Gibbs, 1992; Rips et al. 2006; Simons & Levin, 1998). Thus, in our scene, a temporary representation would be created for the object that disappeared behind the screen. When the object came back into view, it would be compared with the contents of this temporary representation; if similar, it would be viewed as the same object; if different, it would be viewed as a numerically distinct object, and a new representation would be built for it.

Would *any* change to the object lead to the creation of a new representation? In the initial research on change detection in adults, it was assumed that adults build detailed representations of attended objects and can use all of the information in these representations to detect changes to the objects. Over time, however, it became clear that adults often fail to detect even salient changes to attended objects that go briefly out of view, in both laboratory and real-world settings (Archambault et al., 1999; Rensink, 2002; Levin & Simons, 1997; Simons, Chabris, Schnur, & Levin, 2002). In our simple scene, it seems likely that adults would detect most categorical and featural differences between the objects shown on either side of the screen; under more challenging conditions, however, adults might well fail to do so. In an experiment inspired by the work of Xu and Carey (1996), Simons and Levin (1998) embedded an occlusion event in a novel social interaction on a college campus. An actor who carried a map and was dressed as a construction worker (e.g., in a plain hard hat, black shirt, and white pants) approached individual students and asked for directions. In each case, the interaction between the actor and the student was interrupted by two confederates who passed between them, carrying a door. While occluded, the actor surreptitiously switched positions with one of the confederates, another young White man who also carried a map and was dressed as a construction worker, though in different clothing (e.g., a hard hat with a

logo, a tool belt, a light blue shirt, and tan pants). Most subjects failed to notice the change to the actor.

Several different factors may cause adults to overlook changes to attended objects in a scene, including failure to form sufficiently detailed representations of the objects, failure to retain these representations over time, and failure to compare the pre- and post-change objects (Beck & Levin, 2003; Hollingworth, 2003; Noë, Pessoa, & Thompson, 2000; Mitroff, Simons, & Levin, 2004; Simons & Ambinder, 2005; Varakin, Levin, & Collins, 2007). Yet another factor, which is particularly germane to our model of early individuation, is that adults may compare the pre- and post-change objects but (particularly in challenging situations, where information processing resources are taxed) focus on only a small subset of the available information. In many cases, this subset will correspond to the objects' categorical descriptors; provided these are maintained across views, adults will fail to notice changes to the objects' individual features (unless, of course, these changes are perceptually highly salient). When applied to the experiment described above, this last factor suggests that adults selectively compared the pre- and post-change actors' categorical descriptors (e.g., young, White, male construction worker requesting directions); because these descriptors remained constant across views, subjects mistakenly inferred object continuity.

Evidence consistent with this suggestion comes from experiments showing preserved representations of pre-change objects' individual features despite failures to detect changes to these features (e.g., Angelone, Levin, & Simons, 2003; Hollingworth, Williams, & Henderson, 2001; Mitroff et al., 2004; Simons, et al., 2002). In one experiment (Angelone et al., 2003), for example, adults watched a videotaped event in which a female actor, who was wearing a basketball jersey and an art smock and was carrying a basketball and an art portfolio, asked someone for directions. When the camera cut to a different angle, she was replaced with another actor of the same age, race, and gender who wore the same clothes and carried the same objects. Only about half of the subjects noticed the change to the actor. Next, all subjects were asked to pick out the pre-change actor from a four-choice photographic lineup; strikingly, subjects who failed to notice the change were significantly above chance at selecting the pre-change actor, and their performance did not differ from that of subjects who did notice the change. Here again, one possible interpretation of these findings is that subjects who missed the change selectively compared the pre- and post-change actors' categorical descriptors (e.g., young, White, female student requesting directions); because these descriptors remained constant across views, they mistakenly inferred object continuity, despite a preserved representation of the pre-change actor.

Together, these results suggest that in some cases at least, adults may overlook a featural change to an object even though the information necessary to detect this change has been encoded, is maintained, and can be accessed experimentally via photographic lineups, probing questions, or more implicit measures (Angelone et al., 2003; Hollingworth et al., 2001; Mitroff et al., 2004; Simons et al., 2002). Such change-detection errors occur because adults focus exclusively on the object's categorical descriptors when checking for object continuity across views and falsely conclude they are facing the same object if these descriptors remain the same.

Causally relevant event-specific featural information.

Even if the spatiotemporal and categorical information available in our simple scene both signaled a single object moving back and forth behind the screen, adults could still posit two objects if they were led by their physical knowledge (Rips et al., 2006; Strickland & Scholl, 2015) to attend to causally relevant featural information and to notice changes incompatible with object continuity.

Adults possess a great deal of knowledge about simple causal interactions between objects (e.g., occlusion, containment, collision, and support events) and about simple causal transformations (e.g., cutting, compressing, painting, baking, melting, and burning events); in particular, they know which features are causally relevant for monitoring the progress of different events and for predicting their outcomes. Because adults' physical knowledge guides inferences about what changes to objects are possible or impossible in a specific event, it can also contribute to their ability to individuate objects. To illustrate, consider the following example, adapted from Rips et al. (2006). Imagine that an ice cube is placed behind a screen and, after some time, is taken out again. When taken out, the ice cube is either smaller or larger than before. Based solely on the spatiotemporal and categorical information available, one might have little reason to distinguish between these two outcomes. Based on a causal analysis of the scene, however, one might infer that the smaller ice cube was the same one as before (it had begun to melt), but that the larger ice cube was a different one than before (ice cubes cannot grow larger at room temperature). Thus, adults' physical knowledge can guide their attention to causally relevant featural information (e.g., size information) and can help them determine whether changes across views are compatible with object continuity.

Summary.

The evidence and arguments reviewed in this section suggest that adults may use spatiotemporal information, object-representation information, and causally relevant event-specific featural information to individuate objects in a scene. As we discuss in the next section, our model of early individuation assumes that young infants use the same three types of information, but can be stymied when conflicts arise between them.

2.2. A New Model of Object Individuation in Young Infants—As a brief overview, our model has three key assumptions. First, we assume that two cognitive systems help infants individuate objects in physical events: the object-file (OF) system (Kahneman et al., 1992) and the physical-reasoning (PR) system (Baillargeon et al., 2009a). In infants, the OF system uses primarily spatiotemporal and categorical information for this purpose, whereas the PR system also uses whatever featural information it has identified as causally relevant for the event category involved. Because the two systems rely on somewhat different information to individuate objects in physical events, conflicts can arise between them. Second, during an event, the PR system has priority over the OF system in predicting how the event will unfold. Third, when an event comes to an end, infants successfully track the objects to the next event if the two systems agree on how many objects are present (e.g., both systems agree that there are two objects behind the screen or inside the box). If the two systems disagree, infants' individuation performance depends on the nature of this

disagreement. When the two systems disagree quantitatively, in that both infer that objects are present but disagree on their number (e.g., the OF system signals that one object is present, whereas the PR system signals that two objects are present), this disagreement cannot be resolved and results in a catastrophic individuation failure: Infants hold no expectation at all about how many objects will be present in the next event. However, if the two systems disagree only qualitatively (e.g., the OF system signals that no object is present, whereas the PR system signals that one object is present), this disagreement can be resolved, leading to successful individuation.²

Below, we describe our model more fully. For ease of communication, we present each assumption in turn and discuss its implications for infants' performance in individuation tasks.

2.1. Two Systems—We assume that two different systems help infants individuate and track objects from event to event: the OF and PR systems.

OF system.

The overarching goal of the OF system is that of representing “where” and “what” information about objects, drawing on incoming perceptual information as well as on stored knowledge. The OF system thus includes two separate yet closely related processes, one that represents spatiotemporal information (*spatiotemporal* process), and one that represents categorical and featural information (*object-representation* process).³

When objects emerge in alternation on either side of a screen, the spatiotemporal process specifies that a single object is present if it can establish a continuous trace for this object; if it cannot, then it specifies that two objects are present. The object-representation process encodes both categorical and featural information about the objects seen on either side of the screen, but (in young infants at least) it relies primarily on the categorical information to decide whether these are the same object or two separate objects (this selective reliance on categorical information could be due to several factors, which we consider in the General Discussion). Thus, the object-representation process specifies that one object is present if it assigns the same categorical descriptors each time an object comes into view, and it specifies that two objects are present if it assigns different categorical descriptors.

Both processes gradually develop during infancy, resulting in the use of more detailed spatiotemporal information (Huttenlocher & Lourenco, 2007; Wilcox & Schweinle, 2003)

²In our previous models of early individuation, we focused on conflicts between categorical and featural information within the PR system (Baillargeon et al., 2012) and on conflicts (due largely to working-memory limitations) between the object-tracking and the PR systems (Levine & Baillargeon, 2016). The new model proposed here does a better job of explaining infants' successes and failures at various individuation tasks, relates better to findings on adults' ability to individuate and track objects, and provides a sounder basis for the predictions tested in the present experiments.

³Broadly speaking, accounts of object perception differ in how they integrate spatiotemporal (“where”) and object-representation (“what”) information. In object-file accounts, object-representation information is stored in an object's file, and spatiotemporal information is used to *address* this file (Gordon & Irwin, 1996; Kahneman et al., 1992). In object-tracking accounts, an object is assigned an index or spatiotemporal pointer that “sticks” to it as it moves, and object-representation information can be *bound* to this index (Leslie et al., 1998; Pylyshyn, 1989). The debate about how best to integrate spatiotemporal and object-representation information is far from settled (Pertzov, Dong, Peich, & Husain, 2012; Naughtin, Mattingley, & Dux, 2016). For present purposes, we gloss over these distinctions and treat the spatiotemporal and object-representation processes as separate yet closely related processes that both reside within the OF system.

and more fine-grained categorical descriptors (Pauen, 2002; Xu & Carey, 1996) over time. For example, at about 4.5 months of age, the spatiotemporal process begins to use speed information, in addition to path information, to individuate objects. If an object disappears at one end of a screen and immediately reappears at the other end (i.e., too soon, given its current speed, to have traveled the distance behind the screen), infants age 4.5 months and older infer that two objects are present, but younger infants do not (Wilcox & Schweinle, 2003). Similarly, at about 12 months of age, as discussed earlier, the object-representation process begins to use basic-level descriptors, in addition to broad ontological descriptors, to individuate objects (Xu & Carey, 1996).

Starting early in life, disagreements between the spatiotemporal and object-representation processes are resolved in favor of whichever process posits two objects. Thus, 4-month-olds infer that two objects are present when an object disappears behind one screen and reappears from behind a different screen without appearing in the gap between them: The object-representation process posits one object because it assigns the same categorical descriptors each time the object comes into view; the spatiotemporal process posits two objects because it detects the path discontinuity in the event; and the spatiotemporal process overrides the object-representation process, resulting in successful individuation (Aguiar & Baillargeon, 2002; Spelke Kestenbaum, Simons, & Wein, 1995). Similarly, as was discussed earlier, 4-month-olds infer that two objects are present when objects previously shown to belong to two distinct functional categories emerge in alternation on either side of a screen: The spatiotemporal process posits one object because it can establish a continuous trace for it; the object-representation process posits two objects because it assigns different categorical descriptors to the objects on either side of the screen; and the object-representation process overrides the spatiotemporal process, resulting in successful individuation (Stavans & Baillargeon, 2018). This explains the positive results of the different-locations and different-categories tasks discussed earlier (Fig. 2A and 2B): In each case, the process that specifies two objects overrides the process that specifies only one object.

From the preceding descriptions, it follows that when two objects infants encode as merely featurally distinct emerge in alternation from behind a screen, the spatiotemporal and object-representation processes mistakenly infer that a single object is present—the spatiotemporal process because it can establish a continuous trace for the object, and the object-representation process because it assigns the same categorical descriptors each time the object comes into view (the object’s “changing” individual features are simply updated at each new emergence). This can lead to disagreements with the PR system, as we discuss next.

PR system.

The PR system is a causal-reasoning system that becomes engaged whenever infants observe or produce a physical interaction between objects; the PR system then attempts to predict and interpret the progress of this interaction over time (Baillargeon et al., 2009a, 2012; Mascialoni et al., 2013; Wang & Goldman, 2016). Even simple physical interactions may call for a whole slew of predictions, and infants must become adept at these predictions both to keep up with the world as it unfolds and to act adaptively on objects.⁴

When infants see objects emerge in alternation on either side of a screen, the PR system begins to build a specialized representation of this event. It first uses the spatiotemporal and categorical information from the OF system to categorize the event as an occlusion event and to assign event roles (e.g., occluder, occludee). Next, the PR system taps the OF system for information about the features that have been identified as causally relevant to occlusion events (e.g., size and shape by 4 months, pattern by 7 months, and color by 11 months; Wilcox, 1999). This selected featural information is then added to the event's representation. All of the information included in the representation—the spatiotemporal, categorical, and causally relevant featural information—is used by the PR system's physical knowledge to predict and interpret the progress of the event.

If similar categorical and causally relevant featural information is included for the objects that emerge on either side of the screen, the PR system concludes that a single object is involved. This conclusion will be correct in many cases, but it will be incorrect when the event involves two objects that differ only in features not yet identified as causally relevant to occlusion events (e.g., when infants under 11 months, who have not yet identified color as an occlusion feature, see two balls that differ only in color emerge alternately from behind the screen; Wilcox et al., 2012; Wilcox, Haslup, & Boas, 2010).

If similar categorical information, but different causally relevant featural information, is included for the objects that emerge on either side of the screen, the PR system must decide whether these featural differences are compatible with the presence of a single object or whether they signal the presence of two distinct objects. In some cases, the PR system may conclude that a single object is present. For example, infants realize that a soft, deformable object may change shape when out of view (Schaub et al., 2013), and that a self-propelled object may re-arrange the orientation of its parts when out of view (e.g., Baillargeon et al., 2009b). In many cases, however, such as those commonly studied in individuation tasks, the PR system infers that two objects are present. For example, if a large ball and a small ball emerge in alternation from behind a screen, the PR system infers (beginning at about 4 months of age, when size has been identified as an occlusion feature) that two different balls are involved in the event. According to the persistence principle, a large ball cannot spontaneously change into a small ball when passing behind a screen (Baillargeon, 2008).

From the preceding descriptions, it should be clear that conflicts will arise between the OF and PR systems when (a) the OF system mistakenly infers that a single object is present in an event, based on the spatiotemporal and categorical information at its disposal, but (b) the PR system correctly infers that two objects are present, based on the causally relevant featural information at its disposal.⁵

⁴.To clarify what we mean by predicting and interpreting the progress of a physical event, imagine that we saw an object being lowered toward the opening of a container. As the event unfolded, we might consider: whether the object would fit through the opening of the container; whether the object, when resting inside the container, would remain partly visible above its rim or through its sidewalls; whether the object would remain inside the container if the latter was slid to the side, lifted, or turned upside-down; and whether the object, when removed from the container, had the same appearance as before or a different appearance. From a young age, infants begin to master these various predictions about containment events (Baillargeon, Li, Luo, & Wang, 2006; Baillargeon et al., 2009a; Hespos & Baillargeon, 2001a, 2001b; Wang et al., 2004, 2005; see also Strickland & Scholl, 2015, for related results with adults).

⁵.Readers may wonder why we describe the OF and PR systems as separate systems in our model. Could the two instead correspond to distinct operations or computations within a single system? We think not, in part because of how we generally construe causal

2.2. During an Event—While a physical event is unfolding, the PR system has priority over the OF system in predicting and interpreting the progress of the event. This means that if a disagreement arises between the OF and PR systems, as was just discussed (e.g., the OF system mistakenly infers that a single object is present behind a screen, whereas the PR system correctly infers that two objects are present), infants may still give evidence that they have established separate representations for the two objects as long as these representations are probed *during* the event itself. Because the PR system has control over infants' responses while the event is unfolding, infants can succeed at detecting individuation violations. This explains the positive result of the one-event task devised by Wilcox and Chapa (2002). Recall that if at the end of a different-objects event the screen is lowered to reveal only one of the objects standing behind a transparent panel (Fig. 2C), young infants detect this individuation violation. The PR system (a) uses the causally relevant featural information at its disposal to infer that two objects are present, (b) monitors the occlusion event as it continues to unfold, first with the opaque and then with the transparent occluder, and (c) detects an individuation violation when one of the two objects spontaneously disappears.

2.3. When an Event Ends—When the OF and PR systems agree on how many objects are present as an event comes to an end, infants successfully track the objects to the next event. When the systems disagree, however, performance depends on the nature of this disagreement: Qualitative disagreements (e.g., none vs. one) can be resolved, but quantitative disagreements (e.g., one vs. two) cannot, leading to catastrophic individuation failures.

To illustrate various situations in which the OF and PR systems might agree or disagree (Fig. 3), imagine that infants ages 4–12 months are watching an occlusion event. If the same object disappears and reappears from behind the screen (Fig. 3A), both systems agree that a single object is present, as there is no information to suggest otherwise. If at some point during the occlusion event two objects emerge simultaneously from behind the screen, as in a different-locations task, both systems agree that two objects are present, based on the spatiotemporal information available (Fig. 3B). If two objects appear alternately on either side of the screen, and these objects receive different categorical descriptors, as in a different-categories task, both systems again agree that two objects are present, based on the categorical information available (Fig. 3C). However, if the two objects receive identical categorical descriptors and differ in one or more features identified as relevant to occlusion events (Fig. 3D), then the two systems disagree. The OF system assumes that a single object is present, based on the spatiotemporal and categorical information available, whereas the PR system assumes that two objects are present, based on the causally relevant featural information available. This quantitative disagreement (one vs. two objects behind the screen)

reasoning in core domains. We assume that in the domains of physical, psychological, and sociomoral reasoning, a separate skeletal framework of principles and concepts guides how infants represent, reason about, and learn about events in each domain (Baillargeon et al., 2015; Buyokozer Dawkins, Ting, Stavans, & Baillargeon, in press). From this perspective, it seems unlikely that the OF and PR system would be part of a single system. The OF system creates temporary representations for objects, whether in events or in static scenes (e.g., it can represent trees in a yard, dots on a page, and so on). In contrast, the PR system only becomes involved when objects interact; its main function is that of predicting how this causal interaction will unfold and, when predictions fail, that of revising its domain knowledge to better predict outcomes in the future. Because the two systems have such different purposes and requirements, it seems unlikely that they would be components of the same system. Nevertheless, the two systems must certainly interact a great deal, and our model focuses on specific difficulties that arise in the context of these interactions.

results in a catastrophic individuation failure, leading infants to hold no expectation about how many objects should be revealed when the screen is removed. This explains the negative results of the standard and search tasks described earlier (Fig. 1).

One exception to these catastrophic failures was discussed in the last section and has to do with one-event tasks (Fig. 2C). Because a transparent panel stands behind the screen, the occlusion event does not come to an end. The PR system continues to monitor the ongoing event and detects an individuation violation if one of the objects spontaneously disappears (Fig. 3E).

Another exception to these catastrophic failures has to do with remainder tasks (Fig. 2D). The OF system assumes, based on the spatiotemporal and categorical information at its disposal, that the occlusion event involves only one object that is now resting in plain view next to the screen, thus leaving nothing behind the screen. In contrast, the PR system assumes, based on the featural information at its disposal, that the occlusion event involves two objects, the one that is resting next to the screen and another one that is still hidden behind the screen. Because the disagreement between the two systems is qualitative (no object vs. one object behind the screen), it can be resolved, leading infants to expect one object when the screen is removed (Fig. 3F).⁶

In sum, catastrophic individuation failures occur when an event comes to an end and there is a quantitative conflict between the OF and PR systems: The two agree that objects are present behind the screen (standard tasks) or inside the box (search tasks), but disagree on their number. In contrast, no individuation failure occurs when the conflict between the two systems is qualitative in nature: The OF system assumes that there is nothing left behind the screen or inside the box, whereas the PR system assumes that one or more objects remain hidden there.

Why do quantitative conflicts result in catastrophic individuation failures, but qualitative conflicts do not? First, let us consider what might happen when a quantitative conflict arises. Imagine that the OF system signals that there is one object present behind a screen, but the PR system signals that there are in fact two objects present. How does the OF system respond to this information? One possibility is that it entirely *discards* its object file. After all, an object file is meant to refer to a particular object in the world and to hold information about its properties—it can never refer to two or more objects. The OF system may thus discard its faulty object file and start over (as though “rebooting”), resulting in a catastrophic individuation failure. Another possibility is that the OF system *attempts to split* its one object file into two separate files, one for each hidden object. However, the processing demands associated with this split may exceed young infants’ limited information-

⁶Readers may wonder at the difference between different-objects trials in search and remainder tasks. After all, after infants retrieve the first object in a search task, doesn't the task essentially become a remainder task? The answer is no. At the end of the containment event in the search task, the OF and PR systems disagree quantitatively about how many objects are present; as a result, infants hold no expectation about what they will see or do next, and hence they have no expectation to update as events ensue. At the end of the containment event in the remainder task, however, the OF and PR systems agree that a single object is present in the box, thus leading infants to hold a clear expectation for one object. When a featurally distinct object is removed from the box, the two systems update their representations. The OF system assumes there is now nothing left in the box, while the PR system infers that the original object must still be in the box. Because this amounts to a qualitative disagreement (no object vs. one object in the box), it can be successfully resolved, leading infants to search for the original object.

processing resources, resulting once again in a catastrophic individuation failure (akin to an “error message”).

Now, consider what might happen when a qualitative conflict arises. Imagine that the OF system signals that there is one object present, resting next to a screen, but the PR system signals that there are in fact two objects present, the one next to the screen and another one behind it. How might the OF system respond? In this case, there is no single object file erroneously pointing to two different objects; the OF system’s object file is pointing to the object next to the screen, and the PR system is signaling that in addition to that object, there is another object elsewhere in the scene, behind the screen. In response to this signal, there is no need for the OF system to discard its object file or to split it in two—it can simply *add* a new file pointing to the object behind the screen, resulting in successful individuation.

3. Findings from Other Tasks Supporting Our Model

In this section, we discuss a few findings from infant-cognition tasks that support the assumptions of our model. A full consideration of the research related to the OF and PR systems is beyond the scope of this article, so our discussion is illustrative, rather than exhaustive.

3.1. Two Systems, with Priority to the PR System during Physical Events

We have seen that when the PR system represents a physical event and has identified a given feature as causally relevant for predicting outcomes in the event category involved, it requests information about that feature from the OF system. In such cases, both the OF and the PR systems end up having access to the information about the feature. If the PR system has not yet identified the feature as relevant, however, there is a dissociation between the two systems: The OF system possesses featural information that the PR system lacks. Because the PR system is in charge of predicting how the event will unfold, this means that infants will fail to detect violations involving the feature, even though the OF system has encoded the necessary information for detecting these violations.

Evidence for these claims comes from experiments by Wang and her colleagues (Wang & Goldman, 2016; Wang & Mitroff, 2009). These experiments built on prior findings that the feature height is identified at different ages in covering events (12 months) and tube events (14 months) (Wang & Baillargeon, 2006; Wang et al., 2005). In one experiment (Wang & Goldman, 2016), 12-month-olds saw an experimenter’s hand lower a tall cover (cover condition) or a tall tube (tube condition) over a short object. Next, the hand lifted the cover or tube to reveal either the same object as before (no-change event) or a much taller object (change event). As expected, infants in the cover condition detected the change to the object’s height, whereas those in the tube condition did not. Strikingly, however, infants did detect this change in a modified-tube condition in which they were briefly turned away to face a plain wall while the tube was lowered over the object and then lifted back again. Wang and Goldman (2016) concluded that (a) the OF system encoded the object’s height in all conditions and events; (b) when infants observed the interaction between the cover or tube and the object, the PR system became involved and had priority in guiding infants’ expectations, so infants detected the change to the object’s height in the cover condition

(height information was included in the PR system's event representation), but not the tube condition (no height information was included); and (c) when infants did not observe the interaction between the tube and the object in the modified-tube condition, the PR system was not involved, so the OF system alone guided infants' responses, leading to enhanced attention to the novel object shown in the change event.

In another experiment, Wang and Mitroff (2009) first showed 11-month-olds a covering event in which a short object was much taller after being hidden under a tall cover (infants failed to detect this change violation, in line with prior findings). Following the covering event, infants saw static displays pairing either the pre-change (short) or post-change (tall) object with a novel, medium-sized object. Infants looked preferentially at the medium-sized object, and this effect was eliminated if they were shown only the static displays. These results suggest that (a) the OF system had encoded the heights of the pre- and post-change objects during the covering event and viewed the medium-sized object as relatively novel, and (b) infants could exhibit this novelty preference because the PR system was not involved during the static displays.

The preceding results support several assumptions of our model: The OF and PR systems are two distinct systems; while a physical event is unfolding, the PR system is in charge of predicting and interpreting the progress of the event; and the OF system readily passes on the categorical and spatiotemporal information at its disposal to the PR system (e.g., whether the hand is holding a cover or a tube, and whether it is lowering it next to, behind, or over the object), but it supplies information about individual features (e.g., height information) only if the PR system specifically requests it.

3.2. Other Catastrophic Failures

Readers familiar with the research on early numerical cognition may have noticed intriguing parallels between that research and the individuation model and findings we have described. In a series of forced-choice experiments (Feigenson, Carey & Hauser, 2002; Feigenson & Carey, 2005; vanMarle, 2013), 10–12-month-olds watched as an experimenter placed crackers one-by-one into two containers; one container always had more crackers than the other. Infants approached the container with more crackers in trials with 1 vs. 2, 2 vs. 3, and 4 vs. 8 crackers, but not in trials with 1 vs. 4, 3 vs. 6, or 2 vs. 8 crackers, where they selected a container at random.

One possible interpretation of these results (vanMarle, 2013) is that infants experienced a catastrophic numerical-comparison failure when forced to compare representations produced by two distinct systems: the parallel-individuation system, which focuses on individual items and can represent up to three simultaneously, and the approximate-number system, which focuses on groups of items and represents their approximate cardinal value (Feigenson, Dehaene, & Spelke, 2004; Hyde & Spelke, 2011). Thus, infants succeeded when they could compare representations from the same system (e.g., parallel-individuation: 1 vs. 2, approximate-number: 4 vs. 8), but failed when they had to compare representations from the two systems (e.g., 1 vs. 4, 2 vs. 8). Consistent with this interpretation, additional experiments (vanMarle, Mou, & Seok, 2016) showed that when comparing representations

from the approximate-number system alone, 10–12-month-olds succeeded when the two sets of crackers differed by a 2:3 ratio (e.g., 4 vs. 6), but not a smaller ratio (e.g., 7 vs. 8).

Together, the preceding results make it unlikely that infants failed with sets of 1 vs. 4 or 2 vs. 6 crackers because they could not represent sets of 4 or 6 crackers, could not perform comparisons involving such large sets, or could not compare sets differing by a 1:4 or a 1:3 ratio (infants succeeded with a smaller, 2:3 ratio). Rather, it seems more likely that infants failed because of difficulties in comparing quantitative representations from different systems. Our model of early individuation also focuses on failed interactions between two systems, although the ones we describe concern disagreements about how many objects are involved in an event, rather than “incommensurable” representations of smaller and larger quantities (vanMarle, 2013). In both cases, however, these failed interactions result in catastrophic failures, with infants holding no expectation at all about what they should see or do next.

4. The Present Research

In three experiments with infants ages 9–11 months, we tested key assumptions of our model. All three experiments bore on our first assumption: The OF uses primarily spatiotemporal and categorical information to individuate objects, whereas the PR system also uses causally relevant featural information; as a result, the two systems sometimes posit different numbers of objects in an event. Our experiments tested the consequences of these disagreements for infants’ performance in different individuation tasks.

Experiment 1 focused on our assumption that *event boundaries* matter in infants’ individuation performance. When an event comes to an end, agreement between the OF and PR systems is critical for infants to hold a clear expectation about how many objects are present; during the event itself, however, the PR system has priority in predicting how the event will unfold, so disagreements with the OF system do not matter. To support this assumption, we sought to confirm the finding (never replicated to date) that young infants succeed at one-event tasks even when tested with two objects they view as merely featurally distinct (Wilcox & Chapa, 2002). In a different-objects event, infants saw two different dolls emerge in alternation from behind a screen, which was then lowered to reveal only one of the dolls. For half of the infants, this doll rested behind a transparent panel (*panel* condition); for the other infants, no panel was used (*no-panel* condition). According to our model, infants in the panel condition should detect the individuation violation they were shown, but infants in the no-panel condition should not (as usually found). In the panel condition, the occlusion event did not come to an end when the screen was lowered; instead, it continued on with the panel now serving as the occluder. During this event, the PR system had priority and could infer, based on the causally relevant featural information at its disposal, that two objects were present; infants should therefore detect a violation when one of the objects spontaneously disappeared. Thus, strikingly, the mere addition of a transparent panel should enable infants to detect the violation in the different-objects event, thereby supporting our analysis of one-event tasks and our model more generally.

Experiment 2 focused on our assumption that *quantitative disagreements* between the OF and PR systems give rise to catastrophic individuation failures. When an event comes to an end and the OF and PR systems both posit hidden objects but disagree on their number, infants hold no expectation at all about how many objects will be present in the next event. This assumption leads to a striking prediction. When infants experience a catastrophic failure, they should not merely be uncertain as to whether one or two objects should be present in the next event: They should fail to detect a violation even if *no object* is present. To evaluate this prediction, infants were tested with a standard between-subjects task using two objects that were merely featurally distinct, a dotted and a striped cup. In one condition, infants saw the two cups in alternation (different-objects event) or the same cup (same-object event) emerge on either side of a screen. Next, the screen was lowered to reveal no cup. According to our model, infants who saw the same-object event should detect this violation: The OF and PR systems should agree that one object was present behind the screen, leading to a clear expectation for one object. In contrast, infants who saw the different-objects event should detect no violation: The OF (one object) and PR (two objects) systems should disagree about how many objects were present, leading to a catastrophic individuation failure. Thus, remarkably, infants who saw *more* objects in the occlusion event should be *less* likely to detect a violation when the screen was removed to reveal no object, thereby providing strong evidence for our model.

Experiment 3 again tested our assumption concerning *quantitative disagreements* and catastrophic failures, this time using a novel event sequence composed of a containment event followed by a shaking event. The objects were two rigid toys that infants this age would encode as merely featurally distinct, a toy wolf and a ball. Infants first saw the two toys in alternation (different-objects event) or the same toy (same-object event) lifted from and lowered into a large box. Next, the box was shaken briskly but produced no thumping sounds, as though empty (the box was many times larger than the two toys combined, so the absence of thumping sounds could not be attributed to the toys tightly filling the box). According to our model, infants who saw the same-object event should detect a violation during the silent shaking event: The OF and PR systems should agree that one object was present inside the box, leading to a clear expectation for thumping sounds as the object collided with the box's interior walls. In contrast, infants who saw the different-objects event should detect no violation: The quantitative disagreement between the OF (one object) and PR (two objects) systems about how many objects were present should lead to a catastrophic individuation failure. Thus, remarkably, infants who saw *more* objects in the containment event should be *less* likely to detect a violation when the shaken box seemed to be empty, thereby providing further evidence for our model.

Finally, Experiment 3 also tested our assumption about *qualitative disagreements*. When an event comes to an end and the OF system posits no hidden object but the PR system posits one hidden object, this qualitative disagreement (i.e., no object vs. one object) is readily resolved, leading infants to hold a clear expectation that one object is still present. To test this prediction, additional infants saw different-objects and same-object events identical to those described above, with one exception: In the containment event, the last object lifted from the box was visibly removed from the apparatus, instead of being returned to the box. In the case of the same-object event, the OF and PR systems should agree, when the

containment event ended, that the box was now empty, so infants should expect it to be silent when shaken. In the case of the different-objects event, in contrast, the OF system should infer that the box was empty, whereas the PR system should infer that it still held one object. This qualitative disagreement should then be resolved, leading infants to detect a violation when the shaken box was silent, as though empty. Such evidence would provide clear support for our assumption that quantitative and qualitative disagreements between the OF and PR systems have different consequences for infants' individuation performance.

5. Experiment 1

Experiment 1 built on prior results by Wilcox and Chapa (2002) and examined whether 10-month-old infants would succeed at detecting an individuation violation in a one-event task involving two objects they encoded as merely featurally distinct. Infants faced a puppet-stage apparatus and were randomly assigned to a panel or a no-panel condition (Fig. 4).

Infants in the *panel* condition received one orientation trial, one familiarization trial, and one test trial; each trial had an initial phase and a final phase. At the start of the (4-s) initial phase in the orientation trial, infants faced a large screen standing upright at the center of the apparatus floor; after a few seconds, the entire screen was lowered to reveal a transparent panel standing behind it (before the testing session, an experimenter also showed infants the panel, so they knew it consisted of a clear solid surface). During the final phase of the trial, infants saw the panel until the trial ended (see Procedure for criteria).

The familiarization trial introduced infants to the objects that would be used in the test trial. In their task, Wilcox and Chapa (2002) used a block and a ball (objects that 9-month-olds would encode as merely featurally distinct) and compared infants' ability to individuate and track these objects with and without a transparent panel. For a fuller investigation of this ability, we presented infants with one of three events: different-categories, different-objects, or same-object. Our stimuli were adapted from those of Bonatti et al. (2002), who found that 10-month-olds succeeded at a standard within-subject task when tested with a dog head and a doll head, but not when tested with two different doll heads. These and other results (Bonatti et al., 2005; Kibbe & Leslie, in press) suggested that young infants spontaneously assign distinct categorical descriptors to a toy animal (non-human) and a doll (human-like) shown one at a time, but not to two different dolls shown one at a time, even when they differ in skin, hair, and eye color. In line with these results, our stimuli consisted of various toy heads mounted on identical supports; across events, infants saw a toy animal and a doll (different-categories event), two different dolls (different-objects event), or a single doll (same-object event).

Thus, in the (28-s) initial phase of the *different-categories* familiarization event, infants first saw a toy animal (a light rabbit or a dark cat, counterbalanced) emerge to the right of the screen, pause, and then return behind the screen (6 s). After a 1-s pause, a doll (a light-skinned blonde or a dark-skinned brunette, with different hairdos, counterbalanced) emerged to the left of the screen, paused, and then returned behind the screen (6 s). After a 1-s pause, this sequence was repeated one more time, ending with another 1-s pause. Each object thus emerged into view twice in total. During the final phase of the trial, infants saw the upright

screen until the trial ended. The *different-objects* event was similar except that the two different dolls (counterbalanced) emerged on either side of the screen. Finally, the *same-object* event was again similar except that the same doll (counterbalanced) emerged on either side of the screen (for better experimental control, two identical dolls were used to produce the event).

In the test trial, infants saw the same event as in the familiarization trial, with one exception: During the last second of the (28-s) initial phase, the screen was lowered to reveal the last object to return behind the screen (i.e., the doll in the different-categories event, and the last-seen doll in the different-objects and same-object events), resting centered behind the transparent panel. During the final phase of the trial, infants watched this final display until the trial ended. Infants in the *no-panel* condition were tested using the same procedure as in the panel condition except that there was no orientation trial and no transparent panel in the test trial.

Our predictions were as follows (using the order established in Fig. 3 as a general guide). In the case of the same-object test event, both the OF and the PR systems should agree that a single object was present, as there was no information to suggest otherwise (Fig. 3A). Infants should therefore detect no violation when the screen was lowered to reveal that object, and this should be true with or without the transparent panel. In the case of the different-categories test event, the OF system should assign different categorical descriptors to the toy animal and the doll (non-human, human-like), so the OF and PR systems should agree that there were two objects present, as in prior different-categories tasks (Fig. 3C). Infants should therefore detect a violation when the screen was lowered to reveal only one of the objects, and this should again be true with or without the transparent panel. In the case of the different-objects test event, however, different predictions held for the two conditions. In the no-panel condition, the OF and PR systems should disagree about the number of objects present (Fig. 3D). The OF system should assign the same categorical descriptors to each doll and thus should infer that a single object was present.⁷ In contrast, the PR system should assume, based on the causally relevant featural information at its disposal (e.g., the distinct sizes and shapes of the dolls' hairdos), that two objects were present. This quantitative disagreement should lead to a catastrophic individuation failure, so infants should detect no violation when the screen was lowered to reveal only one object. In the panel condition, in contrast, the PR system should still be in charge of monitoring the occlusion event, which continued on with the transparent panel now serving as the occluder. Because the PR system assumed that two different objects were involved in the event, infants should detect an individuation violation when one of the objects spontaneously disappeared (Fig. 3E).

Together, these predictions suggested that (a) in the no-panel condition, infants would look significantly longer if shown the different-categories as opposed to the different-objects or

⁷Readers might wonder why infants would not encode the two different dolls as categorically distinct, given the apparent differences in their races (one doll was Caucasian, while the other doll's race was indeterminate but clearly non-Caucasian). After all, infants in the first year of life do form race-based categories: After seeing multiple exemplars from one race (e.g., White), they dishabituate to exemplars from another race (e.g., Black) (Anzures, Quinn, Pascalis, Slater, & Lee, 2010; Waxman & Grace, 2012). When shown a single exemplar from each racial category (as in Experiment 1), however, there is no evidence that infants encode them as members of distinct racial categories (Bonatti et al., 2002; Kinzler & Spelke, 2011).

same-object event; (b) in the panel condition, infants would look significantly longer if shown the different-categories or the different-objects event than if shown the same-object event; and (c) only the different-objects event would differ between the two conditions, with infants looking significantly longer when the transparent panel was used. Finding the predicted results would confirm those of Wilcox and Chapa (2002), would support our analysis of one-event tasks, and more generally would provide robust evidence for our model.

5.1. Method

5.1.1. Participants—Previous standard between-subjects tasks have used numbers of participants ranging from 6 to 14 per cell (i.e., per combination of condition and event; Stavans & Baillargeon, 2018; Wilcox & Baillargeon, 1998; Wilcox & Chapa, 2002; Wilcox & Schweinle, 2003). An a-priori power analysis using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007), based on the averaged Condition \times Event effect size ($\eta^2 = .148$) obtained by Stavans and Baillargeon (2018), indicated that with power set at .80 and alpha set at .05, the minimum cell size required was 10 participants for 2×3 designs (as in Experiments 1 and 3), and 12 participants for 2×2 designs (as in Experiment 2). For consistency, all experiments reported in this article used 12 infants per cell.

Participants were 72 full-term infants (36 male, $M = 10$ months, 27 days, range = 10;0–11;19). Another 16 infants were tested but excluded (9 in the no-panel condition and 7 in the panel condition), because they were active (6) or fussy (5), peeked behind the screen (4), or had a test looking time over 3 standard deviations from the condition mean (1; this infant saw the same-object event in the panel condition). Twelve infants were randomly assigned to each combination of condition (no-panel, panel) and event (same-object, different-objects, different-categories).

Infants' names in this and the following experiments were obtained from a university-maintained database of parents interested in participating in child development research. Written informed consent was obtained from each infant's parent prior to the testing session, and all protocols were approved by the Institutional Review Board at the University of Illinois at Urbana-Champaign.

5.1.2. Apparatus and Stimuli—The apparatus consisted of a brightly lit display booth (201 cm high \times 100 cm wide \times 46 cm deep) with a large opening (55 cm \times 95 cm) in its front wall; between trials, a supervisor lowered a curtain in front of this opening. Inside the apparatus, the sidewalls were painted white, and the back wall (made of foam board) and floor were covered with adhesive paper.

The screen (25 cm \times 30.5 cm \times 0.5 cm) was blue and was mounted on a thin rod that extended out of the right wall of the apparatus; at the end of the orientation and test trials, a hidden first assistant rotated the rod to lower the screen. When upright, the screen stood centered on the apparatus floor. The transparent panel (22.5 cm \times 30.5 cm \times 0.5 cm) was made of clear Plexiglas, was outlined with blue tape, and had a horizontal support at the back (0.5 cm \times 30.5 cm \times 5.5 cm) that allowed it to stand upright. When used, it was placed 3 cm behind the screen.

The objects consisted of six different toy heads mounted on identical green supports (12.5 cm × 10.5 cm × 10.5 cm) decorated with red dots. Two of the heads belonged to toy animals, a light brown rabbit (12 cm × 6 cm × 7.5 cm; all stimuli measurements are at the largest points) and a dark brown cat (5.5 cm × 8 cm × 6.5 cm). Two of the heads were identical and depicted a light-skinned and blue-eyed doll (6.5 cm × 14 cm × 8.5 cm), with long blonde hair decorated with blue streaks and worn in two braids, each tied with a blue hair tie. The remaining two heads were identical and depicted a dark-skinned and brown-eyed doll (5.5 cm × 12 cm × 5.5 cm), with long loose brown hair decorated with a light blue bow. During the familiarization and test events, the objects moved 2.5 cm in front of the back wall. At the back of each object was a thin flat handle that extended through a thin gap (concealed by white muslin fabric) at the bottom of the back wall; behind the wall, a hidden second assistant slid the handle along a track (with markings and stoppers) to move the object back and forth. To avoid producing sound cues that might indicate the number of hidden objects present, each object's bottom surface was covered with felt, and each object was moved silently along a clear plastic track glued to the apparatus floor. Finally, the back wall had a concealed trapdoor, centered behind the screen, which was used by a hidden third assistant to surreptitiously remove the penultimate object to return behind the screen at the end of the test trial.

During each testing session, a metronome beat softly to help the assistants adhere to the events' second-by-second scripts. One camera captured an image of the events, and another camera captured an image of the infant. The two images were combined, projected onto a monitor located behind the apparatus, and watched by the supervisor to confirm that the events followed the prescribed scripts. Recorded sessions were also checked off-line for accuracy.

5.1.3. Procedure—Infants sat on a parent's lap centered in front of the apparatus; parents were instructed to remain silent and to close their eyes during the test trial. Two observers hidden on either side of the apparatus monitored each infant's looking behavior; from their perspective, the observers could not determine which objects were used in the familiarization and test trials. Looking times during the initial and final phases of each trial were computed separately, using the primary observer's responses. Inter-observer agreement during the final phase of each trial was calculated by dividing the number of 100-ms intervals in which the two observers agreed by the total number of intervals in the final phase. Agreement was measured for 69/72 infants (only one observer was present for the other infants) and averaged 92% per trial per infant.

Infants were highly attentive during the initial phases of the trials they received; in the no-panel condition, they looked, on average, for 94% of each initial phase, and in the panel condition, they looked, on average, for 92% of each initial phase. The final phase of each trial ended when infants (a) looked away for 1 consecutive second after having looked for at least 2 (orientation and familiarization) or 10 (test) cumulative seconds, or (b) looked for a maximum of 30 cumulative seconds. A longer minimal looking value was used in the test trial to give infants the opportunity to process the final display revealed when the screen was lowered.

To reduce any positive skewness in the test data, all looking times in this article were log-transformed, and analyses were conducted on the log-transformed data (Csibra, Hernik, Mascaro, Tatone, & Lengyel, 2016). For ease of communication, however, raw looking times are provided in the text and figures. Preliminary analyses of the test data in Experiment 1 revealed no significant interaction of condition and event with either infants' sex or last-seen object (light-skinned blonde or dark-skinned brunette), both p s > 0.346; the data were therefore collapsed across these latter two factors.

5.2. Results

Looking times during the final phase of the *familiarization* trial were compared by means of an analysis of variance (ANOVA) with condition (no-panel or panel) and event (same-object, different-objects, or different-categories) as between-subjects factors. The analysis yielded no significant effects, all p s > 0.110, suggesting that across conditions and events, infants tended to look equally at the upright screen during the final phase of the familiarization trial (for descriptive statistics, see Experiment 1 Dataset in Supplementary Online Material, or SOM).

Looking times during the final phase of the *test* trial were analyzed as above (Fig. 5). The analysis yielded significant main effects of condition, $F(1, 66) = 9.14, p = 0.004$, and event, $F(2, 66) = 22.48, p < 0.001$, as well as a significant Condition \times Event interaction, $F(2, 66) = 3.33, p = 0.042, \eta^2 = 0.09$. Five planned comparisons (with a Bonferroni-adjusted α -level of $.05/5 = 0.01$) were conducted. In the no-panel condition, as predicted, infants who saw the different-categories event ($M = 23.48, SD = 5.01$) looked significantly longer at the final display than did infants who saw the different-objects event ($M = 13.87, SD = 3.19$) or the same-object event ($M = 14.53, SD = 4.79$), $F(1, 66) = 36.28, p < 0.001$. In the panel condition, again as predicted, infants who were shown the different-categories event ($M = 23.95, SD = 4.78$) or the different-objects event ($M = 19.78, SD = 5.42$) looked significantly longer at the final display than did infants who were shown the same-object event ($M = 16.17, SD = 3.56$), $F(1, 66) = 11.22, p = 0.001$. Comparing between conditions, significantly different looking times were found, as predicted, for the different-objects event, $F(1, 66) = 14.14, p < 0.001$, Cohen's $d = 1.39$, but not the different-categories event, $F(1, 66) = 0.04, p = 0.836, d = 0.11$, or the same-object event, $F(1, 66) = 1.60, p = 0.210, d = 0.54$. Non-parametric Wilcoxon rank-sum tests comparing conditions confirmed the results found with the different-objects ($Z = 3.01, p = 0.003$), different-categories ($Z = 0.17, p = 0.862$), and same-object ($Z = 1.59, p = 0.112$) events.

5.3. Discussion

As predicted, no difference was found between the no-panel and panel conditions for the same-object and different-categories events. In the case of the same-object event, the OF and PR systems agreed that a single object was involved in the occlusion event (Fig. 3A), so infants simply expected to see that object when the screen was lowered, with or without the transparent panel. In the case of the different-categories event, the OF system assigned distinct ontological descriptors to the toy animal and the doll (non-human vs. human-like), so the OF and PR systems agreed that two objects were involved in the occlusion event (Fig. 3C). Infants thus expected to see two objects when the screen was lowered, and they

detected an individuation violation when shown only one object, again with or without the transparent panel.

Also as predicted, responses for the different-objects event differed significantly between the two conditions. As the two dolls emerged in alternation from behind the screen, the OF system inferred that a single object was present, based on the spatiotemporal and categorical information at its disposal, whereas the PR system inferred that two objects were present, based on the causally relevant featural information at its disposal (e.g., the distinct sizes and shapes of the dolls' hairdos). In the no-panel condition, the occlusion event ended when the screen was lowered. Because the two systems disagreed quantitatively about the number of hidden objects present, infants experienced a catastrophic individuation failure, leading them to hold no expectation about how many objects would be revealed (Fig 3D). In the panel condition, in contrast, the occlusion event continued when the screen was lowered, with the panel now serving as the occluder, making this a one-event task. Because the PR system remained in charge, infants expected to see two objects behind the panel, and they detected an individuation violation when shown only one object (Fig. 3E).

Could infants in the panel condition who saw the different-objects event have succeeded not because they saw an ongoing occlusion event, as our model suggests, but for entirely different reasons? For example, it might be suggested that infants had established weak representations for the two dolls, and the transparent panel provided something akin to a "Watch this space!" sign that helped infants remain focused on the area behind the lowered screen and the two dolls that should be there. Alternatively, perhaps infants had formed a three-way association between the two dolls and the opaque screen, and the final display with both the doll and the transparent "screen" (panel condition) provided a stronger retrieval cue for the missing doll than the final display with the doll alone (no-panel condition). Though logically possible, these alternative interpretations are unlikely, for two reasons. First, recall that in the study of Wilcox and Chapa (2002), only the central portion of the screen was lowered at the end of the different-objects event (see Fig. 2C), leaving a thin frame that was either empty (no-panel condition) or filled with a transparent panel (panel condition). Importantly, the empty frame in the no-panel condition did not lead to successful individuation, even though it would also have helped draw attention to the area behind the lowered screen or provided a cue for the retrieval of the missing object. Second, all events in the no-removal condition of Experiment 3 ended with the box being shaken, which would be expected to provide both a "Watch this space!" sign and a retrieval cue for the objects in the box; nevertheless, as predicted by our model, infants who saw the different-objects event again failed to correctly individuate the objects, making these alternative interpretations unlikely.⁸

⁸This is not to say, of course, that memory considerations play little role in infants' physical reasoning. Indeed, there are many different ways in which working-memory limitations have been found to affect infants' ability to represent and reason about events. In a series of experiments (Kibbe & Leslie, 2011, 2013), for example, 6- and 9-month-olds first saw an object (e.g., a disk) being placed behind a screen, and then they saw a second object (e.g., a triangle) being placed behind a second screen. When the first screen was lifted to reveal the last-hidden object, 9-month-olds detected this swap. However, 6-month-olds did not, most likely because they lacked sufficient working-memory capacity to fully represent two separate occlusion events shown in quick succession (e.g., infants did not have time to retrieve or bind the featural information relevant to the object in the first occlusion event because they had to go on to represent the second occlusion event). When three separate occlusion events were shown in succession, and swaps now involved lifting the second screen to reveal the first-hidden object, 9-month-olds now failed, but 12-month-olds still succeeded (9-month-olds succeeded at detecting the swap in the second of three objects when only two occlusion events occurred in succession). In another

In sum, the results of Experiment 1 provide additional evidence for those of prior standard (Xu et al., 2004) and different-categories (Bonatti et al., 2002, 2005) tasks, and they confirm the finding by Wilcox and Chapa (2002) that young infants succeed at one-event tasks even when tested with objects they view as merely featurally distinct (e.g., a block and a ball, or two different dolls). Infants' contrasting responses to the different-objects event in the no-panel condition (standard task) and the panel condition (one-event task) were particularly striking: The simple addition of a transparent panel enabled infants to detect the individuation violation they were shown. This counter-intuitive finding supports two assumptions of our model. First, the OF and PR systems use somewhat different information to individuate objects and track their identity: The OF system uses primarily spatiotemporal and categorical information, whereas the PR system also uses causally relevant event-specific featural information. Second, event boundaries matter in infants' individuation performance. As long as a physical event is ongoing, the PR system is in charge of predicting how it will unfold; when it comes to an end, however, the OF and PR systems both weigh in on how many objects are present, and quantitative disagreements result in catastrophic individuation failures.

6. Experiment 2

According to our model, when infants fail at a standard or a search task, they are not merely uncertain as to whether one or two hidden objects are present. Rather, the quantitative disagreement between the OF system (which infers that one object is present) and the PR system (which infers that two objects are present) results in a catastrophic individuation failure, leading infants to hold no expectation at all about how many objects are present. Our model thus predicts that infants should fail to detect an individuation violation even if *no object* were present. Experiment 2 tested this prediction with 11-month-olds. Infants received a standard between- subjects task using two red cups that differed only in pattern: One had yellow stripes and the other had green dots (recall that pattern is typically identified as a causally relevant occlusion feature by about 7 months of age; Wilcox, 1999). Infants were randomly assigned to a no-object or a one-object condition (Fig. 6).

Infants in the *no-object* condition received one familiarization and one test trial. Half of the infants saw a different-objects event in each trial, and half saw a same-object event. These events were identical to those in Experiment 1 with two exceptions. First, as noted above, the two objects were a striped and a dotted cup; which cup was seen last in the different-objects event, and which cup was shown in the same-object event, were counterbalanced across infants (as before, two identical cups were used to produce the same-object event, for better experimental control). Second, the initial phase of the test trial was extended from 28 to 31 s, to allow the last-seen cup to be surreptitiously removed before the screen was lowered (the first-seen cup was removed after its last emergence). During the final phase of

experiment (Kibbe & Feigenson, 2016), 12-month-olds detected a violation when two different objects were placed, one at a time, behind a large screen, which was then lowered to reveal two identical objects (i.e., one of the objects had become identical to the other). However, infants failed to detect a violation when four different objects were hidden and two pairs of identical objects were revealed, again suggesting that infants lacked sufficient working-memory capacity to retrieve or bind the featural information relevant to four different objects hidden one after the other. Together, these results make clear that with increases in either the number of events or the number of objects in a single event, limits to infants' working memory seriously affect their performance.

the test trial, infants thus watched a final display with no cup: No object stood behind the lowered screen.

Infants in the *one-object* condition saw identical familiarization and test events except that the last-seen cup was not removed in the test trial and stood centered behind the lowered screen during the final phase of the trial.

Because infants in the no-object and one-object conditions saw final displays with different numbers of objects, which could create interpretive difficulties due to different baseline preferences (recall that in standard within-subject tasks, infants often show a baseline preference for final displays with two objects vs. one object; Bonatti et al., 2002; Futó et al., 2010; Surian & Caldi, 2010; Xu & Carey, 1996; Xu et al., 2004), our predictions focused on infants' responses *within* each condition, following the logic of standard between-subjects tasks. Thus, in the no-object condition, we predicted that infants who saw the same-object and different-objects events would respond differently. In the case of the same-object event, the OF and PR systems should agree that a single object was present, as there was no information to suggest otherwise (Fig. 3A). Infants should therefore expect to see one object when the screen was lowered, and they should detect an individuation violation when shown no object. In the case of the different-objects event, however, the OF system should infer that one object was present, based on the spatiotemporal and categorical information at its disposal, whereas the PR system should infer that two objects were present, based on the causally relevant featural information at its disposal (i.e., the cups' different patterns). This quantitative disagreement should lead to a catastrophic individuation failure: Infants should hold no expectation at all about the number of objects that should be revealed when the screen was lowered (Fig. 3D), and they should therefore fail to detect a violation when shown no object. Infants in the no-object condition should thus look significantly longer if shown the same-object as opposed to the different-objects event. Strikingly, seeing *more* objects emerge from behind the screen should render infants *less* likely to detect a violation when the screen was lowered to reveal no object.

In the one-object condition, we predicted that infants who saw the same-object and the different-objects events would respond similarly, albeit for different reasons. In the case of the same-object event, as explained above, the two systems should agree that one object was present; infants should thus view the final display of that object, when the screen was lowered, as expected. In the case of the different-objects event, however, the two systems should disagree quantitatively about how many objects were present, leading to a catastrophic individuation failure: Infants should hold no expectation about how many objects should be revealed, and they should therefore fail to detect a violation when shown one object. Infants in the one-object condition should thus look about equally, and equally short, whether they were shown the same-object or the different-objects event.

6.1. Method

6.1.1. Participants—Participants were 48 full-term infants (23 male, $M = 11$ months, 26 days, range = 10;29–12;24). Another 10 infants were tested but excluded (6 in the no-object condition and 4 in the one-object condition) because they were inattentive (4), distracted (1), or fussy (2), or had a test looking time over 3 standard deviations from the condition mean

(3; all saw the different-objects event, 1 in the no-object condition and 2 in the one-object condition). Twelve infants were randomly assigned to each combination of condition (no-object or one-object) and event (same-object or different-objects).

6.1.2. Apparatus and Stimuli—The experimental set-up was similar to that in Experiment 1 with the following exceptions. First, the stimuli were two red plastic cups (each 11×7 cm), one decorated with thin horizontal yellow stripes and the other with large green dots. Second, the screen ($34.5 \text{ cm} \times 38.5 \text{ cm} \times 0.5 \text{ cm}$) was beige, positioned 19 cm from the back wall (now 59 cm deep), and supported at the back by a 1-cm dowel that extended across the apparatus; as before, a hidden first assistant rotated the right end of the rod to raise and lower the screen. Third, because the cups had narrow bottoms (5 cm in diameter), a different mechanism was used to move them across the apparatus, to make sure their handles remained hidden from infants' view. Each cup had a vertical handle (a rod $16 \text{ cm} \times 1 \text{ cm}$) affixed to the center of its bottom surface. This handle was inserted into a long slit ($70 \text{ cm} \times 2.5 \text{ cm}$) in the apparatus floor, centered between the screen and the back wall; beneath the floor, a hidden second assistant slid the handle along a track (with markings and stoppers) to move the cup back and forth. To avoid producing sound cues that might indicate the number of hidden objects present, the slit was lined with felt, and the bottom of each cup and the top of each handle were also covered with felt. Finally, a concealed trapdoor in the apparatus floor, centered behind the screen, was used by a hidden third assistant to surreptitiously remove the first-seen cup (one-object condition) or both cups (no-object condition) before the screen was lowered in the test trial.

6.1.3. Procedure—The procedure was similar to that in Experiment 1. Infants were attentive during the initial phases of the familiarization and test trials and looked, on average, for 90% (no-object condition) and 91% (one-object condition) of each initial phase. The criteria used to end the final phases of the familiarization and test trials were the same as in Experiment 1, with one exception: The test trial ended when the infant looked away for 2 consecutive seconds, instead of 1 consecutive second. In the no-object condition, in particular, the use of a slightly longer look-away value gave infants more time to register and process the empty final display that was revealed when the screen was lowered (i.e., an empty display meant there was no object to attract and retain infants' attention). Inter-observer agreement during the final phase of each trial averaged 92% per trial per infant. Finally, preliminary analyses of the test data revealed no significant interaction of condition and event with either infants' sex or last-seen object (dotted or striped cup), both p s > 0.538 ; the data were therefore collapsed across these latter two factors.

6.2. Results

Looking times during the final phase of the *familiarization* trial were compared by means of an ANOVA with condition (no-object or one-object) and event (same-object or different-objects) as between-subjects factors. The analysis yielded no significant effects, all p s > 0.182 , suggesting that across conditions and events, infants tended to look equally at the upright screen during the final phase of the familiarization trial (for descriptive statistics, see Experiment 2 Dataset in SOM).

Looking times during the final phase of the *test* trial were analyzed as above (Fig. 7). The main effect of event was significant, $F(1, 44) = 4.75, p = 0.035$, as was the Condition \times Event interaction, $F(1, 44) = 4.39, p = 0.042, \eta^2 = 0.09$. Planned comparisons revealed that, as predicted, (a) infants in the no-object condition looked significantly longer if shown the same-object event ($M = 19.20, SD = 5.47$) than if shown the different-objects event ($M = 13.99, SD = 3.05$), $F(1, 44) = 9.14, p = 0.004, d = 1.16$, and (b) infants in the one-object condition looked about equally whether they were shown the same-object event ($M = 14.57, SD = 2.22$) or the different-objects event ($M = 14.74, SD = 4.14$), $F(1, 44) = 0.00, p = 0.953, d = 0.00$. Wilcoxon rank-sum tests confirmed the results of the no-object ($Z = 2.31, p = 0.021$) and one-object ($Z = 0.75, p = 0.452$) conditions.⁹

6.3. Further Results

In the no-object condition, infants who saw the same-object event looked significantly longer than those who saw the different-objects event. This finding suggested that (a) infants who saw the same-object event detected an individuation violation when the screen was lowered to reveal no cup, but (b) infants who saw the different-objects event did not detect this violation, consistent with the claim that they were experiencing a catastrophic individuation failure. To confirm this finding, which provided key evidence for our model, we tested an additional group of 11-month-olds using the same procedure as in the no-object condition ($n = 24, M = 11$ months, 21 days, range = 10;28–12;22); another 5 infants were excluded because they were inattentive (2) or fussy (1), peeked behind the screen (1), or were subject to parental interference (1). Half of the infants in this no-object replication saw the same-object event, and half saw the different-objects event. Infants were attentive during the initial phases of the trials, looking, on average, for 89% of each initial phase. Inter-observer agreement averaged 89% per trial per infant. Looking times during the final phase of the test trial were analyzed using an ANOVA with event (same-object or different-objects) as a between-subjects factor. This effect was significant, $F(1, 22) = 9.29, p = 0.006, d = 1.29$, indicating that infants who were shown the same-object event ($M = 16.43, SD = 2.87$) looked significantly longer than those who were shown the different-objects event ($M = 12.99, SD = 2.75$). A Wilcoxon rank-sum test confirmed this result ($Z = 2.40, p = 0.016$).

Lastly, we compared test looking times in this no-object replication to those in the no-object condition of Experiment 2 ($n = 48$), using an ANOVA with experiment (replication or Experiment 2) and event (same-object or different-objects) as between-subjects factors. The analysis yielded only a significant main effect of event, $F(1, 44) = 15.96, p < 0.001, \eta^2 = 0.27$, indicating that infants who saw the same-object event ($M = 17.81, SD = 4.50$) looked

⁹As noted earlier, comparisons of final displays involving different numbers of objects can be difficult to interpret, as infants sometimes display a baseline tendency to look longer at displays containing more objects. Nevertheless, for the sake of completion, we also compared infants' responses to each event between the no-object and one-object conditions. For the same-object event, infants looked significantly longer at the final display in the no-object condition ($M = 19.20, SD = 5.47$) than in the one-object condition ($M = 14.57, SD = 2.22$), $F(1, 44) = 6.40, p = 0.015, d = 1.05$, and this result was confirmed by a Wilcoxon rank-sum test, $Z = 2.37, p = 0.018$. Infants thus held a clear expectation that a single cup was hidden behind the screen, and they detected a violation when no cup was revealed. For the different-objects event, infants looked about equally at the final display in the no-object condition ($M = 13.99, SD = 3.05$) and the one-object condition ($M = 14.74, SD = 4.14$), $F(1, 44) = 0.19, p = 0.666, d = 0.20$, and this result was again confirmed by a Wilcoxon rank-sum test, $Z = 0.41, p = 0.685$. Thus, in line with our model, infants held no expectation about what they would see when the screen was lowered, so they looked about equally whether they saw no cup or one cup.

significantly longer than those who saw the different-objects event ($M = 13.49$, $SD = 2.89$). A Wilcoxon rank-sum test again confirmed this result ($Z = 3.45$, $p < 0.001$).

6.4. Discussion

As predicted, infants in the no-object condition and its replication looked significantly longer if shown the same-object as opposed to the different-objects event. In the case of the same-object event, the OF and PR systems agreed that a single object was present behind the screen (Fig. 3A), so infants expected to see that object when the screen was lowered, and they detected an individuation violation when shown no object, confirming prior findings with occlusion events involving a single object (Kibbe & Leslie, 2011; Luo et al., 2009). In the case of the different-objects event, however, the OF and PR systems posited different numbers of objects behind the screen, leading to a catastrophic individuation failure (Fig. 3D). These interpretations were supported by the results of the one-object condition, in which infants tended to look equally, and equally short, at the two events. In the same-object event, infants simply saw what they expected to see when the screen was lowered; in the different-objects event, infants again experienced a catastrophic individuation failure and held no expectation about what they would see.

Thus, infants who saw the different-objects event failed to detect an individuation violation not only when the screen was lowered to reveal one object, as in prior research, but also when the screen was lowered to reveal no object. These results provide evidence for two assumptions of our model. First, the OF and PR systems use somewhat different information to individuate objects, so they sometimes posit different numbers of objects in an event. Second, when an event comes to an end, a quantitative disagreement between the OF system (one object) and the PR system (two objects) leads infants to hold no expectation at all about the number of objects present. Infants cannot easily align or reconcile the two systems' divergent representations, leading to a catastrophic individuation failure.

7. Experiment 3

Experiment 3 had two goals. One was to provide converging evidence, using a novel task, for our assumption that quantitative disagreements between the OF and PR systems result in catastrophic individuation failures. The other goal was to provide new evidence, using a modified version of this same task, for our assumption that qualitative disagreements, unlike quantitative ones, are easily resolved, resulting in successful individuation.

In our novel task, 9-month-olds saw a containment event followed by a shaking event: Rigid objects were lifted from and lowered into a large box, which was then shaken briskly. When shaken, the box produced no thuds or thumping sounds but instead was silent, as though empty. The rationale was that if infants held a clear representation that the box contained one or two rigid objects, then (a) they should expect the object(s) to collide with the box's interior walls during the shaking event, (b) they should expect to hear thumping sounds when these impacts occurred (Bahrick, 1983; Setoh et al., 2013), and (c) they should detect a violation when the box was silent. If, however, infants held no clear representation of how many objects were in the box, due to a quantitative disagreement between the OF and PR systems (i.e., one vs. two objects in the box), then they should fail to detect a violation when

the shaken box was silent. Finally, if the task was modified to bring about a qualitative, as opposed to a quantitative, disagreement between the two systems (i.e., no object vs. one object in the box), this disagreement should then be resolved, leading infants to detect a violation when the shaken box was silent.

Infants were randomly assigned to a no-removal or a removal condition (Fig. 8). Infants in the *no-removal* condition received one orientation trial, one familiarization trial, and one test trial; each trial had an initial phase and a final phase (Fig. 7). At the start of the (5-s) initial phase in the orientation trial, a large box rested upright on the apparatus floor, and an experimenter's gloved right hand held the back of the box. The hand rotated the box toward the infant, to show that it was empty, and then returned it to its original orientation. During the final phase of the trial, the hand (a) shook the box by sliding it briskly back and forth on the apparatus floor (2 s), (b) paused briefly (2 s), and then (c) repeated this sequence of actions until the trial ended. As would be expected, the empty box produced no internal sounds when shaken.

The familiarization trial introduced infants to the objects that would be used in the test trial. Infants saw one of three events: different-objects, same-object, or different-locations. Stimuli were two rigid toys that 9-month-olds would encode as merely featurally distinct, a toy wolf and a soccer ball; each object was mounted on an identical wooden support. At the beginning of the (35-s) initial phase in the *different-objects* event, the box rested upright on the apparatus floor; the experimenter's right hand rested at a window in the back wall, and her left hand held the back of the box (to keep it in place). To start, the right hand reached into the box, lifted one of the objects (e.g., the wolf; counterbalanced), and placed it on the apparatus floor next to the box. The hand then tapped the object first against the apparatus floor and then against the left side of the box; this tapping served to demonstrate that the object produced impact sounds, as would be expected, when it collided with other rigid surfaces. The hand then returned the object to the box (16 s). After a 2-s pause, the hand repeated these actions with the other object (e.g., the ball; 16 s), and then withdrew to the window and paused (1 s). During the final phase of the trial, infants watched this paused scene until the trial ended. The *same-object* event was identical except that the same object (e.g., the wolf; counterbalanced) was shown twice (as before, two identical toys were used to produce the event, for better experimental control). Finally, the *different-locations* event was identical to the different-objects event except that the second object to be manipulated (e.g., the ball) rested on the apparatus floor at the start of the trial. Thus, when the hand lifted the first object (e.g., the wolf) from the box, infants could see both objects simultaneously, in different locations. After returning the first object to the box, the hand grasped the second object, tapped it, and then placed it inside the box.

The procedure in the test trial was identical to that in the familiarization trial, with one exception: At the end of the (35-s) initial phase, the experimenter grasped the back of the box and then shook it back and forth until the trial ended, as in the orientation trial. Critically, the box was always silent when shaken (the objects were surreptitiously removed from it), so that all infants watched the same shaking event.

Infants in the *removal* condition underwent the same procedure as in the no-removal condition with one exception: After tapping the second object in the (35-s) initial phases of the familiarization and test trials, the experimenter's right hand removed it from the apparatus through the window in the back wall, instead of placing it in the box. The hand then either rested at the window (familiarization trial) or shook the box (test trial), as before.

Our predictions were as follows. In the case of the same-object test event, different predictions held for the no-removal and removal conditions. In the no-removal condition, the OF and the PR systems should agree that one object was in the box (there was no information to suggest otherwise; Fig. 3A); infants should thus expect to hear thumping sounds during the shaking event as the object collided with the box's interior walls, and they should detect a violation when the box was silent, suggesting that it was empty. In the removal condition, the OF and the PR systems should agree that the box was now empty, and hence they should detect no violation when the shaken box was silent. In the case of the different-locations event, predictions were similar for the no-removal and removal conditions. Based on the spatiotemporal information available (Fig. 3C), the two systems should agree that the box contained either two objects (no-removal condition) or one object (removal condition). Thus, in either case, infants should expect to hear thumping sounds from the shaken box, and they should detect a violation when they heard no such sounds, as though the box was empty. Finally, in the case of the different-objects event, predictions differed for the two conditions. In the no-removal condition, the OF and PR systems should disagree about the number of objects present. The OF system should assume, based on the spatiotemporal and categorical information at its disposal, that one object was in the box, whereas the PR system should assume, based on the causally relevant featural information at its disposal (e.g., the toys' different shapes and sizes) that two objects were in the box. This quantitative disagreement (i.e., one vs. two objects in the box; Fig. 3D) should lead to a catastrophic individuation failure, so infants should hold no particular expectation about what happened next, and hence they should detect no violation when the shaken box was silent. In the removal condition, the OF system should assume that no object was left in the box, following the last-seen object's removal from the apparatus, whereas the PR system should assume that the first-seen object was still in the box. As in other remainder tasks (Fig. 2D), this qualitative disagreement (i.e., no object vs. one object in the box; Fig. 3E) should be easily resolved. Upon receiving the signal from the PR system that in addition to the object removed from the apparatus, there was an object remaining in the box (i.e., the first-seen object), the OF system should add a new object file pointing to that object, leading to a clear expectation for one object in the box. When the shaking event began, infants should thus expect to hear thumping sounds, and they should detect a violation when they did not.

Together, these predictions suggested that (a) in the no-removal condition, infants should look significantly longer if shown the same-object and the different-locations events as opposed to the different-objects event; (b) in the removal condition, infants should look significantly longer if shown the different-locations or the different-objects event than if shown the same-object event; and (c) both the same-object and the different-objects event should differ between conditions, though with reverse looking patterns: Infants who saw the same-object event should look significantly longer in the no-removal (one object) than in the removal (no object) condition, whereas infants who saw the different-objects event should

look significantly longer in the removal (one object) than in the no-removal (two objects) condition. Finding the predicted results would confirm those of Experiment 2, would support our analysis of remainder tasks, and more generally would provide robust evidence for our assumptions concerning the consequences of quantitative and qualitative disagreements between the OF and PR systems for infants' ability to individuate and track objects.

7.1. Method

7.1.1. Participants—Participants were 72 full-term infants (36 male, $M = 9$ months, 18 days, range = 9;0 – 10; 12). Another 5 infants were tested but excluded (2 in the no-removal condition, and 3 in the removal condition) because they were fussy (2) or inattentive (1), or had a test looking time over 3 standard deviations from the condition mean (2; 1 who saw the different-objects event in the no-removal condition, and 1 who saw the same-object event in the removal condition). Twelve infants were randomly assigned to each combination of condition (no-removal, removal) and event (same-object, different-objects, different-locations).

7.1.2. Apparatus and Stimuli—The apparatus was the same as in Experiment 2 except that the back wall had a narrow window (25 cm × 46 cm) located 24.5 cm above the floor, 30 cm from the right wall. An experimenter wearing a white shirt and white gloves stood behind the window. Stimuli included two identical brown boxes (each 23 cm × 22 cm × 22 cm and lined with a brightly patterned adhesive paper). One box was used in the orientation trial. It initially rested 3.5 cm in front of the back wall, centered between the two sidewalls; its bottom was covered with felt so that it made very little noise when the experimenter repeatedly shook it briskly along the floor, toward the infant and back, over a distance of about 20 cm. The other box was used in the familiarization and test trials, and it was shifted to the right so that its right edge aligned with that of the window. Unbeknownst to infants, this box had a hole in its bottom, which matched an opening in the apparatus floor; beneath the floor, a hidden assistant used this opening to silently pass an object (when the experimenter reached into the box to lift an object) or take away an object (when the experimenter lowered an object into the box). The objects were two identical brown toy wolves (each 13 cm × 8.5 cm × 8 cm) and two identical yellow soccer balls (each 8 cm in diameter). The objects were rigid and mounted on identical wooden bases (each 1 × 8.5 × 4.5 cm), so that they produced similar sounds when their bases were tapped against the apparatus floor or the side of the box. Finally, the experimenter always lifted and lowered objects at the center of the box, to prevent giving spatial cues as to how many objects were present.

7.1.3. Procedure—The procedure was similar to that in Experiment 1. Infants were attentive during the initial phases of the orientation, familiarization, and test trials; in each condition, they looked, on average, for 96% of each initial phase. The criteria used to end the final phases of the trials were the same as in Experiment 1. Inter-observer agreement during the final phase of each trial averaged 95% per trial per infant. Finally, preliminary analyses of the test data revealed no significant interaction of condition and event with either infants' sex or last-seen object (wolf or ball), both p s > 0.925; the data were therefore collapsed across these latter two factors.

7.1.4. Adult Data—To confirm our analysis of our shaking task, we asked 20 adults (15 female, $M = 21$ years, range = 19–23) whether they would expect to hear thumping sounds from the box if the experimenter were to shake it with one of the toys in it. Participants sat in front of the apparatus. To start, the experimenter showed them that the box was empty and that it produced no internal sounds when briskly shaken (as in the orientation trial). Next, the experimenter grasped one of the toys (e.g., the wolf; counterbalanced), tapped its base against the apparatus floor and the side of the box (as in the familiarization and test trials), and then placed the toy in the box. Adults rated on a 7-point scale (1 = no sounds, 7 = very loud sounds) whether they would now expect to hear sounds from the box if the experimenter were to shake it as before. The experimenter then removed the toy from the box and repeated the same procedure with the other toy (e.g., the ball). For each toy, ratings were compared against the possibility of expecting no sounds (rating = 1) using a one-sample t -test. Adults significantly expected to hear sounds from the shaken box when either the wolf ($M = 4.75$, $SD = 0.91$, $t(19) = 18.42$, $p < 0.001$) or the ball ($M = 4.50$, $SD = 1.10$, $t(19) = 14.23$, $p < 0.001$) was placed inside it; ratings for the two toys did not differ, $F(1, 19) = 1.00$, $p = 0.330$. Thus, not surprisingly, adults expected each toy to produce thumping sounds when the box was shaken, and they expected equally loud sounds for the two toys.

7.2. Results

Looking times during the final phase of the *orientation* trial were compared by means of an ANOVA with condition (no-removal or removal) and event (same-object, different-objects or different-locations) as between-subjects factors. The analysis yielded no significant effects, all $ps > 0.445$. Looking times during the final phase of the *familiarization* trial were analyzed in the same manner, and the analysis again produced no significant results, all $ps > 0.104$. Thus, across conditions and events, infants tended to look equally during the orientation and familiarization trials (for descriptive statistics, see Experiment 3 Dataset in SOM).

Looking times during the final phase of the *test* trial were analyzed as above (Fig. 9). The only significant effect was the Condition \times Event interaction, $F(2, 66) = 11.06$, $p < 0.001$, $\eta^2 = 0.25$. Five planned comparisons (with a Bonferroni-adjusted α -level of 0.01) were conducted. In the no-removal condition, infants who saw the same-object event ($M = 22.84$, $SD = 5.88$) or the different-locations event ($M = 22.48$, $SD = 7.13$) looked significantly longer when the box was shaken than did infants who saw the different-objects event ($M = 14.37$, $SD = 4.54$), $F(1, 66) = 16.38$, $p < 0.001$. In the removal condition, infants who saw the different-objects event ($M = 22.85$, $SD = 7.21$) or the different-locations event ($M = 20.26$, $SD = 6.30$) looked significantly longer when the box was shaken than did infants who saw the same-object event ($M = 15.54$, $SD = 4.71$), $F(1, 66) = 8.29$, $p = 0.005$. Finally, comparing between the two conditions, different looking times were found for the different-objects event, $F(1, 66) = 12.28$, $p < 0.001$, and for the same-object event, $F(1, 66) = 9.34$, $p = 0.003$, but not for the different-locations event, $F(1, 66) = 0.52$, $p = 0.475$. Wilcoxon rank-sum tests also showed significant differences between conditions for the different-objects event ($Z = 2.52$, $p = 0.012$), and the same-object event, ($Z = 2.75$, $p = 0.006$), but not the different-locations event ($Z = 0.58$, $p = 0.559$).

7.3. Discussion

As predicted, infants who saw the same-object event looked significantly longer at the shaking of the box in the no-removal than in the removal condition. In the no-removal condition, the OF and PR systems agreed that the box held one object (no information suggested otherwise; Fig. 3A). Thus, when the box was shaken, infants expected the object to produce thumping sounds as it collided with the box's interior walls, and they detected a violation when they heard no such sounds, as though the object had spontaneously vanished from the box.¹⁰ In the removal condition, both systems agreed that the box was now empty, as the object had been visibly removed from the apparatus instead of being returned into the box, so infants did not detect a violation when the shaken box was silent. Together, these results demonstrate that infants could reason correctly about our novel task: They expected the shaken box to produce thumping sounds when it contained at least one object (as did our adult participants), but to be silent otherwise.

Also as predicted, infants who saw the different-locations event tended to look equally, and equally long, in the no-removal and removal conditions. Based on the spatiotemporal information available, both systems assumed that the box held either two objects (no-removal condition; Fig. 3B) or one object (removal condition). Infants therefore expected to hear thumping sounds when the box was shaken, and they detected a violation when they heard no such sounds. The results of the no-removal condition are particularly important because they make clear that when infants represented two objects in the box, they expected to hear thumping sounds during the shaking event as the objects collided with each other and with the box's interior walls (i.e., infants did not assume the two objects would fit tightly within the box and produce no sounds).

Finally, also as predicted, infants who saw the different-objects event looked significantly longer at the shaking of the box in the removal than in the no-removal condition. In the no-removal condition, the OF system assigned similar categorical descriptors to the wolf and the ball and therefore inferred that the box held a single object; in contrast, the PR system represented some of the featural differences between the two toys and therefore inferred that the box held two objects. This quantitative disagreement (Fig. 3D) resulted in a catastrophic individuation failure, so infants held no particular expectation about what would happen next. In the removal condition, however, the visible removal of the last-seen toy from the apparatus resulted in a qualitative, as opposed to a quantitative, disagreement between the two systems. The OF system signaled that the object it had been tracking was now gone, leaving the box empty, but the PR system signaled that the box in fact still held an object. The OF system then added an object file pointing to this object, bringing the two systems in alignment. As a result, infants expected thumping sounds from the shaken box and detected a violation when they heard no such sounds, as though the box was empty. These results are

¹⁰We assume that upon noticing that the box was silent when shaken, infants concluded that the object had inexplicably vanished (a continuity violation). However, infants could have posited more sophisticated types of persistence violations to make sense of what they saw and heard. For example, they could have concluded that the object had magically passed through the back of the box (a solidity violation), had become attached to one of the box's interior walls (a boundedness violation), or had changed into a soft object that caused no sound when colliding with the box's interior walls (an unchangeableness violation). Based on our data, we cannot determine whether or how infants attempted to explain the violation they detected, but it is not necessary that we do so for testing our model. The critical finding is that when infants successfully individuated the object, they expected to hear it bounce around during the shaking event and detected a violation when they did not.

particularly useful because they illustrate the difference between quantitative and qualitative disagreements under highly similar conditions. Infants failed to detect the individuation violation they were shown when the first and second objects were each returned to the box prior to the shaking event (resulting in a quantitative disagreement), but they succeeded if the second object was removed from the apparatus instead (resulting in a qualitative disagreement).¹¹

As was alluded in section 5.3., the results of Experiment 3 also help rule out alternative interpretations of the positive result obtained with the different-objects event in the panel condition of Experiment 1. Recall that according to these interpretations, the transparent panel behind the screen provided either an attentional or a mnemonic cue that helped infants detect the individuation violation they were shown. In Experiment 3, infants saw the experimenter shake the box at the end of every event, and this repeated shaking no doubt helped infants focus on the box and the objects it contained. Nevertheless, as predicted by our model, infants correctly individuated the objects in the box only when there was agreement, or qualitative disagreement, between the OF and PR systems; quantitative disagreement, as usual, led to a catastrophic individuation failure.

8. General Discussion

For over two decades, the study of young infants' ability to individuate objects has captivated developmental researchers' interest, because this ability sheds light on our earliest representations of objects. However, the large body of findings generated by these investigations has yielded an unwieldy picture, with infants succeeding at some individuation tasks but failing at others. To date, various explanations have been offered for subsets of these findings, but none so far has been able to account for infants' positive and negative performances across tasks. In this article, we proposed a new model of early object individuation that sought to do just that.

Because object individuation is part and parcel of infants' ability to represent and reason about physical events, we situated our work within the broader framework of early physical reasoning. We strove to integrate findings from object individuation and physical reasoning into a single model of how infants individuate, represent, and reason about objects in physical events.

Our model has three main assumptions. First, when different objects emerge in alternation from behind a screen or from inside a box, two cognitive systems help infants individuate the objects: the object-file (OF) and physical-reasoning (PR) systems. The OF system uses primarily spatiotemporal and categorical information for this purpose, whereas the PR system also uses whatever featural information has been identified as causally relevant for the event category involved. Because the two systems rely on somewhat different

¹¹An anonymous reviewer suggested a clever way to adapt our box task to offer not only a standard task (no-removal condition) and a remainder task (removal condition), as in Experiment 3, but also a one-event task. In this new task, infants would again see either two objects in alternation (different-objects event) or one object (same-object event) lifted from and returned into the box. Next, instead of shaking the box, the experimenter would rotate it to reveal a large transparent window in its back sidewall, with only one object visible through this window. In this case, the containment event would still be ongoing, with the PR system in control of predicting its progress, and infants should now detect the individuation violation in the different-objects event.

information to individuate the objects, conflicts sometimes arise between them. Second, during the event, the PR system has priority over the OF system in predicting how the event will unfold. Third, when the event comes to an end, and the two systems agree on how many objects are present behind the screen, infants successfully track the objects to the next event. If the systems disagree, however, infants' performance depends on the nature of this disagreement. Quantitative disagreements (in which both systems agree that objects are present but disagree on their number) lead to catastrophic individuation failures: Infants hold no expectation at all about how many objects are present (not even a default expectation that there should be at least one object present!). In contrast, qualitative disagreements (in which the OF system assumes that there is no hidden object left but the PR system assumes that one or more hidden objects still remain) are easily resolved, leading to successful performance.

The three experiments with 9–11-month-olds ($N = 216$) reported in this article supported our model. Experiment 1 replicated the findings of Wilcox and Chapa (2002) with one-event tasks, supporting the second assumption listed above. After two different dolls emerged in alternation from behind a screen, infants detected a violation if the screen was lowered to reveal only one of the dolls standing behind a transparent panel. The addition of this panel meant that the occlusion event continued on, with the PR system remaining in charge of predicting how the event would unfold; because it had established separate representations for the two dolls, based on their featural differences, infants detected a violation when one spontaneously disappeared.

Experiment 2 supported the third assumption listed above and more specifically the claim that quantitative disagreements between the OF and PR systems lead to catastrophic individuation failures. After two different cups emerged in alternation from behind a screen, infants failed to detect a violation if the screen was lowered to reveal no cup at all. Experiment 3 provided additional support for this claim using a novel event sequence. After seeing two different toys being lifted in alternation from a large box, infants failed to detect a violation if the box was then shaken and produced no internal sounds, as though empty. In each experiment, the OF system inferred that a single object was hidden, behind the screen or inside the box, whereas the PR system inferred that two objects were hidden; infants could not resolve this quantitative disagreement and hence held no expectation about what they would see or hear next.

Finally, Experiment 3 also provided support for the claim that qualitative disagreements between the OF and PR systems are easily resolved and hence do not lead to catastrophic individuation failures. In the removal condition, the second toy that was lifted from the box was removed from the apparatus. This resulted in a qualitative disagreement about how many objects were left in the box, with the OF system signaling that no object was left and the PR system signaling that one object was left. The OF system then created an object file for this object, thereby bringing the two systems into agreement. As a result, infants held a clear expectation that one object was still in the box, and they detected a violation when the shaken box was silent, as though empty.

Together, these findings provide robust evidence for our model and demonstrate that it can predict and explain infants' performance in a variety of individuation tasks, including novel tasks. Still, many questions remain about the development of the cognitive systems that underlie early object individuation, and about the continuity of these systems from the first year of life to adulthood. We conclude our article by suggesting three directions for future research.

8.1. Scaling up Our Model: Infants' Individuation of Three Objects

In prior different-locations tasks, 13–18-month-olds could remember the presence of three different objects that were first displayed simultaneously and then hidden in a box (Zosh & Feigenson, 2012, 2015). Keeping track of three objects in a physical event thus appears to be within the limits of infants' working-memory capacity. Given these findings, one might ask how well our model would scale up in predicting young infants' performance in individuation tasks with three featurally distinct objects, such as three different balls. Our predictions would in fact be the same as with two objects. In standard or search tasks, the OF system would infer that a single object was present (based on the categorical information at its disposal); the PR system would infer that three objects were present (based on the causally relevant featural information at its disposal); and this quantitative disagreement (one object vs. three objects) would result in a catastrophic individuation failure: Infants would hold no expectation at all about how many objects were present.

Now, imagine a remainder task involving the three different balls. At the end of the occlusion or containment event, the third ball, instead of being returned behind the screen or inside the box, would be left in plain view or removed from the scene. In either case, the OF system would assume there was no hidden object left; the PR system would assume there were two hidden objects left; and this qualitative disagreement (no object vs. two objects) would be resolved by the OF system adding object files as needed, leading infants to expect two objects.

8.2. Types of Categorical Information

We have seen that infants succeed at individuation tasks when they encode the two objects as categorically distinct, either spontaneously or via experimental manipulations. According to our model, the nature of these categorical distinctions does not matter: *Any* categorical distinctions should enable the OF system to correctly individuate the two objects. Prior individuation tasks have tended to focus on only a few types of categorical distinctions, leaving open several avenues for future research.

First, prior manipulations used in standard tasks to induce young infants to assign objects to distinct categories have all focused on deep conceptual categories formed via linguistic, functional, and/or pedagogical evidence (Futó et al., 2010; Stavans & Baillargeon, 2018; Xu, 2002). According to our model, however, simple manipulations focusing on *shallow perceptual categories* should be equally effective. One such manipulation is suggested by experiments on early object segregation conducted by Needham and her colleagues (Dueker, Modi, & Needham, 2003; Needham, Dueker, & Lockhead, 2005). Four-month-olds were found to correctly parse a static test display composed of two adjacent objects, a block and a

cylinder, if they were first exposed to an array of three different blocks similar to the test block. According to Needham, infants (a) formed a “block” category when shown the array of three blocks, (b) recognized the test block as a novel member of this category, and (c) inferred that the test block and cylinder were two separate objects. Building on these results, one could examine whether young infants who were first shown a static array of three blocks and a static array of three cylinders (in separate trials) would then succeed at a standard task involving one object from each array. Positive results would support the claim that any categorical distinctions available to the OF system—even ones involving shallow perceptual categories—can lead to success in standard tasks.

Second, prior individuation tasks in which infants successfully assigned the two objects to distinct categories (either spontaneously or via manipulations) have all made use of categorical distinctions related to the objects’ intrinsic properties (e.g., a ball and a toy duck could not be interchangeably assigned to each other’s category, nor could a masher and a pair of tongs). However, there are other types of categorical distinctions that are more extrinsic to the objects involved, such as *event roles*. Recall that when the PR system builds a representation of a physical event, it categorizes the event (e.g., collision) and assigns distinct categorical roles to the objects (e.g., hitter, hittee; Leslie & Keeble, 1987). It might be possible to take advantage of these event roles to induce young infants to succeed at standard tasks. Imagine that 12-month-olds first received two trials involving collision events. In one trial, a blue ball would roll across an apparatus floor and hit a toy car; in the other trial, the toy car would roll across the floor and hit a red ball. In the test trial, infants would see a standard between-subjects different-objects event involving the blue and the red balls. Evidence that infants now detected the violation in this event would suggest that (a) as each ball emerged from behind the screen, the OF system gathered information from memory as well as from incoming perceptual information about the ball; (b) the information from memory included each ball’s role in the collision events (e.g., formerly a hitter vs. formerly a hittee); (c) the OF system used these distinct categorical descriptors to correctly individuate the two balls; and (d) the OF and PR systems then agreed that two balls were present, leading infants to expect to see both balls when the screen was lowered.

Finding that any categorical distinction (whether conceptually rich or shallow, and whether intrinsic or extrinsic to the objects’ properties) leads to success in individuation tasks would provide additional evidence for our model and for the privileged role of categorical information in the OF system, at least in the first year of life.

8.3 The OF System and Featural Information

We have argued that 12-month-olds fail at standard or search tasks involving objects that are merely featurally distinct in large part because the OF system considers only the categorical information at its disposal when individuating objects. How can we explain this selective reliance on categorical information? One possibility is that it reflects a *competence* limitation. In this view, the OF system would at first be *unable* to use featural information to individuate objects. Another possibility is that this initial reliance on categorical information reflects a *performance* limitation. In this view, the OF system would be *capable* of using

featural information for individuation purposes but would at first do so spontaneously only if the information was highly salient.

One way to evaluate these two possibilities would be to examine whether *priming* the OF system to attend to a particular feature could result in a successful individuation performance. Imagine that 12-month-olds were first primed to attend to color information in a context chosen to activate only the OF system (i.e., no physical event took place). For example, infants might see a single static array of three balls differing only in color, to highlight color as an individuating feature. Next, in the test trial, infants would see a standard between-subjects different-objects event involving two novel balls that again differed only in color. Evidence that infants now detected a violation when the screen was lowered to reveal only one ball would suggest that (a) the OF system used the highlighted color information to individuate the two balls, (b) the OF and PR systems agreed that two balls were present, and hence (c) infants expected to see both balls when the screen was lowered. Such positive results would suggest that the OF system is *capable* of using featural information to individuate objects in physical events; however, it tends not to do so, at least in the first year of life, unless that information is made particularly salient. These results would dovetail with many other findings on young infants' reasoning about objects in physical events. For example, recall that although young infants' OF system does not spontaneously consider objects' basic-level categories when individuating objects, it can be primed to do so (Stavans & Baillargeon, 2018; Xu, 2002). Similarly, although young infants' PR system does not spontaneously consider pattern or color information when predicting how an occlusion event will unfold, it can be primed to do so (Baillargeon et al., 2009a; Wilcox & Chapa, 2004).

Future research will also need to establish at what age infants begin to spontaneously succeed at standard or search tasks involving merely featurally distinct objects, and by what mechanisms they come to do so. The answer to this last question will depend in part on whether the OF system's initial reliance on categorical information represents a competence or a performance limitation, as discussed above. In the latter case, several factors could contribute to the OF system's emerging use of featural information: (a) increases in infants' information-processing resources could make it possible for the OF system to consider a larger subset of the available information when checking for object continuity during physical events (Cowan, 2016; Diamond, 2013); and (b) infants' acquisition of adjectives for referring to objects' individual features (e.g., "striped", "big") could render these features more salient or accessible for the OF system (Booth & Waxman, 2009; Ferguson & Waxman, 2017), much like the acquisition of count nouns is thought to contribute to the OF system's spontaneous use of basic-level categories beginning around the first birthday (Carey, 2011; Xu & Carey, 1996).

More generally, and perhaps most importantly, all of this research will help bring to light the similarities and differences between the OF and PR systems. For example, it may be that both systems (a) consider more and more information with age when individuating objects in physical events, and (b) can often be primed to consider information that is not yet spontaneously used for this purpose. Nevertheless, the two systems may differ in *how* they come to consider additional information with age. For the PR system, such changes may

occur largely through explanation-based learning, as infants identify features that are causally relevant for predicting outcomes in each event category (Baillargeon & DeJong, 2017; Wang & Kohne, 2007). For the OF system, however, such changes may occur through quite different mechanisms, such as the acquisition of labels for referring to objects' categories and properties (Carey, 2011; Ferguson & Waxman, 2017; Xu & Carey, 1996).

8.4. Concluding Remarks

We began this article with the following question: Why do young infants fail to detect individuation violations but successfully detect other types of persistence violations? We suggested that individuation failures are best understood as catastrophic failures caused by quantitative conflicts between two cognitive systems, the OF and PR systems, at event boundaries. We then provided novel evidence for these catastrophic failures by showing that (a) after seeing two different objects emerge in alternation from behind a screen, infants detected no violation if the screen was lowered to reveal no object, and (b) after seeing two different objects revealed in alternation from inside a box, infants detected no violation if the box was silent when shaken, as though empty.

Many details of our model may turn out to be incorrect. Nevertheless, we hope that our efforts will inspire new behavioral and brain research on the cognitive architecture that underlies infants' ability to represent, individuate, and reason about objects, and on the development of this architecture from infancy to adulthood.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This research was supported by a Fulbright postdoctoral fellowship to M. S. and by grants from NICHD (HD-021104) and the UIUC Research Board to R. B. All authors have contributed significantly to the preparation of the manuscript and have approved its final version. We thank Cowan Nelson and two other anonymous reviewers for their valuable input; Dan Simons for many helpful discussions; and Diane Beck, Jerry DeJong, Cindy Fisher, Dan Hyde, Dave Irwin, Alejandro Lleras, Francesco Margoni, Su-hua Wang, and Teresa Wilcox for helpful suggestions or comments as we developed our model. We also thank Steve Holland for his help with the figures; the staff of the UIUC Infant Cognition Laboratory for their help with the data collection; and the parents and infants who participated in the research.

References

- Aguiar A, & Baillargeon R (1998). Eight-and-a-half-month-old infants' reasoning about containment events. *Child Development*, 69(3), 636–653. doi: 10.1111/j.1467-8624.1998.tb06234.x [PubMed: 9680677]
- Aguiar A, & Baillargeon R (1999). 2.5-month-old infants' reasoning about when objects should and should not be occluded. *Cognitive Psychology*, 39(2), 116–157. doi: 10.1006/cogp.1999.0717 [PubMed: 10462457]
- Aguiar A, & Baillargeon R (2002). Developments in young infants' reasoning about occluded objects. *Cognitive Psychology*, 45(2), 267–336. doi: 10.1016/S0010-0285(02)00005-1 [PubMed: 12528903]
- Angelone BL, Levin DT, & Simons DJ (2003). The relationship between change detection and recognition of centrally attended objects in motion pictures. *Perception*, 32(8), 947–962. doi: 10.1068/p5079 [PubMed: 14580141]

- Anzures G, Quinn PC, Pascalis O, Slater AM, & Lee K (2010). Categorization, categorical perception, and asymmetry in infants' representation of face race. *Developmental Science*, 13(4), 553–564. doi: 10.1111/j.1467-7687.2009.00900.x [PubMed: 20590720]
- Archambault A, O'Donnell C, Schyns PG (1999). Blind to object changes: When learning the same object at different levels of categorization modifies its perception. *Psychological Science*, 10(3), 249–255. doi: 10.1111/1467-9280.00145
- Bahrack LE (1983). Infants' perception of substance and temporal synchrony in multimodal events. *Infant Behavior and Development*, 6(4), 429–451. doi: 10.1016/S0163-6383(83)90241-2
- Baillargeon R (1991). Reasoning about the height and location of a hidden object in 4.5- and 6.5-month-old infants. *Cognition*, 35(1), 13–42. doi: 10.1016/0010-0277(91)90021-U
- Baillargeon R (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science*, 3(1), 2–13. doi: 10.1111/j.1745-6916.2008.00056.x [PubMed: 22623946]
- Baillargeon R, & DeJong GF (2017). Explanation-based reasoning in infancy. *Psychonomic Bulletin & Review*, 24(5), 1511–1526. doi: 10.3758/s13423-017-1334-4 [PubMed: 28698990]
- Baillargeon R, & DeVos J (1991). Object permanence in young infants: Further evidence. *Child Development*, 62(6), 1227–1246. doi: 10.1111/j.1467-8624.1991.tb01602.x [PubMed: 1786712]
- Baillargeon R, Li J, Gertner Y, & Wu D (2011). How do infants reason about physical events? In Goswami U (Ed.), *The Wiley-Blackwell handbook of childhood cognitive development* (second ed., pp. 11–48). Oxford: Blackwell.
- Baillargeon R, Li J, Luo Y, & Wang S (2006). Under what conditions do infants detect continuity violations? In Johnson MH, & Munakata Y (Eds.), *Processes of change in brain and cognitive development: Attention and Performance XXI* (pp. 163–188). Oxford: Oxford University Press.
- Baillargeon R, Li J, Ng W, & Yuan S (2009a). An account of infants' physical reasoning In Woodward A & Needham A (Eds.), *Learning and the infant mind* (pp. 66–116). New York, NY: Oxford University Press.
- Baillargeon R, Needham A, & DeVos J (1992). The development of young infants' intuitions about support. *Early Development and Parenting*, 1(2), 69–78. doi: 10.1002/edp.2430010203
- Baillargeon R, Scott R, He Z, Sloane S, Setoh P, Jin K, Wu D, & Bian L (2015). Psychological and sociomoral reasoning in infancy In Mikulincer M, Shaver PR (Eds.), Borgida E, Bargh e J. A. (Assoc. Eds.), *APA handbook of personality and social psychology: Vol. 1. Attitudes and social cognition* (pp. 79–150). Washington, DC: American Psychological Association.
- Baillargeon R, Stavans M, Wu D, Gertner Y, Setoh P, Kittredge AK, & Bernard A (2012). Object individuation and physical reasoning in infancy: An integrative account. *Language, Learning and Development*, 5(1), 4–46. doi: 10.1080/15475441.2012.630610
- Baillargeon R, Wu D, Yuan S, Li J, & Luo Y (2009b). Young infants' expectations about self-propelled objects In Hood B & Santos L (Eds.), *The Origins of Object Knowledge* (pp. 285–352). Oxford, England: Oxford University Press.
- Beck MR, & Levin DT (2003). The role of representational volatility in recognizing pre-and postchange objects. *Perception & Psychophysics*, 65(3), 458–468. doi: 10.3758/BF03194575 [PubMed: 12785074]
- Bonatti LL, Frot E, & Mehler J (2005). What face inversion does to infants' counting abilities. *Psychological Science*, 16(7), 506–510. doi: 10.1111/j.0956-7976.2005.01565.x [PubMed: 16008780]
- Bonatti LL, Frot E, Zangl R, & Mehler J (2002). The human first hypothesis: Identification of conspecifics and individuation of objects in the young infant. *Cognitive Psychology*, 44(4), 388–426. doi: 10.1006/cogp.2002.0779 [PubMed: 12018939]
- Booth AE, & Waxman SR (2009). A horse of a different color: Specifying with precision infants' mappings of novel nouns and adjectives. *Child Development*, 50(1), 15–22. doi: 10.1111/j.1467-8624.2008.01242.x
- Buyukozer Dawkins M, Ting F, Stavans M, & Baillargeon R (in press). Early moral cognition: A principle-based approach To appear in Poeppel D, Mangun GR, & Gazzaniga MS, *The cognitive neurosciences VI*. Cambridge, MA: MIT Press.

- Cacchione T, Schaub S, & Rakoczy H (2013). Fourteen-month-old infants infer the continuous identity of objects on the basis of nonvisible causal properties. *Developmental Psychology*, 49(7), 1325–1329. doi: 10.1037/a0029746 [PubMed: 22906060]
- Carey S (2011). *The origin of concepts* (paperback ed.). New York, NY: Oxford University Press.
- Casasola M (2008). The development of infants' spatial categories. *Current directions in Psychological Science*, 17(1), 21–25. doi: 10.1111/j.1467-8721.2008.00541.x
- Cowan N (2016). Working memory maturation: Can we get at the essence of cognitive growth? *Perspective on Psychological Science*, 11(2), 239–264. doi: 10.1177/1745691615621279
- Csibra G, Hernik M, Mascaro O, Tatone D, & Lengyel M (2016). Statistical treatment of looking-time data. *Developmental Psychology*, 52(4), 521–536. doi: 10.1037/dev0000083 [PubMed: 26845505]
- DeJong GF (2014). Explanation-based learning In Gonzalez T, Diaz-Herrera J, & Tucker A (Eds.), *CRC computing handbook: Computer science and software engineering* (3rd ed., pp. 66.1–66.26). Boca Raton: CRC Press.
- Diamond A (2013). Executive functions. *Annual Review of Psychology*, 64(1), 135–168. doi: 10.1146/annurev-psych-113011-143750
- Dueker GL, Modi A, & Needham A (2003). 4.5-month-old infants' learning, retention and use of object boundary information. *Infant Behavior & Development*, 26(4), 588–605. doi: 10.1016/j.infbeh.2003.05.002
- Faul F, Erdfelder E, Lang A, & Buchner A (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. doi: 10.3758/BF03193146 [PubMed: 17695343]
- Feigenson L, & Carey S (2005). On the limits of infants' quantification of small object arrays. *Cognition*, 97(3), 295–313. doi: 10.1016/j.cognition.2004.09.010 [PubMed: 16260263]
- Feigenson L, Carey S, & Hauser M (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological Science*, 13(2), 150–156. doi: 10.1111/1467-9280.00427 [PubMed: 11933999]
- Feigenson L, Dehaene S, & Spelke ES (2004). Core systems of number. *Trends in Cognitive Science*, 8(7), 307–314. doi: 10.1016/j.tics.2004.05.002
- Ferguson B, & Waxman SR (2017). Linking language and categorization in infancy. *Journal of Child Language*, 44(3), 527–552. doi: 10.1017/S0305000916000568 [PubMed: 27830633]
- Flombaum JI, Scholl BJ & Santos LR (2009). Spatiotemporal priority: The engine that drives object persistence In Hood BM & Santos LR (Ed.) *The Origins of Object Knowledge: The Yale Symposium on the Origins of Object & Number Representation*. Oxford: Oxford University Press 135–164.
- Futó J, Téglás E, Csibra G, & Gergely G (2010). Communicative function demonstration induces kind-based artifact representation in preverbal infants. *Cognition*, 117(1), 1–8. doi: 10.1016/j.cognition.2010.06.003 [PubMed: 20605019]
- Gelman R (1990). First principles organize attention to and learning about relevant data - Number and the animate-inanimate distinction as examples. *Cognitive Science*, 14(1), 79–106. doi: 10.1207/s15516709cog1401_5
- Gelman SA (2003). *The essential child: Origins of essentialism in everyday thought*. New York: Oxford University Press.
- Gordon RD, & Irwin DE (1996). What's in an object file? Evidence from priming studies. *Perception & Psychophysics*, 58(8), 1260–1277. doi: 10.3758/BF03207558 [PubMed: 8961836]
- Gordon RD, & Irwin DE (2000). The role of physical and conceptual properties in preserving object continuity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(1), 136–150. doi: 10.1037//0278-7393.26.1.136
- Hauf P, Paulus M, & Baillargeon R (2012). Infants use compression information to infer objects' weights: Examining cognition, exploration, and prospective action in a preferential-reaching task. *Child Development*, 83(6), 1978–1995. doi: 10.1111/j.1467-8624.2012.01824.x [PubMed: 22861050]
- Hespos SJ, & Baillargeon R (2001a). Infants' knowledge about occlusion and containment events: A surprising discrepancy. *Psychological Science*, 12(2), 141–147. doi: 10.1111/1467-9280.00324 [PubMed: 11340923]

- Hespos SJ, & Baillargeon R (2001b). Reasoning about containment events in very young infants. *Cognition*, 78(3), 207–245. doi: 10.1016/S0010-0277(00)00118-9 [PubMed: 11124350]
- Hespos SJ, & Baillargeon R (2006). Décalage in infants' knowledge about occlusion and containment events: Converging evidence from action tasks. *Cognition*, 99(2), B31–B41. doi: 10.1016/j.cognition.2005.01.010 [PubMed: 15939414]
- Hollingworth A (2003). Failures of retrieval and comparison constrain change detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2), 388–403. doi: 10.1037/0096-1523.29.2.388 [PubMed: 12760623]
- Hollingworth A, Williams CC, & Henderson JM (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin & Review*, 8(4), 761–768. doi: 10.3758/BF03196215 [PubMed: 11848597]
- Huttenlocher J, & Lourenco SF (2007). Using spatial categories to reason about location In Plumert J & Spencer J (Eds.), *The Emerging Spatial Mind* (pp. 3–24). New York: Oxford University Press.
- Hyde DC, & Spelke ES (2011). Neural signatures of number processing in human infants: Evidence for two core systems underlying non-verbal numerical cognition. *Developmental Science*, 14(2), 360–371. doi: 10.1111/j.1467-7687.2010.00987.x [PubMed: 21399717]
- Kahneman D, Treisman A, & Gibbs BJ (1992). The reviewing of object files: object- specific integration of information. *Cognitive Psychology*, 24(2), 175–219. doi: 10.1016/0010-0285(92)90007-0. [PubMed: 1582172]
- Káldy Z, & Leslie AM (2003). Identification of objects in 9-month-old infants: Integrating 'what' and 'where' information. *Developmental Science*, 6(3), 360–373. doi: 10.1111/1467-7687.00290
- Keil FC (1995). The growth of causal understandings of natural kinds: Modes of construal and the emergence of biological thought In Sperber D, Premack D & Premack AJ; (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 234–262). Oxford, England: Oxford University Press.
- Kibbe MM, & Feigenson L (2016). Infants use temporal regularities to chunk objects in memory. *Cognition*, 146, 251–263. doi: 10.1016/j.cognition.2015.09.022 [PubMed: 26484498]
- Kibbe MM, & Leslie AM (2011). What do infants remember when they forget? Location and identity in 6-month-olds' memory for objects. *Psychological Science*, 22(12), 1500–1505. doi: 10.1177/0956797611420165 [PubMed: 22095976]
- Kibbe MM, & Leslie AM (2013). What's the object of object working memory in infancy? Unraveling 'what?' and 'how many?' *Cognitive Psychology*, 66(4), 380–404. doi: 10.1016/j.cogpsych.2013.05.001 [PubMed: 23770623]
- Kibbe MM, & Leslie AM (in press). Conceptually rich, perceptually sparse: Object representations in six-month-olds' working memory. *Psychological Science*.
- Kinzler KD, & Spelke ES (2011). Do infants show social preferences for people differing in race? *Cognition*, 119(1), 1–9. doi: 10.1016/j.cognition.2010.10.019 [PubMed: 21334605]
- Kotovsky L, & Baillargeon R (1998). The development of calibration-based reasoning about collision events in young infants. *Cognition*, 67(3), 311–351. doi: 10.1016/S0010-0277(98)00036-5 [PubMed: 9775513]
- Krøjgaard P (2000). Object individuation in 10-month-old infants: Do significant objects make a difference? *Cognitive Development*, 15(2), 169–184. doi: 10.1016/S0885-2014(00)00020-4
- Leslie AM (1995). A theory of agency In Sperber D, Premack D & Premack AJ; (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 121–149). Oxford, England: Oxford University Press.
- Leslie AM, & Keeble S (1987). Do six-month-old infants perceive causality? *Cognition*, 25(3), 265–288. doi: 10.1016/S0010-0277(87)80006-9 [PubMed: 3581732]
- Leslie AM, Xu F, Tremoulet PD, & Scholl BJ (1998). Indexing and the object concept: Developing 'what' and 'where' systems. *Trends in Cognitive Sciences*, 2(1), 10–18. doi: 10.1016/S1364-6613(97)01113-3 [PubMed: 21244957]
- Levin DT, & Simons DJ (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, 4(4), 501–506. doi: 10.3758/BF03214339
- Levine SC, & Baillargeon R (2016). Different faces of language in numerical development: Exact number and individuation In Barner D & Baron AS (Eds.), *Core knowledge and conceptual change* (pp. 127–150). New York, NY : Oxford University Press.

- Luo Y, & Baillargeon R (2005). When the ordinary seems unexpected: Evidence for incremental physical knowledge in young infants. *Cognition*, 95(3), 297–328. doi: 10.1016/j.cognition.2004.01.010 [PubMed: 15788161]
- Luo Y, Kaufman L, & Baillargeon R (2009). Young infants' reasoning about physical events involving inert and self-propelled objects. *Cognitive Psychology*, 55(4), 441–486. doi: 10.1016/j.cogpsych.2008.11.001
- Mandler JM, & McDonough L (1998). On developing a knowledge base in infancy. *Developmental Psychology*, 34(6), 1274–1288. doi: 10.1037/0012-1649.34.6.1274 [PubMed: 9823512]
- Mascalzoni E, Regolin L, Vallortigara G, & Simion F (2013). The cradle of causal reasoning: Newborns' preference for physical causality. *Developmental Science*, 16(3), 327–335. doi: 10.1111/desc.12018 [PubMed: 23587033]
- McCurry S, Wilcox T, & Woods R (2009). Beyond the search barrier: A new task for assessing object individuation in young infants. *Infant Behavior & Development*, 32(4), 429–436. doi: 10.1016/j.infbeh.2009.07.002 [PubMed: 19651444]
- Mitroff SR, Simons DJ, & Levin DT (2004). Nothing compares two views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, 66(8), 1268–1281. doi: 10.3758/BF03194997
- Mou Y, & Luo Y (2017). Is it a container? Young infants' understanding of containment events. *Infancy*, 22(2), 256–270. doi: 10.1111/infa.12148
- Naughtin CK, Mattingley JB, & Dux PE (2016). Distributed and overlapping neural substrates for object individuation and identification. *Cerebral Cortex*, 26(2), 566–575. doi: 10.1093/cercor/bhu212 [PubMed: 25217471]
- Needham A, & Baillargeon R (2000). Infants' use of featural and experiential information in segregating and individuating objects: A reply to Xu, Carey, and Welch (1999). *Cognition*, 74(3), 255–284. doi: 10.1016/S0010-0277(99)00075-X [PubMed: 10640572]
- Needham A, Dueker G, & Lockhead G (2005). Infants' formation and use of categories to segregate objects. *Cognition*, 94(3), 215–240. doi: 10.1016/j.cognition.2004.02.002 [PubMed: 15617672]
- Newman G, Herrmann P, Wynn K, & Keil FC (2008). Biases towards internal features in infants' reasoning about objects. *Cognition*, 107(2), 420–432. doi: 10.1016/j.cognition.2007.10.006 [PubMed: 18068151]
- Noë A, Pessoa L, & Thompson E (2000). Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, 7(1), 93–96. doi: 10.1080/135062800394702
- Pauen S (2002). The global-to-basic level shift in infants' categorical thinking: First evidence from a longitudinal study. *International Journal of Behavioral Development*, 26(6), 492–499. doi: 10.1080/01650250143000445
- Pertsov Y, Dong MY, Peich MC, & Husain M (2012). Forgetting what was where: The fragility of object-location binding. *PLOS One*, 7(10), e48214. doi: 10.1371/journal.pone.0048214
- Pylyshyn Z (1989). The role of location indexes in spatial perception - A sketch of the FINST spatial-index model *Cognition*, 32(1), 65–97. doi: 10.1016/0010-0277(89)90014-0 [PubMed: 2752706]
- Pylyshyn ZW (2007). *Things and places: How the mind connects with the world*. Cambridge, MA: MIT Press.
- Rensink RA (2002). Change detection. *Annual Review of Psychology*, 53, 245–277. doi: 10.1146/annurev.psych.53.100901.135125
- Rivera SM, & Zawaydeh AN (2007). Word comprehension facilitates object individuation in 10- and 11-month-old infants. *Brain Research*, 1146, 146–157. doi: 10.1016/j.brainres.2006.08.112 [PubMed: 17010322]
- Rips LJ, Blok S, & Newman G (2006). Tracing the identity of objects. *Psychological Review*, 113(1), 1–30. doi: 10.1037/0033-295X.113.1.1 [PubMed: 16478299]
- Schaub S, Bertin E, & Cacchione T (2013). Infants' individuation of rigid and plastic objects based on shape. *Infancy*, 18(4), 629–638. doi: 10.1111/infa.12023
- Setoh P, Wu D, Baillargeon R, & Gelman R (2013). Young infants have biological expectations about animals. *Proceedings of the National Academy of Sciences of the United States of America*, 110(40), 15937–15942. doi: 10.1073/pnas.1314075110 [PubMed: 24003134]

- Simons DJ, & Ambinder MS (2005). Change blindness: Theory and consequences. *Current Directions in Psychological Science*, 14(1), 44–48. doi: 10.1111/j.0963-7214.2005.00332.x
- Simons DJ, Chabris CF, Schnur T, & Levin DT (2002). Evidence for preserved representations in change blindness. *Consciousness & Cognition*, 11(1), 78–97. doi: 10.1006/ccog.2001.0533 [PubMed: 11883989]
- Simons DJ, & Levin DT (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review*, 5(4), 644–649. doi: 10.3758/BF03208840
- Spelke ES, Breinlinger K, Macomber J, & Jacobson K (1992). Origins of knowledge. *Psychological Review*, 99(4), 605–632. doi: 10.1037//0033-295X.99.4.605 [PubMed: 1454901]
- Spelke ES, Kestenbaum R, Simons DJ, & Wein D (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, 13(2), 113–142. doi: 10.1111/j.2044-835X.1995.tb00669.x
- Spelke ES, Phillips A, & Woodward AL (1995). Infants' knowledge of object motion and human action In Sperber D, Premack D & Premack AJ (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 44–78). Oxford, England: Oxford University Press.
- Stahl AE, & Feigenson L (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94. doi: 10.1126/science.aaa3799 [PubMed: 25838378]
- Stavans M, & Baillargeon R (2018). 4.5-month-old infants individuate and track simple tools following functional demonstrations. *Developmental Science*, 21(1), e12500. doi: 10.1111/desc.12500
- Strickland B, & Scholl BJ (2015). Visual perception involves event-type representations: The case of containment versus occlusion. *Journal of Experimental Psychology. General*, 144(3), 570–580. doi: 10.1037/a0037750 [PubMed: 25867222]
- Surian L, & Caldi S (2010). Infants' individuation of agents and inert objects. *Developmental Science*, 13(1), 143–150. doi: 10.1111/j.1467-7687.2009.00873.x [PubMed: 20121870]
- Taborda-Osorio H, & Cheries EW (2018). Infants' agent individuation: It's what's on the insides that counts. *Cognition*, 175, 11–19. doi: 10.1016/j.cognition.2018.01.016 [PubMed: 29454257]
- Ullman TD, Spelke ES, Battaglia P, & Tenenbaum JB (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21(9), 649–665. doi: 10.1016/j.tics.2017.05.012 [PubMed: 28655498]
- Van de Walle GA, Carey S, & Prevor M (2000). Bases for object individuation in infancy: Evidence from manual search. *Journal of Cognition and Development*, 1(3), 249–280. doi: 10.1207/S15327647JCD0103_1
- vanMarle K (2013). Infants use different mechanisms to make small and large number ordinal judgments. *Journal of Experimental Child Psychology*, 114(1), 102–110. doi: 10.1016/j.jecp.2012.04.007 [PubMed: 22608189]
- vanMarle K, Mou Y, & Seok JH (2016). Analog magnitudes support large number ordinal judgments in infancy. *Perception*, 45(1–2), 32–43. doi: 10.1177/0301006615602630 [PubMed: 26562862]
- Varakin AD, Levin DT, & Collins KM (2007). Comparison and representation failures both cause real-world change blindness. *Perception*, 36(5), 737–749. doi: 10.1068/p5572 [PubMed: 17624119]
- Wang S (2011). Priming 4.5-month-old infants to use height information by enhancing retrieval. *Developmental Psychology*, 47(1), 26–38. doi: 10.1037/a0021060. [PubMed: 21058828]
- Wang S, & Baillargeon R (2006). Infants' physical knowledge affects their change detection. *Developmental Science*, 9(2), 173–181. doi: 10.1111/j.1467-7687.2006.00477.x [PubMed: 16472318]
- Wang S, & Baillargeon R (2008). Can infants be “taught” to attend to a new physical variable in an event category? The case of height in covering events. *Cognitive Psychology*, 56(4), 284–326. doi: 10.1016/j.cogpsych.2007.06.003 [PubMed: 18177635]
- Wang S, Baillargeon R, & Brueckner L (2004). Young infants' reasoning about hidden objects: Evidence from violation-of-expectation tasks with test trials only. *Cognition*, 93(3), 167–198. doi: 10.1016/j.cognition.2003.09.012 [PubMed: 15178376]
- Wang S, Baillargeon R, & Paterson S (2005). Detecting continuity violations in infancy: a new account and new evidence from covering and tube events. *Cognition*, 95(2), 129–173. doi: 10.1016/j.cognition.2002.11.001 [PubMed: 15694644]

- Wang S, & Goldman EJ (2016). Infants actively construct and update their representations of physical events: Evidence from change detection by 12-month-olds. *Child Development Research*, 2016, article ID 3102481, 1–11. doi:10.1155/2016/3102481
- Wang S, & Kohne L (2007). Visual experience enhances infants' use of task-relevant information in an action task. *Developmental Psychology*, 43(6), 1513–1522. doi: 10.1037/0012-1649.43.6.1513 [PubMed: 18020828]
- Wang S, & Mitroff SR (2009). Preserved visual representations despite change blindness in infants. *Developmental Science*, 12(5), 681–687. doi: 10.1111/j.1467-7687.2008.00800.x [PubMed: 19702760]
- Wang S, & Onishi KH (2017). Enhancing young infants' representations of physical events through improved retrieval (not encoding) of information. *Journal of Cognition & Development*, 18(2), 289–308. doi: 10.1080/15248372.2017.1281279
- Wang S, Zhang Y, & Baillargeon R (2016). Young infants view physically possible support events as unexpected: New evidence for rule learning. *Cognition*, 157, 100–105. doi: 10.1016/j.cognition.2016.08.021 [PubMed: 27599219]
- Waxman SR & Grace AD (2012). Developing gender- and race-based categories in infants: Evidence from 7- and 11-month-olds In Hayes G & Bryant M (Eds.), *Psychology of culture: Psychology of emotions, motivations and actions* (pp. 159–175). Hauppauge, NY: Nova Science Publishers
- Welder AN, & Graham SA (2006). Infants' categorization of novel objects with more or less obvious features. *Cognitive Psychology*, 52(1), 57–91. doi: 10.1016/j.cogpsych.2005.05.003 [PubMed: 16246319]
- Wilcox T (1999). Object individuation: infants' use of shape, size, pattern, and color. *Cognition*, 72(2), 125–166. doi: 10.1016/S0010-0277(99)00035-9 [PubMed: 10553669]
- Wilcox T (2007). Sex differences in infants' mapping of complex occlusion sequences: Further evidence. *Infancy*, 12(3), 303–327. doi: 10.1080/15250000701518727
- Wilcox T, & Baillargeon R (1998). Object individuation in infancy: The use of featural information in reasoning about occlusion events. *Cognitive Psychology*, 37(2), 97–155. doi: 10.1006/cogp.1998.0690 [PubMed: 9878104]
- Wilcox T, & Biondi M (2015). Object processing in the infant: Lessons from neuroscience. *Trends in Cognitive Sciences*, 19(7), 406–413. doi: 10.1016/j.tics.2015.04.009 [PubMed: 26008625]
- Wilcox T, & Chapa C (2002). Infants' reasoning about opaque and transparent occluders in an individuation task. *Cognition*, 55(1), B1–B10. doi: 10.1016/S0010-0277(02)00055-0
- Wilcox T, & Chapa C (2004). Priming infants to attend to color and pattern information in an individuation task. *Cognition*, 90(3), 265–302. doi: 10.1016/S0010-0277(03)00147-1 [PubMed: 14667698]
- Wilcox T, Haslup JA, & Boas DA (2010). Dissociation of processing of featural and spatiotemporal information in the infant cortex. *Neuroimage*, 53(4), 1256–1263. doi: 10.1016/j.neuroimage.2010.06.064 [PubMed: 20603218]
- Wilcox T, & Schweinle A (2002). Object individuation and event mapping: Developmental changes in infants' use of featural information. *Developmental Science*, 5(1), 132–150. doi: 10.1111/1467-7687.00217
- Wilcox T, & Schweinle A (2003). Infants' use of speed information to individuate objects in occlusion events. *Infant Behavior and Development*, 26(2), 253–282. doi: 10.1016/S0163-6383(03)00021-3
- Wilcox T, Schweinle A, & Chapa C (2003). Object individuation in infancy In Fagan F & Hayne H (Eds.). *Progress in Infancy Research* (Vol 3, pp. 193–243). Mahwah, NJ: Lawrence Erlbaum Associates
- Wilcox T, Stubbs J, Hirshkowitz A, & Boas DA (2012). Functional activation of the infant cortex during object processing. *Neuroimage*, 62(3), 1833–1840. doi: 10.1016/j.neuroimage.2012.05.039 [PubMed: 22634218]
- Xu F (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, 55(3), 223–250. doi: 10.1016/S0010-0277(02)00109-9
- Xu F (2007). Sortal concepts, object individuation, and language. *Trends in Cognitive Sciences*, 11(9), 400–406. doi: 10.1016/j.tics.2007.08.002 [PubMed: 17698404]

- Xu F, & Baker A (2005). Object individuation in 10-month-old infants using a simplified manual search method. *Journal of Cognition and Development*, 6(3), 307–323. doi: 10.1207/s15327647jcd0603_1
- Xu F, & Carey S (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, 30(2), 111–153. doi: 10.1006/cogp.1996.0005 [PubMed: 8635312]
- Xu F, & Carey S (2000). The emergence of kind concepts: A rejoinder to Needham and Baillargeon (2000). *Cognition*, 74(3), 285–301. doi: 10.1016/S0010-0277(99)00076-1 [PubMed: 10640573]
- Xu F, Carey S, & Quint N (2004). The emergence of kind-based object individuation in infancy. *Cognitive Psychology*, 49(2), 155–190. doi: 10.1016/j.cogpsych.2004.01.001 [PubMed: 15304370]
- Zosh JM, & Feigenson L (2012). Memory load affects object individuation in 18-month-old infants. *Journal of Experimental Child Psychology*, 113(3), 322–336. doi: 10.1016/j.jecp.2012.07.005 [PubMed: 22932473]
- Zosh JM, & Feigenson L (2015). Array heterogeneity prevents catastrophic forgetting in infants. *Cognition*, 136, 365–380. doi: 10.1016/j.cognition.2014.11.042 [PubMed: 25543889]

Early Failures at Individuation Tasks

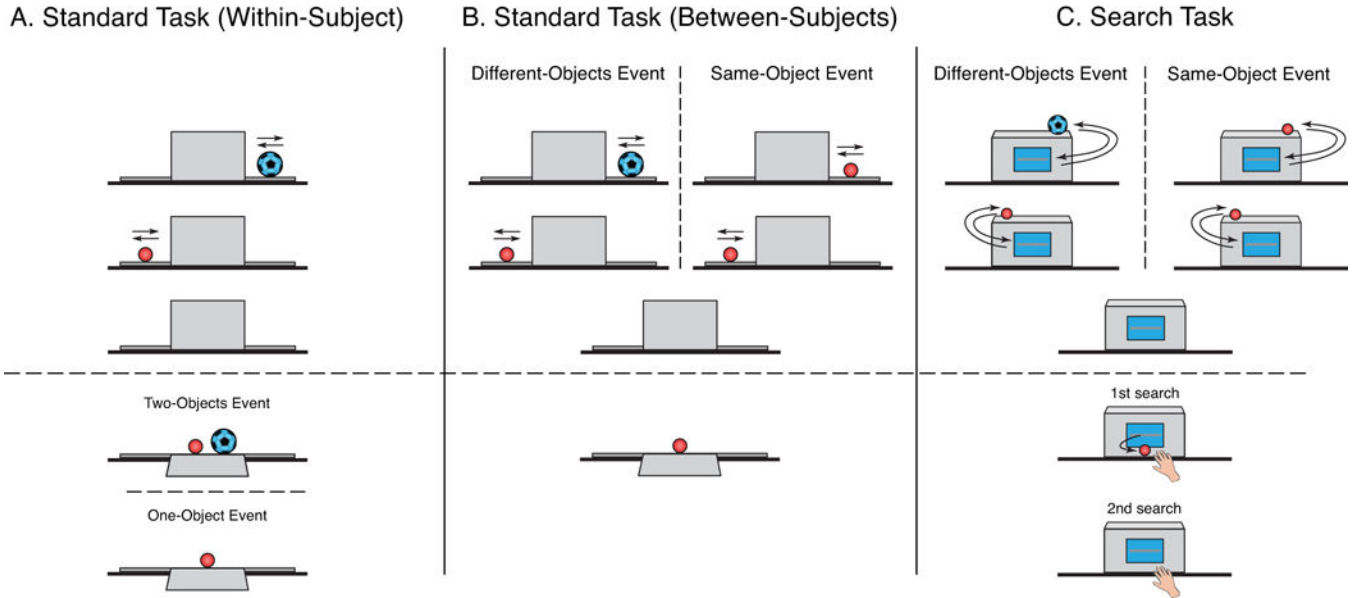


Figure 1. Schematic depiction of test trials in individuation tasks at which infants typically fail. **(A)** In standard within-subject tasks, infants first see two different objects emerge in alternation from behind a screen. On different trials, the screen is lowered to reveal either both objects (two- objects event) or only one object (one-object event). **(B)** In standard between-subjects tasks, infants first see either two different objects in alternation (different-objects event) or the same object (same-object event) emerge on either side of a screen. Next, the screen is lowered to reveal only one object. **(C)** On different trials in search tasks, infants see two different objects in alternation (different-objects event) or the same object (same-object event) removed from and returned to a box. Next, the box is moved within infants’ reach, and they are encouraged to search. Of interest is whether infants continue searching after the first object is retrieved.

Early Successes at Individuation Tasks

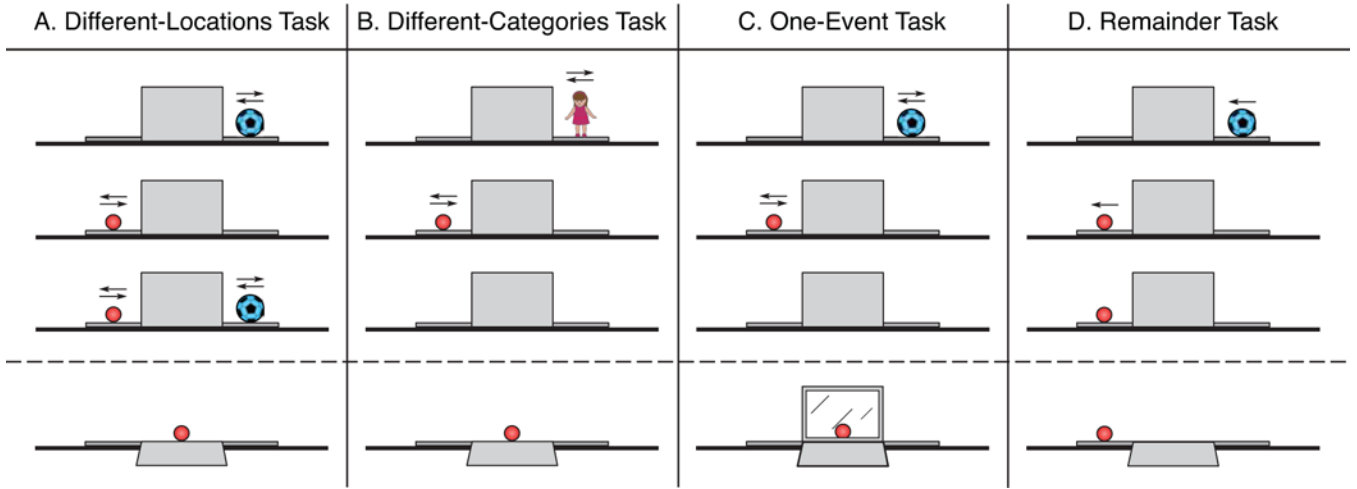


Figure 2. Schematic depiction of test trials in individuation tasks at which infants typically succeed; for ease of communication, these tasks are illustrated using standard between-subjects different-objects events. **(A)** In different-locations tasks, the two objects emerge simultaneously from behind the screen at some point and thus are seen to occupy visibly distinct locations. **(B)** Different-categories tasks make use of objects infants spontaneously assign to distinct categories (more fine-grained categories are used with age). **(C)** In one-event tasks, a transparent panel stands behind the screen. When the screen is lowered, the object is revealed resting behind this transparent occluder. **(D)** In remainder tasks, the second object is not returned behind the screen but is either left in plain view next to the screen or is removed from the scene.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Event		Task	OF Number of Hidden Objects	PR Number of Hidden Objects	OF and PR Agree?	Infants' Expectation
Same-Object	A	Standard and Search	1	1	Agree	1
Different-Objects	B	Different-Locations Standard and Search	2	2	Agree	2
	C	Different-Categories Standard and Search	2	2	Agree	2
	D	Same-Category Standard and Search	1	2	Disagree quantitatively	Catastrophic failure: No expectation
	E	Same-Category One-Event	1	2	Not applicable: Ongoing event, PR has priority	2
	F	Same-Category Remainder	0	1	Disagree qualitatively	1

Figure 3.

Infants' performance in different individuation tasks according to our model. Each row represents a different individuation task. Columns (left to right) describe the event shown (same-object or different-objects) and the task used. The next columns specify the number of hidden objects posited by the object-file (OF) system, the number of hidden objects posited by the physical-reasoning (PR) system, whether these two numbers agree or disagree, and what are the consequences of an agreement or disagreement for infants' expectations about how many objects should be present.

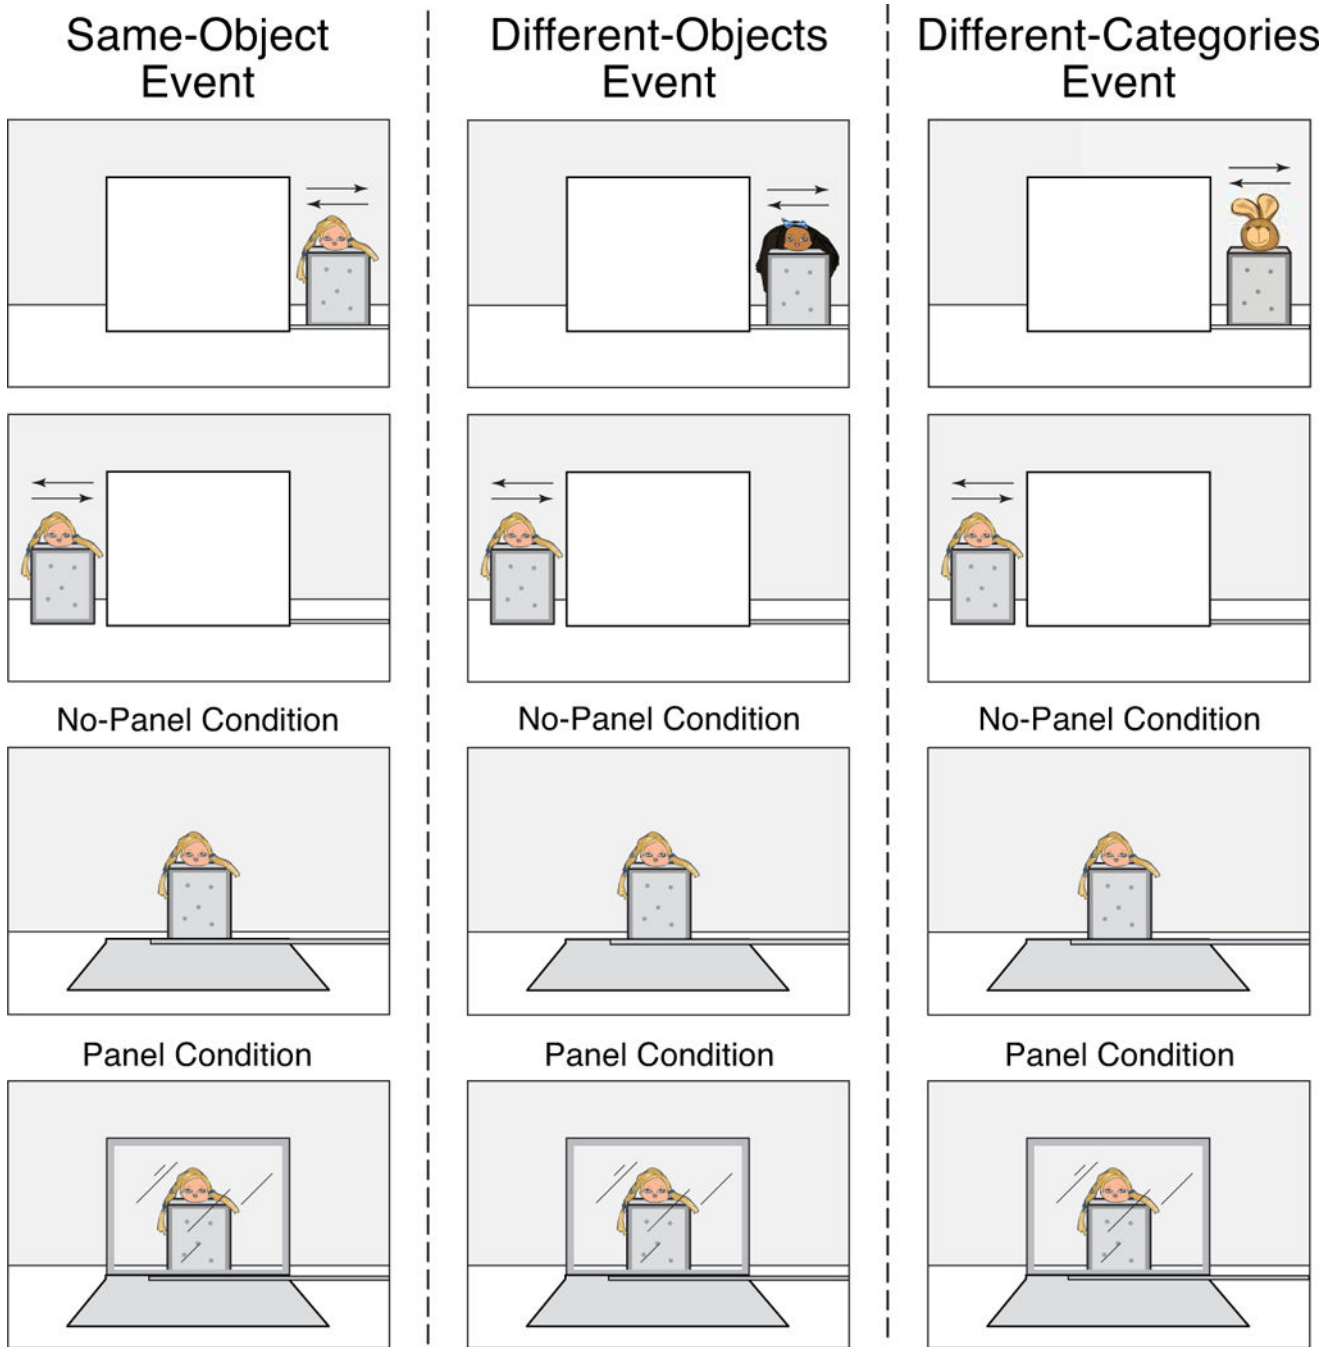


Figure 4. Schematic depiction of the events presented during the test trial of the no-panel and panel conditions in Experiment 1.




Figure 5. Mean looking times in Experiment 1, by event and condition, during the final phase of the test trial. Errors bars represent standard errors, and an asterisk denotes a significant difference between the conditions. Looking times were log-transformed before analysis.




Figure 6. Schematic depiction of the events presented during the test trial of the one-object and no-object conditions in Experiment 2.

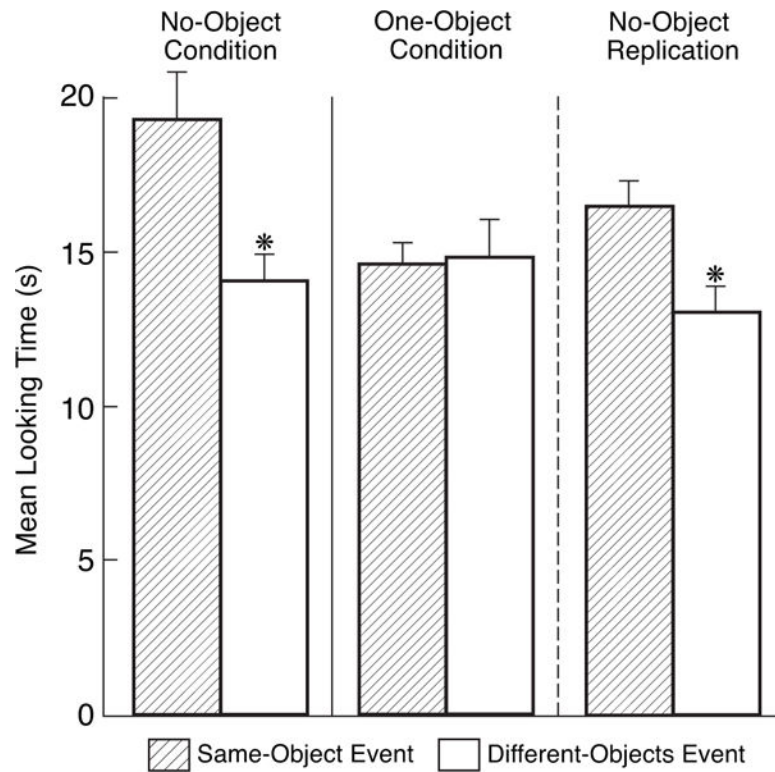


Figure 7. Mean looking times in Experiment 2, by event and condition, during the final phase of the test trial. Errors bars represent standard errors, and an asterisk denotes a significant difference between the conditions. Looking times were log-transformed before analysis.

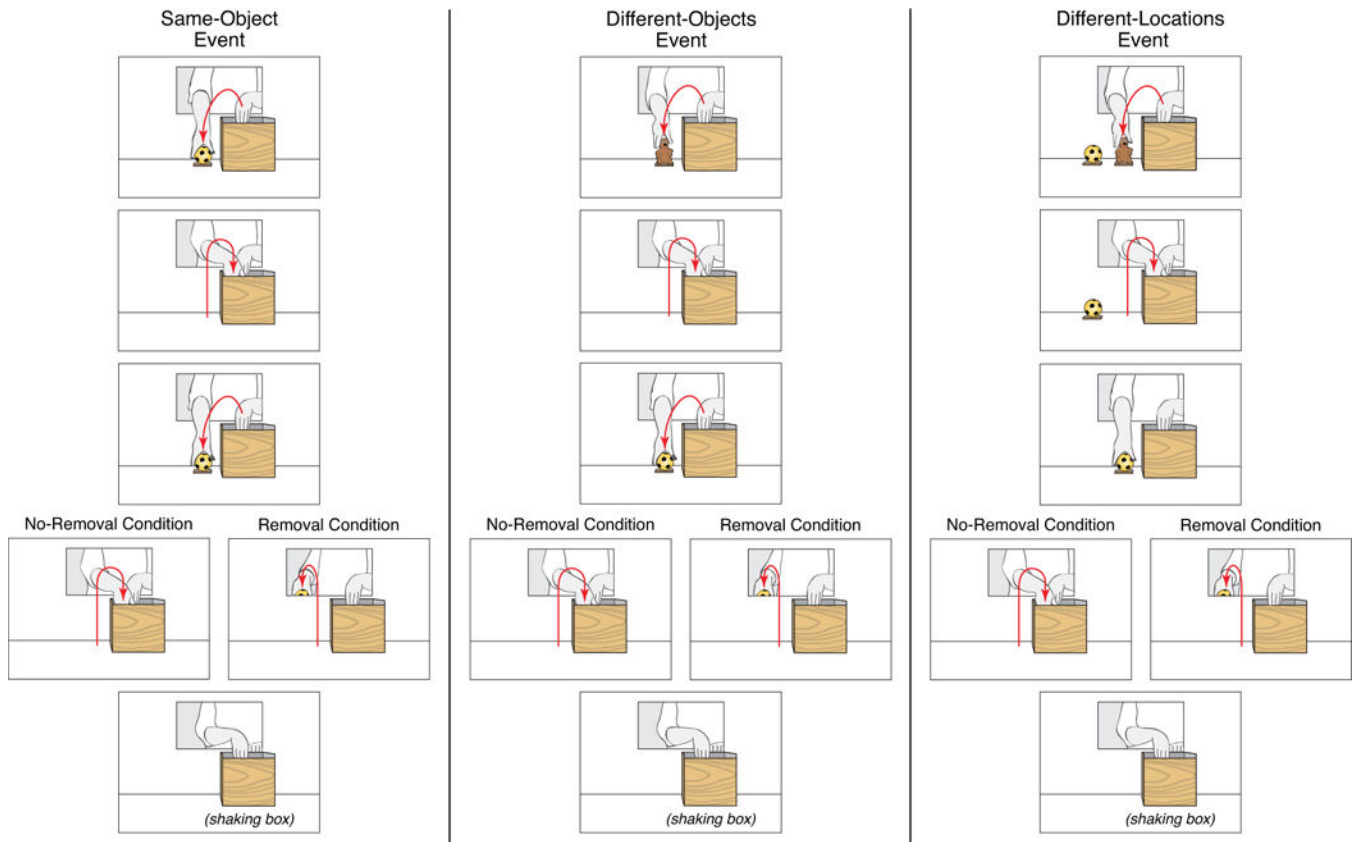


Figure 8. Schematic depiction of the events presented during the test trial of the no-removal and removal conditions in Experiment 3.

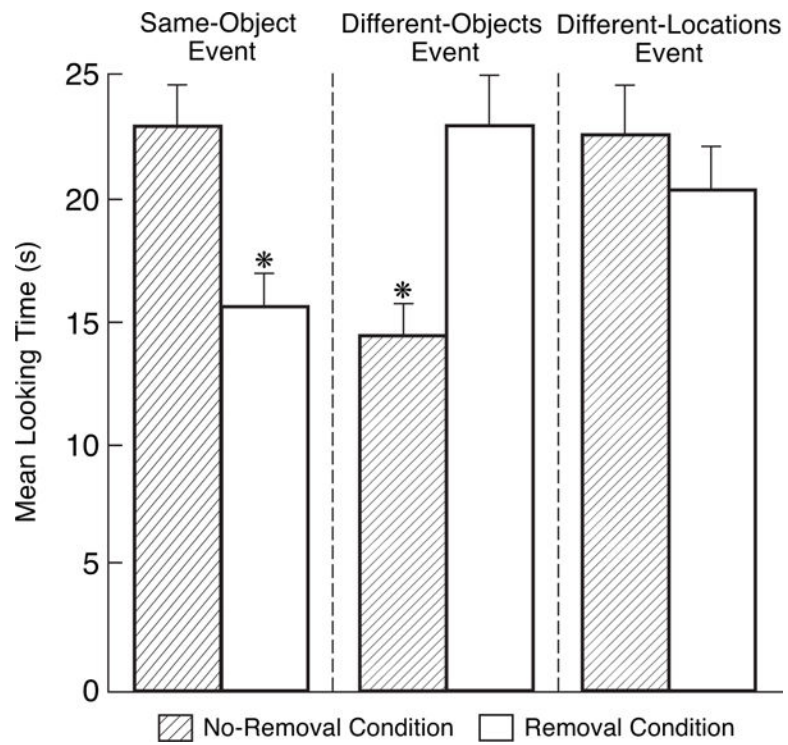


Figure 9. Mean looking times in Experiment 3, by event and condition, during the final phase of the test trial. Errors bars represent standard errors, and an asterisk denotes a significant difference between the conditions. Looking times were log-transformed before analysis.