3

Continuity of Format and Computation in Short-Term Memory Development

LISA FEIGENSON

In their first year of life, infants face the enormous task of making sense of the world around them. Without the ability to store and reason about representations of the individual objects, actions, and sounds in their environment, infants would never accomplish the monumental changes they do. Storing representations of such individuals in memory allows infants to perform computations that, while seemingly simple, are critical to learning about the world. Comparing a scene to one observed earlier, keeping track of the presence of objects even when the objects are temporarily occluded, and making predictions about the outcomes of hidden events are some examples of such computations. The thesis of this chapter is that the short-term memory system that enables infants to store object representations, and many of the computations infants perform over these representations, are continuous throughout the human life span. Infants and adults show similar capacities and similar limitations regarding their ability to represent and reason about objects. At the same time, infants' and adults' short-term memory abilities may differ in some important respects. This chapter explores what is shared and what may differ in early versus mature short-term memory.

First, I would like to lay out some rough definitions. When I talk about short-term memory in this chapter, I am referring to the ability to form and store mental tokens that stand for entities in the outside world. Maintaining these tokens over short durations allows the entities to be thought about even when direct perceptual information is absent, as is the case when objects undergo occlusion. This short-term memory enables infants to represent the presence of entities when those entities might be temporarily hidden (e.g., "There is an object under that blanket") and to store information about those entities (e.g., "The object under the blanket is round and red and striped"). Furthermore, storing representations of more than one item at a time enables infants to compute across an entire scene (e.g., "There are three balls under the blanket"), rather than over just a single item. Here, I suggest that this type of short-term memory in infants corresponds to a system of short-term memory that has been studied in adults (for an argument that *working memory* is a better term for this same system, see chapter 1, this volume). Indeed, one of the most striking observations about this memory system is the extent to which it remains constant throughout development, both in its capacity and in the computations it supports.

I begin with an exploration of the limits on infants' and adults' short-term memory capacity; the evidence I review builds the case that the very same system of memory is relied on across the life span. Second, I examine findings that infants and adults perform similar computations over these short-term memory representations. Both groups compute the continuous and discrete properties of object arrays, and both groups use chunking as a means to recode memory representations into a more efficient format. Third, I address a possible developmental difference in short-term memory, asking whether infants and adults differ in the degree to which their memory computations are driven in top-down versus bottom-up fashion. Finally, I close the chapter by raising some outstanding questions and by suggesting some avenues for future research on short-term memory development.

Short-Term Memory for Object Arrays

Some 25 years ago, pioneers in the newly emerging field of infant cognition demonstrated that, contrary to Piaget's claims, young infants represent and reason about objects. Critically, they do so in ways that go beyond the immediate sensory data available to them. By 5 to 7 months of age, infants represent objects that have been covered by cloths, hidden by screens, or concealed in darkness (e.g., Baillargeon, Spelke, & Wasserman, 1985; Hood & Willats, 1986; Shinskey & Munakata, 2003). That infants have stored representations of these objects in memory is shown by their continued reaching for the objects once hidden, or by their longer looking when objects unexpectedly disappear. Even more impressively, infants reason over these stored representations of objects. For example, when a solid object is placed behind a screen and a second object is launched on a direct path toward it, infants look longer when the second object emerges magically unscathed from the other side of the screen (Baillargeon, 1986). Apparently, infants have reasoned that one solid object cannot pass through another. Because infants' looking times in situations such as these depend on inferred interactions between objects that are hidden rather than visible, we conclude that infants are operating on representations stored in memory rather than operating directly on the immediate sensory data.

Researchers have attempted to characterize the memory systems infants use to reason about such instances of object behavior. One question of interest is whether infants reason over object representations stored in short-term memory or long-term memory.[1] Although the distinction between short- and long-term memory systems has often been controversial, many have suggested that the amount of information stored by each person over long durations is too large and unwieldy to allow sufficiently rapid access for the moment-to-moment comparisons we constantly perform, and which characterize the occlusion events typical of experiments with infants. This problem motivates the existence of a system that is distinct from the larger, long-term memory store. This system holds a limited amount of information in a temporary state of privileged access. Representations held in this short-term memory system can be formed quickly but decay over time, whereas long-term memory representations take longer to form but are far more enduring (see chapters 7–10, this volume). On some models, the information held in short-term memory can come either from the outside world (e.g., storing a representation of an object that is currently visible), or it can come from the activation of a representation previously stored in long-term memory (e.g., thinking about an object that was seen yesterday; Cowan, 2001). In either case, information held in short-term memory is available for immediate processing.

Several factors hint that short-term memory does underlie infants' reasoning about the kinds of object arrays typical of infant cognition research. In such studies, infants receive only brief exposure to a scene before objects are hidden from view, perhaps limiting the extent to which they have the opportunity to store long-term representations. In addition, experiments manipulating the delay between an object's disappearance from view and the moment when infants are allowed to retrieve it reveal that infants' object memories fade rapidly (Diamond, 1990). These factors begin to suggest that short-term memory is the likely locus of infants' object-tracking abilities. However, most studies investigating infants' object representations have not been designed to distinguish the relative contributions of short- versus long-term memory. For example, they have systematically manipulated neither the duration of infants' object exposure nor the interval over which infants must maintain the object representations in memory, at least not in ways that would bear decisively on which memory system is involved. Therefore, the conclusion that many studies of infants' object representations are tapping short-term memory remains tentative.

Short-Term Memory Capacity in Adults

Further evidence is needed. Measuring the capacity of infants' memory is a potential source of such evidence, since capacity differences have traditionally been a distinguishing characteristic of short- versus long-term memory. While long-term memory is usually thought of as unlimited, short-term memory is thought to store only a small amount of information at any one time. This notion of a limited-capacity short-term memory store originated with James (1890) and received significant attention following Miller's (1956) influential postulation of a "magical number seven, plus or minus two." Miller made famous the view of short-term memory as a repository limited by the number of unique items it can hold (approximately seven, in Miller's view), rather than by overall informational load (where information load is influenced by factors such as the complexity of the items). Since then it has been suggested that "seven, plus or minus two" probably overestimates short-term memory capacity and is likely the by-product of additional mental processes, such as chunking, that allow subjects to recode individual items into groups. A more reasonable estimate, obtained when chunking is prevented, is three to four (Cowan, 2001).

Cowan (2001) reviewed a wide range of experiments probing short-term memory capacity, and finds that most of these produce estimates of three to four items. Space precludes presenting that evidence here, but a summary of a classic experimental series by Sperling (1960) illustrates the type of results obtained. Sperling showed adult subjects 3×4 arrays of letters, presented too briefly for the subjects to store all 12 letters in long-term memory. On whole-report trials, subjects reported the names of all of the letters they could remember; they averaged around 4. On partial-report trials, subjects reported only a subset of the array as specified by an auditory cue. When the cue was heard 2 to 5 seconds after the array had disappeared, subjects were able to report an average of 1.3 of the 4 letters in the cued row.[2] Multiplied by the number of rows in the array (3), this again yielded approximately 4 as the upper limit on short-term memory capacity.

More recently, short-term memory tasks produced a similar capacity limit in adults. Halberda, Simons, and Wetherhold (2006) showed subjects a rapidly flashing grid of 32 dots, each of a different luminance value. All but one of the dots maintained its individual luminance value from flash to flash; the remaining dot oscillated between two different values. Subjects had to find the single changing dot. Halberda et al. found that subjects were able to encode the luminance values of a subset of the dots on each flash of the array, store them in memory, and compare them to the dots' luminance values on the next flash. Subjects examined subsets of dots in this way until they located the changing item. By analyzing the number of flashes required to locate the target dot, Halberda et al. estimated that the number of dots subjects could store in short-term memory from a single flash of the array was three.

Other findings strengthen the claim that the capacity limit is defined by the discrete number of items being held in memory, rather than by total information load. Luck and Vogel (1997) used a change detection task in which subjects received a 100-millisecond visual exposure to an array of items, followed 900 milliseconds later by a test array. Subjects had to report whether the two arrays were identical or whether any of the items had changed their features. Luck and Vogel found that performance was at ceiling for arrays containing one to three items, and declined with sets of four or more. Most strikingly, subjects were just as accurate for arrays of items that contained multiple features (color and texture orientation) as for arrays containing just a single feature (color). Thus, the number of items in the array, and not the number of features in the array, determined subjects' memory capacity.

Several researchers have suggested that there is a similar three-to-four-item capacity limit in attention, prior to the storage of any items in short-term memory (Carey, 2004; Scholl & Leslie, 1999; Trick & Pylyshyn, 1994). This view derives largely from results of the multiple object tracking paradigm. This paradigm was developed by Pylyshyn to examine the process by which a subset of the information in a scene achieves priority for further processing, before its transfer into memory (Pylyshyn, 1994, 2001). In the multiple object tracking task, subjects track several moving onscreen targets amid a field of identical distractors. Because no featural cues distinguish the targets from the distractors, the only way for subjects to succeed is to attend to the targets from the start of each trial (when targets flash briefly to indicate their status) and to keep attending to them in parallel as they move haphazardly through the scene. Subjects perform this task effortlessly with one, two, three, and often four targets. But when asked to track more than four, performance plummets (Pylyshyn & Storm, 1988). Because this task was designed to require little or no memory, many have concluded that the observed limit is grounded in attention. However, more recent evidence disputes the view of a three-to-four-item limit on attention. Experiments by Alvarez and Franconeri (2005) suggest that attentional capacity increases to well above four items when the items move more slowly. Thus it remains to be seen whether an item-based limit on attention will hold.

If such a limit does hold up in purely attentional tasks, is the existence of an identical three-to-four-item limit that constrains both attention and short-term memory purely coincidental? An alternative view has been offered by Cowan (2001), who suggests that there is no structural distinction between attention and short-term memory. Instead, Cowan suggests that in order to reason about remembered items, the items need to be pulled from memory storage into the "focus of attention." This focus of attention can be thought of as activation of the stored items, where activation is required for any sort of conscious processing. Cowan suggests that while memory storage itself is unlimited, only three to four items can be brought into the focus of attention at any given time. Cowan's proposal is controversial but serves to highlight the difficulty in distinguishing between capacity limits on attention versus those on short-term memory.

Measuring Short-Term Memory in Infancy: Recent Advances

We now return to the question of which memory system underlies infants' ability to track and reason about hidden objects. Given that adults can maintain three to four items in short-term memory, a similar limit on infants' abilities would be important in illustrating continuity across development. Recent findings have obtained just such a limit. In a modified version of

Luck and Vogel's (1997) change-detection paradigm, Ross-Sheehy, Oakes, and Luck (2003) presented 10-month-old infants with cycles of simple object arrays appearing simultaneously on a pair of adjacent screens. On every cycle, the arrays appeared on each screen for 500 milliseconds, disappeared for 250 milliseconds, then reappeared for 500 milliseconds. On one of the two screens, one of the objects in the array changed color during the 250-millisecond retention period. On the other screen, none of the objects changed. The arrays cycled such that one screen always displayed an array that changed during the retention period, while the other screen always displayed an array that remained constant (for a more detailed description of their experimental design, see chapter 4, this volume).

Ross-Sheehy et al. (2003) reasoned that since infants generally prefer to look at more complex displays rather than at simple ones, they would spend more time looking at the screen with the changing array than at the screen with the constant array. The ability to notice a change in the array from cycle to cycle depended on storing a representation of the objects in the array, maintaining this representation over the 250 milliseconds when no display was visible, and then comparing it to the next array that appeared. Therefore, a preference for looking at the changing screen implies successful memory retention of the features of all of the objects in the array (since which particular object changed varied randomly from cycle to cycle). Ross-Sheehy and colleagues used this paradigm to probe memory limits by comparing infants' performance with arrays containing different numbers of objects. They found that with one-, two-, three-, and four-object arrays, infants preferred to look at the changing screen. But with six-object arrays, infants showed no such preference. Apparently, 10-month-old infants were unable to represent more than four items at a time and therefore did not discriminate the changing from the unchanging array.

This work, using methods that closely resemble those used to study adults' visual short-term memory, suggests that infants, like adults, can store representations of three to four items at a time. But what about memory for the kinds of real, three-dimensional object arrays used in so many experiments on infant cognition, and which constitute the majority of infants' natural daily experience? Objects in such arrays are likely to be more complex in their shape, shading, and features than the simple squares used by Ross-Sheehy et al. (2003). In addition, objects in a natural scene often undergo complex patterns of motion, sometimes involving periods of occlusion during which they might be hidden for several seconds at a time. Will the three-to-four-item limit of short-term memory also be observed when infants are faced with naturalistic object arrays?

Infants' Short-Term Memory for Naturalistic Object Arrays

The question of exactly how many such hidden objects infants can remember and reason over was first raised by Karen Wynn, who demonstrated that by 5 months of age infants already can represent at least two occluded objects. After seeing a screen hide one object, followed by the addition of a second object behind the screen, infants expected to see two objects when the screen was lifted (Wynn, 1992). Work using similar looking time methods to measure infants' representation of occluded objects has confirmed that infants can remember at least two hidden objects at a time (Kaldy & Leslie, 2003; Koechlin, Dehaene, & Mehler, 1997; Simon, Hespos, & Rochat, 1995; Uller, Huntley-Fenner, Carey, & Klatt, 1999).

My colleagues and I have extended these findings by asking just how many such hidden objects infants can remember. We probed the upper limits of infants' ability to track occluded objects by creating a simple task in which the number of objects infants had to store in memory was parametrically varied. In this "cracker choice" task (Feigenson & Carey, 2005; Feigenson, Carey, & Hauser, 2002), 10- and 12-month-old infants saw two quantities of desirable objects (graham crackers) sequentially placed into a pair of opaque buckets and then were allowed to choose between them. Since determining which bucket contained more crackers required maintaining and comparing representations of the hidden objects, and since adults have been shown to store object representations in short-term memory for durations comparable to those we used (approximately 8–10 seconds; Noles, Scholl, & Mitroff, 2005), we reasoned that our procedure would serve as a naturalistic test of preverbal children's short-term memory.

We started by giving infants a choice between one or two, two or three, or three or four crackers. Groups of 16 different 10-month-old infants and 16 different 12-month-old infants participated in each of these numerical comparisons. Infants' spontaneous, untrained abilities were revealed by giving each infant just one opportunity to make a choice; thus, the experiment consisted of a single trial for each participant. In our experimental procedure, infants sat on the floor across from an experimenter. The experimenter produced two opaque plastic buckets, showed infants that they were empty, and placed them on the floor approximately 70 cm from infants' starting location and approximately 70 cm from each other. The experimenter then placed crackers one at a time into the buckets, making sure that infants attended to the placement of each cracker. For example, in a one-or-two choice, the experimenter placed one cracker in one bucket and two crackers one at a time into the other bucket. Which side the presentation began on and which bucket received the greater number of crackers was counterbalanced across participants. The dependent measure was simply which bucket infants chose to walk or crawl to.

Figure 3.1 displays the pattern of infants' spontaneous choices. We found that with choices of one or two and two or three crackers, infants of both age groups successfully chose the bucket containing the greater quantity. Infants failed, however, with a choice of three or four crackers. Infants' failure with this comparison might have been due to either the less discriminable ratio between the quantities, or to the quantities having exceeded the maximum

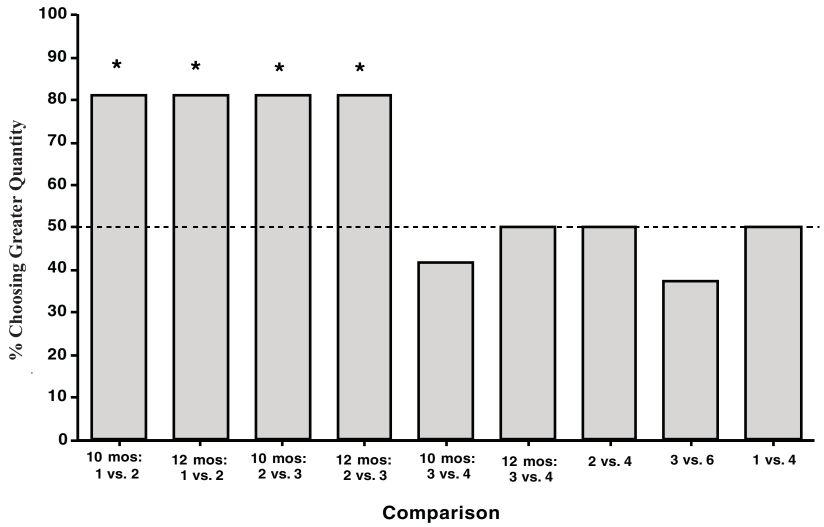

FIGURE 3.1. The percentage of infants choosing the greater of two quantities. Infants chose the greater quantity with small arrays, but failed whenever either array contained four or more objects. Reprinted from Feigenson, L., Carey, S., & Hauser, M. (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological Science, 13*(2), 150–156, copyright (2002), with permission from Blackwell; and from Feigenson, L., & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition, 97,* 295–313, copyright 2005, with permission from Elsevier.

number of items infants could hold in memory. Therefore, we next tested new groups of infants with quantities in a ratio with which we had already observed success (1:2), but which involved more than three items in a single location. We gave infants a choice between two or four, three or six, and one or four crackers. Since no age differences between 10- and 12-month-olds were observed in any of the previous comparisons we tested, each of the comparisons included 16 infants ranging between 10 and 12 months.

Infants failed to choose systematically with any of these quantities (see Figure 3.1). This dramatic breakdown in performance illustrates that infants' ability to remember the hidden objects was determined by the total number of objects seen, and not by ratio of differences between the two quantities we presented. Infants succeeded only when one, two, or three crackers were placed in either bucket, and chose entirely by chance when required to remember larger numbers (Figure 3.1). A variety of control conditions ensured that the abrupt break in infants' performance was, in fact, due to the number of objects presented and not to the total presentation duration or complexity (see Feigenson, Carey, & Hauser, 2002, for details). Thus, it appears that in this task infants were limited to storing up to three items in each hiding location.

This naturalistic cracker choice task, which required infants to track real, three-dimensional objects undergoing motion and occlusion, revealed the same three-item limit on performance as has been observed in tasks showing infants (Ross-Sheehy et al., 2003) and adults (Halberda et al., 2006; Halberda, Sires, & Feigenson, 2006; Luck & Vogel, 1997) simple, computerized arrays. This suggests that the three-to-four-item capacity limit applies to a range of entities, from grayscale dots to real moving objects. However, one difference between our cracker choice task and previous tasks assessing short-term memory lies in the timing of the presentation. Our cracker task involved sequentially presented objects, whereas previous tasks presented infants and adults with objects that were all visible at once. Therefore, our next step was to ask whether the three-to-four-item capacity limit would be found when real, three-dimensional objects are simultaneously presented.

We addressed this question by again measuring the number of hidden objects infants could remember, but with a simultaneous rather than sequential presentation. In our manual search paradigm, infants searched for objects they had seen an experimenter hide in an opaque box (Feigenson & Carey, 2003, 2005). A group of 12- and 14-month-old infants saw one to four identical balls simultaneously visible atop the box; the balls were then picked up and inserted through a cloth-covered opening in the box's front face. Afterward, infants were allowed to reach in and retrieve the balls. Unbeknownst to the infants, on some trials the experimenter surreptitiously removed a subset of the balls from a concealed opening in the back of the box. We measured infants' continued searching and compared it to their baseline level of searching on trials when the box was expected to be empty. Any increased searching suggests that infants successfully represented and were searching for the remaining object or objects inside the box. In this way, our manual search task serves as a measure of the number of occluded items infants can remember over a relatively short duration.

We probed the limit on the number of objects infants could simultaneously remember via a series of *x* versus *y* comparisons. For any *x* versus *y* comparison, infants' searching after they saw *x* balls hidden and had retrieved *x* of them was contrasted with their searching after they saw *y* balls hidden and had retrieved only *x* of them. The logic can be illustrated with a one-versus-two comparison. On one-object trials, infants saw the experimenter hide a single ball in the box, were then allowed to retrieve it, and any subsequent searching into the now-empty box was recorded during the 10 second measurement period that followed (Figure 3.2a). This was compared to the duration of searching on two-object trials, on which infants saw two identical balls hidden and then were allowed to retrieve just one of them. While the experimenter surreptitiously held the remaining ball out of reach for 10 seconds, any searching for the "missing" ball was recorded. After 10 seconds, the experimenter retrieved the remaining ball and showed it to infants, after which the box was once again empty. Any further searching was recorded in a final 10 second measurement period (Figure 3.2b).

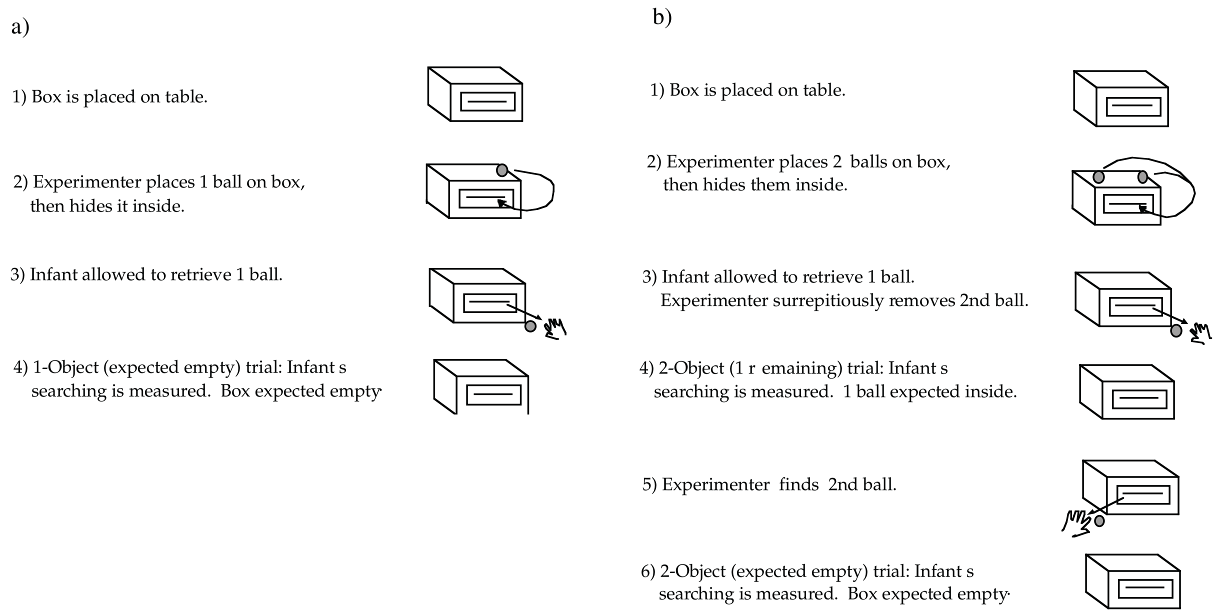

FIGURE 3.2. Presentation sequences illustrating: (a) a one-object trial; (b) a two-object trial. Reprinted from Feigenson, L., & Carey, S. (2003). Tracking individuals via object-files: Evidence from infants' manual search. *Developmental Science, 6,* 568–584, copyright (2003), with permission from Blackwell; and from Feigenson, L., & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition, 97,* 295–313, copyright 2005, with permission from Elsevier.

If infants were able to remember the correct number of objects hidden in the box, they should search the box only when the box was expected to still contain one or more objects. Therefore, we assessed infants' performance by examining the difference in their searching on trials when the box was expected to contain more objects versus trials when the box was expected to be empty. For example, subtracting search time after infants had seen one object hidden and had retrieved it from search time after infants had seen two objects hidden and had retrieved just one of them creates a difference score. If infants represent two as more than one, this difference score should be positive. We found that when 12- and 14-month-old infants were presented with this task, they succeeded (i.e., had positive difference scores) with one-versus-two and two-versus-three comparisons[3], but failed with two-versus-four and one-versus-four comparisons (Feigenson & Carey, 2003, 2005). When infants fail, we observe difference scores that are not different from chance (Figure 3.3).

Taken together, the results from the experiments just reviewed using change detection (Ross-Sheehy et al., 2003), cracker choice (Feigenson & Carey, 2005; Feigenson, Carey & Hauser, 2002), and manual search (Feigenson & Carey, 2003, 2005) have yielded identical patterns of results concerning infants' capacity limits. Whether infants saw food or nonfood objects, two-dimensional or three-dimensional objects, sequential or simultaneous presentation, or were asked to approach the larger of two total quantities, to search for hidden objects, or to

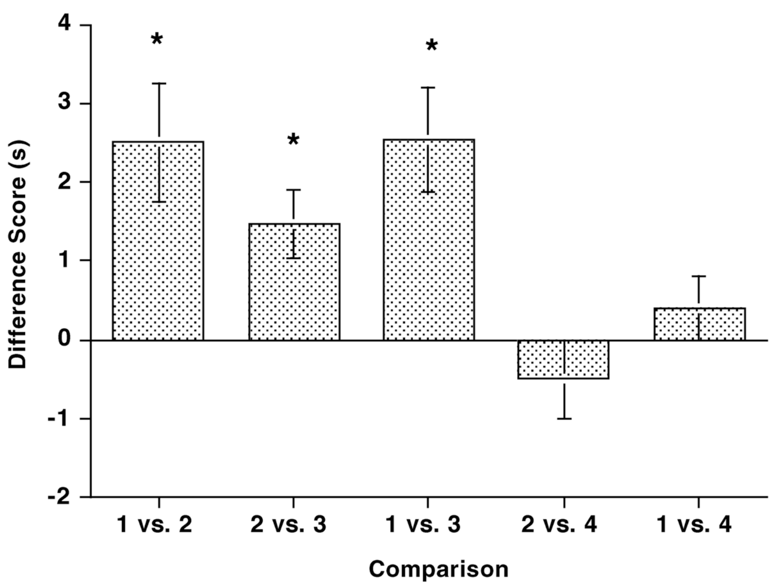

FIGURE 3.3. Difference scores (searching when the box contained more objects minus searching when the box was empty) reflect 12- to 14-month-old infants' capacity to represent and discriminate arrays containing different numbers of objects (Feigenson & Carey, 2003, 2005).

notice changing features in an object array, infants were limited to representing approximately three to four objects at a time in memory.

An important note concerning these observations of short-term memory capacity is that the three-to-four-item limit emerges in situations that require the tracking of objects qua individuals. In contrast, other tasks require infants to represent an array of items as a group whose members are not stored individually, as with a set of 16 dots scattered randomly on a screen. In such tasks, infants have been shown to represent the approximate numerosity of the whole array (Lipton & Spelke, 2003; McCrink & Wynn, 2004; Xu, 2003; Xu & Spelke, 2000; Xu, Spelke, & Goddard, 2005), but likely are not representing the individual dots that comprise the array. It appears that the ability to represent the approximate numerosity of a set containing many items and the ability to represent a small number of discrete individuals are subserved by different mental processes.

Several aspects of infants' performance with small versus large arrays support the view that infants represent them in fundamentally different ways. First, with large numerosities, such as Xu and Spelke's arrays of 16 or 24 dots, infants' success or failure depends on the ratio between to-be-discriminated arrays, rather than on the absolute number of items presented. For example, 6-month-old infants discriminate arrays of 16 from arrays of 32, but not from arrays of 24. We have already seen that for small arrays the reverse is true: It is the absolute number of items that determines performance (Feigenson & Carey, 2003, 2005; Feigenson, Carey, & Hauser, 2002). Second, infants often appear unable to represent the approximate numerosity of arrays containing four or fewer objects when the arrays are controlled for continuous properties that frequently correlate with number (Feigenson, Carey, & Hauser, 2002; Xu, 2003; Xu et al., 2005). In contrast, infants can represent the approximate numerosity of arrays that are controlled for continuous properties when the arrays contain large numbers of items (Lipton & Spelke, 2003; McCrink & Wynn, 2004; Xu, 2003; Xu & Spelke, 2000; Xu et al., 2005). These two divergent patterns of results suggest that two distinct mental systems are available to infants. One of these systems enables the representation of the approximate numerosity of large arrays. The other, which is the focus of the present discussion, allows precise representations of one to four items to be held in short-term memory (for further discussion of this two-system view, see Feigenson, Dehaene, & Spelke, 2004).

Further Evidence of Continuity: Increasing Short-Term Memory Storage via Chunking

Chunking in Adults

Besides claiming that adults have a limited-capacity short-term memory, Miller (1956) also suggested that this capacity could sometimes be increased by condensing information into a more efficient format. Specifically, Miller proposed that individual items could be bound together in memory into chunks whose components were in some way related to each other. If these chunks, rather than the individual items that comprised them, occupied the limited number of available memory slots, then the chunks could later be "unpacked" into their constituent components. Hence, the overall amount of information accessible via short-term memory could effectively be increased.

That chunking can indeed increase memory capacity in this way has been shown in impressive demonstrations of memory enhancement. One particular subject, S.F., was able to increase his memory span to nearly 80 random digits (Ericcson, Chase, & Faloon, 1980). S.F. had an entirely normal memory span at the start of the experiment, but after over 200 hours of laboratory practice he had a span equal to that of professional memory experts. S.F. accomplished this dramatic improvement by developing and perfecting his own idiosyncratic chunking strategy. He associated every three or four digits presented to him with a meaningful unit of information already present in his long-term memory. For example, S.F. remembered the digits 3, 4, 9, and 2 as "3 minutes and 49 point 2 seconds," which he knew was a near-record time to run the mile. The digits 1, 9, 4, and 4 were recalled as "1944, near the end of World War II." Interestingly, even by chunking four-digit strings into discrete chunks, the three-to-four-item limit on short-term memory should have prohibited S.F. from storing any more than three-to-four chunks containing a total of 12 to 16 individual digits. S.F. surpassed this expected limit by creating hierarchical memory entries in which chunks were nested within "superchunks." This extremely efficient collapsing of information accounted for S.F.'s impressive memory abilities. S.F.'s short-term memory enhancement is not an isolated case. Indeed, the finding that adults can use semantic information to increase storage has also been obtained with naive subjects. Typical college students were able to increase their short-term memory for digits over several laboratory sessions by associating groups of digits with preexisting referents, just as S.F. did (Chase & Ericsson, 1981).

A similar chunking mechanism has been found to underlie the exceptional performance of chess experts, who show vastly better memory for the configuration of pieces on a chessboard than do nonexperts. Rather than having a greater number of memory slots in which to store the individual pieces' locations, these experts benefited from the ability to chunk multiple pieces into recognizable formations (Simon & Chase, 1973). Doing so allowed them to store the entire board in terms of only a few formations, the individual components of which could be reconstructed from long-term memory. Support for the explanation that semantic knowledge allowed the formation of chunks comes from experiments testing the memory of expert versus novice players for randomly positioned pieces, as opposed to memory for configurations that might occur in an actual chess game. When presented with random configurations, the recall performance of experts was no better than that of novices (Simon & Chase, 1973).

Thus, the evidence indicates that adults can use existing knowledge to condense individual bits of information into larger chunks. Doing so enables the storage of more total units of information (e.g., more random digits or more locations on a chessboard). This is because short-term memory only has to retain the chunks themselves (e.g., a near-record time for running the mile), since the individual components of these chunks (e.g., 3, 4, 9, 2) already exist in long-term memory. Such a computation is clearly useful, allowing greater speed and efficiency of access to information across a wide variety of situations. But what are the origins of chunking? Is this highly useful strategy a learned one, perhaps acquired during formal education? Or is it available early on, prior to explicit instruction?

Chunking in Infants

My laboratory has addressed this question in a series of studies examining chunking in 14-month-old infants. Infants of this age had previously demonstrated a three-object memory limit in the manual search task, as discussed earlier (Feigenson & Carey, 2003, 2005). The new question was whether infants in this task could be induced to chunk individual items into smaller sets and thereby increase the total number of items remembered. To test this, we presented infants with arrays of identical objects that we then hid inside a box (Feigenson & Halberda, 2004). We used the same one-versus-two and two-versus-four object comparisons as in previous experiments (Feigenson & Carey, 2003, 2005). On one-versus-two comparisons, we asked whether infants searched the box more after seeing two objects hidden and retrieving just one of them (the other was surreptitiously withheld) than they did after seeing one object hidden and retrieving one. Success would indicate the ability to remember at least two objects, and to recognize that two is more than one. On two-versus-four comparisons, we asked whether infants searched the box more after seeing four objects hidden and retrieving just two of them (the other two were surreptitiously withheld) than they did after seeing two objects hidden and retrieving two. Here, success would indicate the ability to remember up to four objects, and the recognition that four is more than two.

Earlier in this chapter, I explained that infants had previously failed at this kind of two-versus-four comparison when four objects were presented in a single line on top of the box. The new manipulation in this study was the spatial arrangement of the objects prior to hiding (see Figure 3.4). On some trials, all of the objects were presented in a single location centered on top of the box (e.g., four objects in a line atop the box). On other trials, the objects were presented on two spatially separated platforms located on either side of the box (e.g., two objects on the left-hand platform, and two on the right-hand platform). Our hypothesis was that this spatial grouping cue would help infants chunk four objects into two sets of two, thereby enabling them to successfully represent a total of four items at once.

We found that this spatial grouping changed the total number of objects infants were able to remember. Although infants succeed at distinguishing

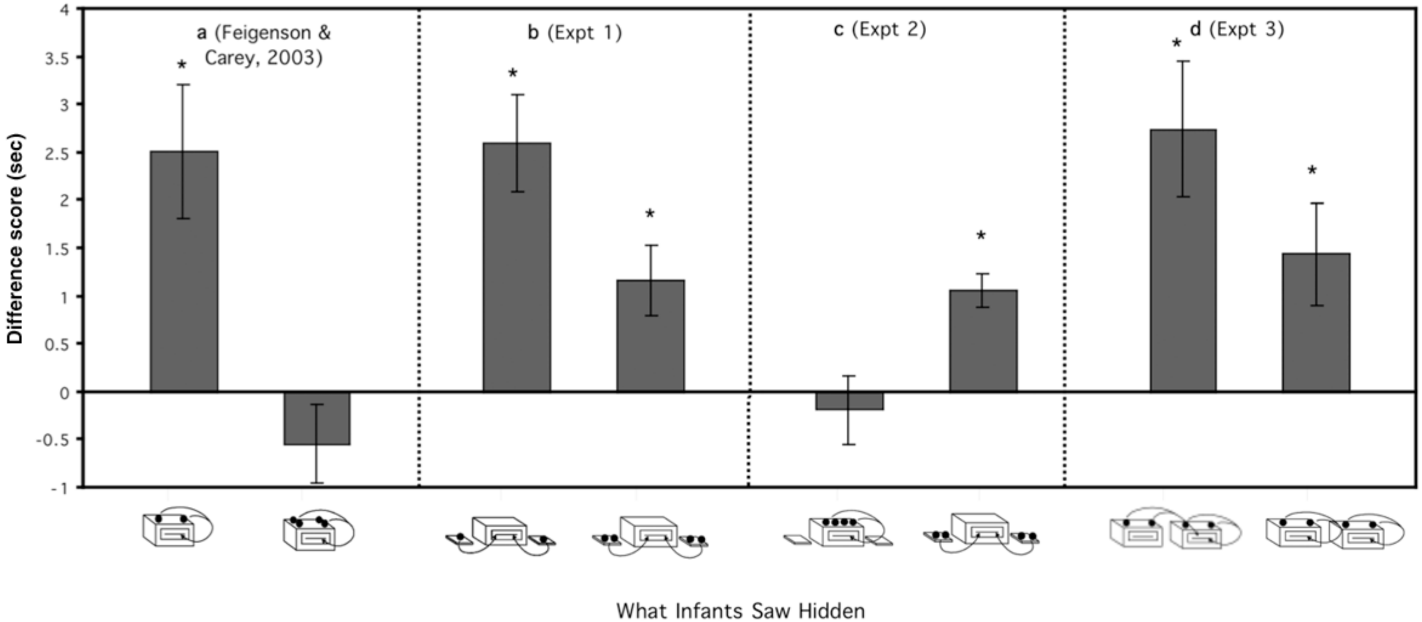

FIGURE 3.4. Difference scores (searching when the box contained more objects minus searching when the box was empty). Difference scores significantly above 0 reflect 14-month-old infants' capacity to remember the correct number of objects in the box. In (a), infants successfully remembered two objects in the box but failed to remember four. In (b), objects were spatially separated into two sets, each containing fewer than three objects. Infants again remembered two objects, and also successfully remembered four objects. In (c), infants' failure to remember four objects and their success at remembering two sets of two was replicated within-subject. In (d), infants remembered the precise location of each set when the sets were hidden in spatially separate locations. Reprinted from Feigenson, L., & Halberda, J. (2004). Infants chunk object arrays into sets of individuals. *Cognition, 91,* 173–190, copyright 2004, with permission from Elsevier.

one versus two no matter how the objects were arranged, they overcame the three-item short-term memory limit only on trials when the objects were spatially grouped. Only when objects were presented in two distinct groups of two did infants distinguish the hiding of two objects from the hiding of four (Figure 3.4). This pattern reveals two things. First, it replicates our previously reported three-item limit on infants' tracking of occluded objects (Feigenson & Carey, 2003, 2005), showing that infants were unable to store a single four-object array in memory. Second, these results also show that this limit can sometimes be overcome. By chunking representations of individual items into smaller units, infants were able to remember more total objects. The 14-month-old infants represented two chunks, each containing two individual objects. To our knowledge, this is the first demonstration of chunking in infants.

Top-Down Versus Bottom-Up Computations

Initiating Chunking

The data reviewed above suggest that the chunking operations that were classically studied by Miller and others may be both independent of formal training and available quite early in life. If so, this clearly strengthens the case for the continuity of short-term memory throughout the life span, as both the limits on short-term storage and the chunking used to overcome these limits appear to be present in infants as well as in adults. However, there is an important difference between the finding that adults can increase short-term memory capacity and the finding that infants can do so. The adult chunking studies show that adults can use semantic information to condense information. For example, adults can bind multiple individual items together based on existing knowledge (as with the race time example) or can recognize multiple individuals as forming a meaningful gestalt (as with the chess expert example). These computations rely on semantic knowledge that is available for association with the objects in the array. The computations thus can be considered top-down, in the sense that they are driven from the internal knowledge to the external, to-be-chunked items in the world.

Can infants use semantic knowledge to drive chunking? Our developmental results (Feigenson & Halberda, 2004) show that the spatial organization of an array into sets, each of which contains three or fewer items, helps infants overcome the three-to-four-item limit on short-term memory. However, unlike S.F. or the chess experts, infants relied on spatiotemporal rather than semantic information. Furthermore, the computation they performed was bottom-up in the sense that the requisite information for chunking was present in the array itself, rather than in infants' existing knowledge. The spatial arrangement that the experimenter imposed on the array led infants to parse it into smaller sets, rather than infants themselves imposing their knowledge to reorganize their representation of the array. Thus far, it is an open question whether infants can also chunk in a top-down fashion on the basis of stored semantic knowledge.

Experiments in my lab are currently exploring which sources of information infants can use to chunk an array. As in our earlier investigations, we take as evidence for chunking the ability to represent a total of four objects in our manual search task. These previous studies revealed that infants could do so only if they saw the four-object array presented as two spatially separated sets, each containing fewer than three items (Feigenson & Halberda, 2004). In a series of new studies, we replace spatial information with semantic information as the potential basis for chunking. Infants see all four objects in a single line atop the box—a spatial arrangement that has previously led infants to fail. However, we now show infants an array of two cars and two cats, instead of the four identical balls we used in our earlier studies (Feigenson & Carey, 2003, 2005; Feigenson & Halberda, 2004). Given that 14-month-old infants are reported by their parents to be familiar with these entities (and given that most 14-month-olds already know the words *car* and *cat*, or *kitty*; Fenson et al., 1994), we hypothesize that infants may be able to use this semantic knowledge to chunk the four-object array into two sets of two. To ask whether any observed success is based on semantic knowledge of the object categories, as opposed to low-level perceptual differences between the two types of objects, on other trials we present infants with two sets of two objects that are perceptually distinct, yet from unfamiliar categories. If infants fail to represent all four objects when presented with unfamiliar objects, such as two toy shrimp and two toy tanks, but succeed with two cars and two cats, then we can more confidently say that infants are able to use semantic knowledge in a top-down fashion to chunk the array.

Initiating Computations of Discrete and Continuous Quantity

The question of whether infants can initiate top-down computations over object representations is not exclusive to chunking, but also arises for other operations performed over representations being held in memory. An example comes from the work on infants' computations of quantity. Infants have been shown to be capable of computing the discrete number of individual objects in object arrays, showing different looking patterns to expected versus unexpected numbers of objects (Cheries, DeCoste, & Wynn, 2003) or searching a box until the expected number of objects has been retrieved (Feigenson & Carey, 2003). These findings obtain when the total area or volume of the arrays is controlled for. Infants also can compute the total continuous extent contained in an array, showing increased looking when the overall summed area or perimeter of the items in the array changes (Clearfield & Mix, 1999, 2001; Feigenson, Carey, & Hauser, 2002) or choosing to approach an array containing a greater total volume of food over an array containing a smaller total volume, regardless of the number of individual objects involved (Feigenson, Carey, & Spelke, 2002). That these quantity computations operate over the same short-term memory representations discussed above is shown by the conditions under which infants successfully perform them. Infants can compute discrete or continuous quantity over arrays containing small numbers of objects. But when arrays contain four or more objects, infants fail to compute either number or total extent (Feigenson & Carey, 2003, 2005; Feigenson, Carey, & Hauser, 2002; Xu, 2003; Xu et al., 2005). Thus, recognizing the number of objects in an array and recognizing the total extent contained in the array are both the output of computations performed over short-term memory representations of objects. When there are too many objects to be represented in short-term memory, infants fail to compute either number or total extent.

What prompts infants to respond to the discrete (e.g., number of individual objects) versus continuous (e.g., color, total extent) properties of a given object array? Although this question will likely be the focus of many future experiments, one recent set of findings suggests that the features of the objects themselves play a role in determining which dimension of quantity infants represent. A group of 7-month-old infants was habituated to object arrays, then tested with arrays in which either the number of objects or the total surface area had changed (Feigenson, 2005; Feigenson, Carey, & Spelke, 2002). The results revealed that when the array contained objects that were identical in color, pattern, and texture, infants dishabituated to changes in the total area of the array, but not to changes in the number of objects in the array. When the array contained objects that contrasted with each other in color, pattern, and texture, however, infants did just the reverse. They dishabituated to changes in the number of objects in the array but not to changes in total area. In other words, infants appeared able to compute either number or surface area but unable to perform both computations over the same array. Figure 3.5 depicts this double dissociation.

In these experiments, the computation that infants performed (number versus total extent) appeared to be under exogenous control, influenced in a bottom-up fashion based on whether objects in the array had identical properties or not. In contrast, adults can exert top-down control over these computations, choosing whether to represent the number or extent contained in an array even on a trial-by-trial basis (Feigenson & Halberda, in preparation). Thus, infants and adults may differ in the endogenous versus exogenous nature of the quantity computations they can perform.

What Develops?

In the preceding pages, I have tried to build the case that infants and adults share a system for maintaining object representations in short-term memory. This system is capacity limited and can only store representations of three to

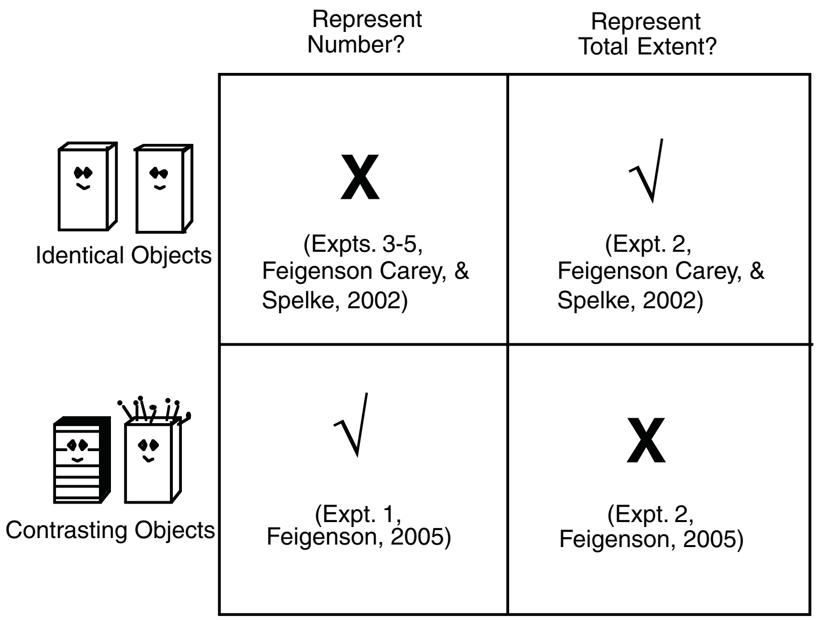

FIGURE 3.5. Double dissociation between array heterogeneity and the computation infants perform over the objects in the array. Reprinted from Feigenson, L. (2005). A double dissociation in infants' representation of object arrays. *Cognition, 95,* B37–B48, copyright 2005, with permission from Elsevier.

four individuals at any one time. Once stored in short-term memory, infants and adults can perform a range of computations over these representations. For example, we have begun to understand the ways in which infants and adults compute quantity, both discrete and continuous, over object arrays. In addition, both infants and adults appear able to reorganize representations of individuals into a hierarchical structure. This chunking of object representations provides a link between limited-capacity short-term memory and the greater storage capacity of long-term memory, and has been shown to enable both infants and adults to overcome the three-to-four-item limit on simultaneous representation. In all of these ways, the representations stored by and the computations performed by infants and adults are strikingly similar.

Where, then, is the development in short-term memory? While I have focused on the respects in which short-term memory may be continuous over development, there may also be important ways in which early versus mature short-term memory differs. I now point to some existing research, some of which is addressed by other chapters in this volume, as well as to avenues for future investigation of short-term memory development.

First, we have already explored some possible differences in short-term memory computations in terms of bottom-up versus top-down processing.

Adults can use semantic knowledge to endogenously initiate the chunking of representations held in short-term memory, and have volition over which dimension of quantity to represent. It remains unclear whether infants also have this ability. To date, infants' computations over object arrays appear to be driven from the bottom up by perceptual information present in the array itself. Even if new research finds that infants can initiate the chunking of an object array in a top-down fashion (for example, by using conceptual knowledge of animals versus vehicles to parse an array into these two categories), the comparative richness of adult knowledge about the world will likely be reflected in developmental differences. If infants do have some top-down control over the chunking of object arrays, this control will almost certainly increase over time. As they come to refine their knowledge of object kinds and categories, infants may gain more ways to parse arrays into chunks and therefore gain more avenues for motivating top-down chunking. Exploring developmental changes in the top-down versus bottom-up execution of short-term memory computations is a promising direction for future research.

In addition, much remains to be understood about the nature of infants' and adults' capacity limits. Although both groups appear able to store up to three to four items in short-term memory, information capacity within each of these three to four available slots is probably not fixed. For example, representing three very complex objects with many features and articulated parts may impose a higher informational load than representing three simple geometric shapes. Alvarez and Cavanaugh (2004) measured this load empirically using a change detection task, and confirmed that the number of objects adults can store depends on the objects' complexity. Adults were able to maintain a larger number of items in visual short-term memory when those items were simple colored squares than when the items were more complex letters or shapes. These results show that although short-term memory can store a maximum number of about four items, the information load of the array can significantly reduce this capacity. This question has yet to be systematically explored in infants, and it raises the possibility that infants and adults may differ in the amount of information they can store in each memory slot. Can infants represent multiple features of three to four complex objects, or are they limited to representing just a few salient properties? Systematically manipulating object complexity will help characterize the subtle limits of short-term memory development.

Finally, while the eventual upper limit on short-term memory appears fixed at three to four items, infants' memory capacity may not reach this limit for some time. In their change detection experiments, Ross-Sheehy et al. (2003) found that 10- and 13-month-old infants detected a changing item contained within a three- or four-item array. But 4- and 6.5-month-old infants could only detect a change with a one-item array, failing with arrays of two and three. A similar pattern has been obtained by Kaldy and Leslie (in press), using a quite different paradigm in which infants are asked to track the shape of three-dimensional objects that move behind occluding screens.

Based on their results, Kaldy and Leslie suggest that 6.5-month-old infants' short-term memory capacity is limited to just one slot. These findings raise the possibility that the capacity of short-term memory increases over the first year of life and reaches asymptote by 10 to 12 months.

This interesting developmental proposal may conflict with previous findings that even by 5 months, infants can successfully track and remember at least two hidden objects at a time (Koechlin et al., 1997; Simon et al., 1995; Wynn, 1992). Ross-Sheehy et al. (2003) suggested that these previous results might not be tapping short-term memory and that, because infants view objects over much longer durations than in the change detection task (several seconds, compared with 500 milliseconds), long-term memory might also be involved. This issue merits deeper exploration. The three-to-four-item limit is observed in tasks involving a wide range of durations, from 500 milliseconds to 30 seconds or more. Does this commonality implicate a single system of memory representation encompassing a wide span of durations? Might there be multiple levels of memory storage that are all limited by a single bottleneck on information processing?

These questions return us to the issues raised by Cowan's (2001) controversial proposal regarding the distinctions between attention, short-term memory, and long-term memory. The developmental evidence reviewed here does not provide definitive answers. Nonetheless, characterizing infants' memory—both its capacities and its limitations—may offer a window into understanding these systems and the interactions between them. As we identify the continuities and discontinuities in representational ability across the life span, we add to the emerging portrait of memory development over time.

Notes

1. Previous work (Feigenson, Carey, & Hauser, 2002; Scholl & Leslie, 1999; Simon, 1997; Uller et al., 1999) has suggested that infants and adults share a system that is dedicated to tracking objects per se, and to storing information about their properties. This "object-file" system enables the creation of a token, or file, that stands for an object in the world and allows it to be represented over changes in spatial location or changes in properties (Kahneman, Treisman, & Gibbs, 1992). While such a system may indeed be in place throughout development, here I make the more general claim that infants and adults share a system for representing discrete items in short-term memory. These items may be objects, but may also be nonobject entities such as sounds or events that are perceived in any sensory modality.

2. Presenting the cue less than 1 second after the array disappeared resulted in much higher capacity limits, which Sperling took as evidence for the existence of a very short-lived, iconic memory store. Iconic memory, which lasts for less than a second, is distinct from the short-term memory that is the focus of the present discussion.

3. Note that this method does not allow us to be certain of the exact number of objects infants represented in the box. For example, for infants who saw three objects hidden, retrieved one of them, and then continued to search, infants may have believed there to be exactly two objects still remaining. Alternatively, it is possible that they represented just one more object in the box or an unspecified number of objects remaining in the box.

References

Alvarez, G. A., & Cavanaugh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15,* 106–111.

Alvarez, G. A., & Franconeri, S. L. (2005). *How many objects can you track?* Paper presented at the annual meeting of the Vision Sciences Society, Sarasota, FL.

Baillargeon, R. (1986). Representing the existence and location of hidden objects: Object permanence in 6- and 8-month old infants. *Cognition, 23,* 21–41.

Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition, 20,* 191–208.

Carey, S. (2004). Bootstrapping and the origins of concepts. *Daedalus, 133,* 59–68.

Chase, W. G., & Ericsson, K. A. (1981). Skilled memory. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 141–189). Mahwah, NJ: Erlbaum.

Cheries, E. W., DeCoste, C., & Wynn, K. (2003). *Number Not Area: Infants Use Property Contrasts for Quantifying Objects.* Poster presented at the Society for Research in Child Development, Tampa, FL.

Clearfield, M. W., & Mix, K. S. (1999). Number versus contour length in infants' discrimination of small visual sets. *Psychological Science, 10*(5), 408–411.

Clearfield, M. W., & Mix, K. S. (2001). Infants' use of area and contour length to discriminate small sets. *Journal of Cognition and Development, 2,* 243–260.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–185.

Diamond, A. (1990). Rate of maturation of the hippocampus and the developmental progression of children's performance on the delayed non-matching to sample and visual paired comparison tasks. *Annals of the New York Academy of Sciences, 608,* 637–676.

Ericsson, K. A., Chase, W. G., & Faloon, S. (1980). Acquisition of a memory skill. *Science, 208,* 1181–1182.

Feigenson, L. (2005). A double dissociation in infants' representation of object arrays. *Cognition, 95,* B37–B48.

Feigenson, L., & Carey, S. (2003). Tracking individuals via object-files: Evidence from infants' manual search. *Developmental Science, 6,* 568–584.

Feigenson, L., & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition, 97,* 295–313.

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.

Feigenson, L., Carey, S., & Spelke, E. (2002). Infants' discrimination of number vs. continuous extent. *Cognitive Psychology, 44*(1), 33–66.

Feigenson, L., Dehaene, S., & Spelke, E. S. (2004). Core systems of number. *Trends in Cognitive Sciences 8*(7), 307–314.

Feigenson, L., & Halberda, J. (2004). Infants chunk object arrays into sets of individuals. *Cognition, 91,* 173–190.

Feigenson, L., & Halberda, J. (in preparation). *Computations made over set representations.*

Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D., Pethick, S. J., et al. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development, 242.*

Halberda, J., Simons, D., & Wetherhold, J. (2006). *You can never attend to more than three items at once: Gestalt grouping principles explain changes in capacity*. Manuscript submitted for publication.

Halberda, J., Sires, S. F., & Feigenson, L. (2006). Multiple spatially overlapped sets can be enumerated in parallel. *Psychological Science, 17*(7), 572–576.

Hood, B., & Willats, P. (1986). Reaching in the dark to an object's remembered position: Evidence for object permanence in 5-month-old infants. *British Journal of Developmental Psychology, 4,* 57–65.

Kahneman, D., Treisman, A., & Gibbs, B. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology, 24*(2), 175–219.

Kaldy, Z., & Leslie, A. M. (2003). Individuation of objects in 9-month-old infants: Integrating "what" and "where" information. *Developmental Science, 6*(3), 360–373.

Kaldy, Z., & Leslie, A. M. (2005). A memory span of one? Object identification in 6.5-month old infants. *Cognition, 97*(2), 153–157.

Koechlin, E., Dehaene, S., & Mehler, J. (1997). Numerical transformations in five month old human infants. *Mathematical Cognition, 3,* 89–104.

James, W. (1890). *The principles of psychology.* New York: Henry Holt.

Lipton, J. S., & Spelke, E. S. (2003). Origins of number sense: Large number discrimination in human infants. *Psychological Science, 15,* 396–401.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281.

McCrink, K., & Wynn, K. (2004). Large-number addition and subtraction in infants. *Psychological Science, 15,* 776–781.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.

Noles, N., Scholl, B. J., & Mitroff, S. R. (2005). The persistence of object file representations. *Perception and Psychophysics, 67*(2), 324–334.

Pylyshyn, Z. W. (1994). Some primitive mechanisms of spatial attention. *Cognition, 50,* 363–384.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition, 80,* 127–158.

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision, 3*(3), 179–197.

Ross-Sheehy, S., Oakes, L. M., & Luck, S. J. (2003). The development of visual short-term memory capacity in infants. *Child Development, 74,* 1807–1822.

Scholl, B. J., & Leslie, A. M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In E. Lepore & Z. Pylyshyn (Eds.), *What is cognitive science?* (pp. 26–73). Oxford, UK: Blackwell.

Shinskey, J. L., & Munakata, Y. (2003). Are infants in the dark about hidden objects? *Developmental Science, 6,* 273–282.

Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist, 61,* 393–403.

Simon, T., Hespos, S. J., & Rochat, P. (1995). Do infants understand simple arithmetic? A replication of Wynn (1992). *Cognitive Development, 10,* 253–269.

Simon, T. J. (1997). Reconceptualizing the origins of number knowledge: A "non-numerical" account. *Cognitive Development, 12*, 349–372.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74.*

Trick, L., & Pylyshyn, Z. (1994). Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision. *Psychological Review, 101*(1), 80–102.

Uller, C., Huntley-Fenner, G., Carey, S., & Klatt, L. (1999). What representations might underlie infant numerical knowledge? *Cognitive Development, 14,* 1–36.

Wynn, K. (1992). Addition and subtraction by human infants. *Nature, 358,* 749–750.

Xu, F. (2003). Numerosity discrimination in infants: Evidence for two systems of representation. *Cognition, 89,* B15–B25.

Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month old infants. *Cognition, 74,* B1–B11.

Xu, F., Spelke, E. S., & Goddard, S. (2005). Number sense in human infants. *Developmental Science, 8,* 88–101.