

**Learning words from pictures: 15- and 17-month-old infants appreciate the referential and symbolic links among words, pictures, and objects**

**1. Introduction**

As developmental psychologists, we extol the benefits of reading books with infants, focusing not only on the benefits to the infant-caregiver bond, but also on the advantages that book reading confers on infant language development, and word-learning in particular. Although research reveals that infants can, indeed, map novel words to two-dimensional representations (presented in picture books or on screens), remarkably little is known about whether and how infants link those words to the real, 3-dimensional objects that they may later encounter as part of their daily experience. This is an important question, because despite the consensus that infants are resoundingly successful at word-learning, the representational status of infants' early words remains a topic of spirited debate (Booth & Waxman, 2003, 2006; Ganea, Allen, Butler, Carey, & DeLoache, 2009; Gelman & Waxman, 2009; Graham, Booth & Waxman, 2012; Plunkett, 1997; Sloutsky, 2009; Sloutsky, Lo & Fisher 2001; Smith, Jones, Yoshida & Colunga 2003; Smith & Samuelson, 2006; Waxman & Gelman, 2009). This debate engages a classic tension, ubiquitous throughout the psychological and developmental sciences, surrounding the fundamental content of human knowledge and the kinds of processes that underlie its acquisition. Within the arena of early word learning, at issue is the content of infants' word meanings and the process by which they are acquired. The current

## LEARNING WORDS FROM PICTURES

experiment was designed to shed new light on this topic, focusing specifically on the referential status of novel words that infants learn within the context of a picture book (e.g., “whisk”, applied exclusively to pictures). We ask whether, at 15 to 17 months, infants’ interpretation of a newly-learned word is sufficiently abstract to permit them to extend it to new exemplars from the same object category (e.g., other whisks) that differ in color and in representational medium (real, three-dimensional whisks).

### **1.1 Words as abstract and referential entities: overview of relevant literature**

We begin with a brief overview of the two alternative frameworks for word-learning. Within an associationist framework, early word meanings are built on the bedrock of exclusively sensory and perceptual content, and this content is bound together by the process of association (Rakison & Lupyan, 2008; Sloutsky, 2009; Sloutsky et al., 2001; Smith et al., 2003; Smith & Samuelson, 2006). Proponents of this view argue that the very nature of word meaning undergoes a frank developmental shift: Initially, infants’ word meanings incorporate perceptual content alone; only later do they come to incorporate conceptual content into word meaning as well. Associationist accounts differ considerably in where they locate this developmental shift – in some accounts as early as 2 years (Samuelson & Smith, 1999) and in others as late as 8 – 11 years (Sloutsky et al., 2001, p. 1707) – but these accounts all share the same underlying assumption that for infants and young children, words are “...features of objects that contribute to their overall similarity, rather than symbols denoting category membership” (Sloutsky et al., 2007, p. 180). Put differently, the claim is that early in development, a word is nothing

## LEARNING WORDS FROM PICTURES

more than a feature of the perceptual and sensory experience(s) with which it has co-occurred, just as a black beret is a feature of the experience we associate with Jean Piaget.

We have argued for a different view, one that is based on fundamentally different assumptions about words, concepts, and development (see Waxman and Gelman, 2009 for extended discussion). On this view, early word meaning incorporates conceptual as well as sensory and perceptual content from the start, and in the *process* of establishing the meaning of a word, infants interpret this content within the context of their expectations about words, concepts, and reference (Baillargeon, 2008; Baldwin, 1995; Booth, Waxman, & Huang, 2005; Carey, 2009; Gelman & Waxman, 2009; Gelman & Williams, 1998; Putnam, 1973; Spelke, 2000; Xu, Cote, & Baker, 2005). In essence, then, development is characterized by continuities rather than frank developmental shifts. Considerable evidence is consistent with this continuity view: From infancy, words (and the concepts to which they refer) are more than collections of sensory and/or perceptual features, bound together by purely associative processes. Moreover, by roughly their first birthdays, infants extend novel nouns in a principled fashion, extending them specifically to object categories and not to other similarity-based groupings (Booth et al., 2005; Waxman, 1999), and using them to support inferences about new objects (Graham, Kilbreath, & Welder, 2004; Graham et al., 2012; Keates & Graham, 2008; Xu et al., 2005). Results like these indicate that infants appreciate that words are symbols whose referential scope extends beyond the particular objects with which it has been paired.

This distinction between associative processes and reference was illuminated in an experimental paradigm introduced by Preissler and Carey (2004). An experimenter introduced 18- and 24-month-old infants to a picture of a novel entity (a whisk) and

## LEARNING WORDS FROM PICTURES

named it (“a whisk”). She then presented infants with (a) another picture of a whisk and (b) an actual, three-dimensional whisk, asking, “Can you show me a whisk?” Preissler and Carey reasoned as follows: On an associative account, infants should select the picture because it is perceptually more similar to the picture with which the word “whisk” had been previously introduced. In contrast, on a referential account, infants should select the real three-dimensional object as readily as the picture, because both are instantiations of the same concept. The results were striking: infants not only selected the real object, but favored it over the picture. This reveals that 24-month-old infants understood something both subtle and profound: words refer to concepts, and are not tethered tightly to the perceptual impressions with which they have been previously paired. In sharp contrast, when children with Autistic Spectrum Disorder (ASD) participated in the very same paradigm, they almost exclusively selected the picture, and not the three-dimensional object (Preissler, 2008). This response pattern, which aligns well with the predictions of the associative account, is consistent with evidence that word learning is impaired in children with ASD, and that when they do acquire words, their meanings tend to center around sensory and perceptual associations (Baldwin, 1995; Baron-Cohen, Baldwin, & Crowson, 1997; Frith & Happé, 1994; Preissler, 2008).

More recently, Ganea and her colleagues (2009) built upon these findings to explore the developmental roots of the referential status of words in younger, typically-developing infants. They taught infants younger than 24 months a novel word for a *picture* of a novel object, and asked whether infants would then extend that word to *real three-dimensional* objects. At issue was whether these younger infants could appreciate the referential status of newly-learned words (revealing developmental continuity) or

## LEARNING WORDS FROM PICTURES

would respond in an associative fashion (revealing a developmental shift from an initially associative to a later referential pattern). To address this question, they modified the original paradigm in several clever ways to accommodate infants as young as 15 and 18 months. For example, they introduced the novel noun in conjunction with realistic colored photographs (rather than line-drawings) within the context of a naturalistic book-reading session. Second, they introduced a training phase to support the younger infants' abilities to respond systematically to the kinds of test questions that would follow. This training phase also permitted the researchers to identify infants who were unable to establish the word-picture pairing during the book-reading session. After completing this training, infants responded to a series of test questions, patterned after those of Preissler and Carey (2004).

This experiment revealed intriguing developmental effects. At 24 months, although many infants selected both the real object and the picture (replicating Preissler & Carey, 2004), they did not *favor* the object over the picture (in contrast to Preissler & Carey, 2004). At 18 and 15 months, infants' responses were more equivocal: infants divided their responses evenly among the three possible response patterns: roughly one-third chose both the real object and the picture, one-third chose the object alone, and one-third chose the picture alone.

To gather more compelling evidence from infants at these younger ages, Ganea et al. (2009) conducted a second experiment, this time changing the color of the real object (but not the picture) presented at test (Figure 1). The goal of this modification was to increase the *perceptual* distance between the pictures with which the novel noun was paired during book-reading (two pictures of a white whisk) and the real object presented

## LEARNING WORDS FROM PICTURES

at test (a blue three-dimensional whisk). The logic was straightforward. If infants' interpretation of a newly-learned word rests primarily on its associative link to the exemplars with which it was previously paired, then increasing the perceptual distance between these named exemplars and the test object should have a strong effect: infants should rarely (if ever) select the three-dimensional test object. But if infants appreciate that words refer to concepts, then they should readily accept the test object as a referent of the newly-learned word, because despite its difference in color, it is nonetheless a member of the same previously-named object category.

Once again, there were intriguing developmental effects. Most 15-month-olds were unable to master the task; as a result, their data could not be analyzed further. Although 18-month-olds fared better, they tended to select the test picture alone, and not the real three-dimensional test object. At 24 months, performance was equivocal: infants favored neither the picture nor the object, and rarely selected both.









Study	Training		Test	
<i>Ganea et al., (2009) Study 2</i>				
<i>Current Study</i>				

Figure 1. Target pictures and objects presented in Ganea et al., (2009) and in the current experiment. Pictures are represented black frames; three-dimensional objects are represented with no frames.

What is the best interpretation of this pattern of results? When infants hear a novel word applied to a picture of an object, are they unable to extend it to real 3-dimensional exemplars of the same object category? We suspect that, on the contrary, infants can

## LEARNING WORDS FROM PICTURES

indeed extend a word from a picture-book to real three-dimensional objects, but that something in this paradigm inadvertently led infants to interpret the novel word narrowly. (See Figure 1). What might have engendered this effect?

We know that infants are sensitive to the range of objects to which a novel word is applied (Waxman, 1990). We suspect that when the experimenter repeatedly paired the novel noun with pictures of a *white* whisk, infants may have adopted a narrow interpretation: in the absence of evidence to the contrary, they may have expected “whisk” to apply to white whisks in particular. After all, many nouns refer to object categories whose members tend to be uniform in color (e.g., “lemon”, “raven”). Notice, too, that in the test phase, infants could maintain this narrow interpretation by selecting the (white) picture, rather than extending the word to the (blue) three-dimensional object. Of course, adopting a narrow interpretation of a novel noun in these circumstances does not constitute evidence that infants failed to grasp its referential status. Rather, Ganea et al.’s infants may simply have opted for a restricted range of extension, in the absence of any information to the contrary.

### 1.2. The current experiment

In the current experiment, our goal was to delve deeper into infants’ representational capacities for newly-learned words. Using Ganea et al.’s (2009) paradigm as a foundation, we introduced three design modifications. First, and most important from a theoretical perspective, we introduced the novel noun in conjunction with pictures of two whisks that differed in color (one purple, one orange). At test, we asked whether infants would extend the noun to an exemplar of the same category that

## LEARNING WORDS FROM PICTURES

differed in color (a picture of a silver whisk; a real three-dimensional silver whisk). (See Figure 1.) If infants' interpretation of a novel noun is influenced by the variability of exemplars to which it has been applied, then infants in the current experiment should adopt a broader interpretation – one that includes exemplars that vary in color. The key question is whether they will extend the novel noun to exemplars that differ not only in color, but also in representational medium (from pictures to objects).

Our second modification was designed to minimize any *a priori* preferences for selecting objects over pictures at test. Ganea et al. (2009) found that infants, especially at 15 months, exhibited a strong 'object bias' (a tendency to handle the three-dimensional objects rather than pictures, independent of the experimenter's request). The authors pointed out that if an infant reveals this *a priori* bias, it is impossible to interpret their performance at test. After all, if these 'object biased' infants selected the real object over its picture at test, this might reflect a response bias for objects, rather than their interpretation of the novel word, *per se*. To reduce such a bias, before the experiment proper began, we permitted infants to play with the test pictures and objects until their interest waned.

Third, we introduced the training phase *before* children looked through the picture book with the experimenter. This modification offered two advantages. First, it permitted us to use the training phase to identify any infants with a persistent 'object bias' (a bias to select objects over pictures) and replace them. Second, it permitted infants to move smoothly from training to test, without an intervening period that might disrupt their learning of the new word (Horst & Samuelson, 2008).



## LEARNING WORDS FROM PICTURES

With these modifications in place, the predictions of the two alternative theoretical accounts can be teased apart. If infants' interpretation of a novel word's meaning is guided only by its association with the perceptual features of the individuals with which it has co-occurred, then infants should be more likely to extend "whisk" at test to the picture of a silver whisk (than to the real silver whisk) because the picture shares more perceptual features with the exemplars with which the word had been introduced. But if infants appreciate the referential status of the novel word, they should extend "whisk" at test to both the picture and to the real object because, although these differ from the picture-book exemplars in color (the picture of silver whisk) and in representational medium (the three-dimensional silver whisk), they are both members of the same object kind to which "whisk" refers.

## 2. Method

### 2.1 Participants

Twenty-four English-acquiring primarily Caucasian infants were recruited, via mailings, from middle-class families in the Chicago area. These included twelve 15-month-olds (9 females;  $M = 15;21$ , range = 14;16-16;15; mean productive vocabulary (MacArthur Communicative Development Inventory or MCDI) = 41 and twelve 17-month-olds (6 females;  $M = 17;10$ , range = 16;16-18;15; mean productive vocabulary (MCDI) = 39). Additional infants (12 15-month-olds; 11 17-month-olds) were excluded for failure to respond correctly on one or more training questions (16). An additional 7 infants were excluded for fussiness (6), or parental interference (1). Excluded infants did not differ from those included in the final sample either by age or the MCDI measure.<sup>1</sup>

## LEARNING WORDS FROM PICTURES

### 2.2 Materials

The objects were eight familiar toys (dog, keys, baby-bottle, cup, car, phone, ball and apple) and two novel objects (silver whisk, silver garlic press). The pictures were high-quality color photographic images printed on cardstock and laminated (13cm x 10cm). They depicted the eight familiar objects (above) and six novel objects (three whisks: one orange, one purple, and one silver; three garlic presses: one green, one pink, and one silver).

For the book-reading phase, we created a book in which two pictures were displayed side-by-side on facing pages. The book contained 10 two-picture displays. Pictures of each of the familiar and novel objects (described above) were included in the book, with two exceptions: the picture of the silver whisk and the silver garlic press were reserved for the test phase. We constructed the book as follows: On the first eight displays, a picture of one novel and one familiar object appeared, counterbalanced for side. The picture of each of the four novel objects (orange whisk, purple whisk, green garlic press, pink garlic press) appeared twice, once on each side. On the final two displays, only novel objects were depicted (i.e., orange whisk paired with pink garlic-press, green garlic-press paired with purple whisk). In total, then, each of the four novel pictures (orange whisk, purple whisk, green garlic press, pink garlic press) appeared three times in the book, yielding six depictions of whisks (three orange and three purple) and six depictions of garlic presses (three green and three pink).

## LEARNING WORDS FROM PICTURES

### 2.3 Procedure

Infants and caregivers were welcomed into a laboratory playroom. While the infant played, the caregiver signed a consent form and completed the MacArthur Communicative Development Inventory (MCDI) (Fenson et al., 1993). They were then escorted into an adjoining test room where infants were seated at a small table on their caregiver's lap, across from the experimenter. Caregivers were instructed not to talk to or influence infants' attention in any way. The procedure included four distinct phases: familiarization, training, book-reading, and then test.

#### 2.3.1 Familiarization Phase

First, the experimenter presented infants with the two real objects that would be presented later at test (silver whisk, silver garlic-press) along with two familiar objects (selected randomly from those objects that would be used later in training). Infants played with these freely until their interest waned (roughly 3 min.). Our goal was to minimize the possibility that infants might favor objects over pictures at test simply because they had not yet had an opportunity to handle the novel test objects.

#### 2.3.2 Training Phase<sup>2</sup>

This phase was designed with three goals in mind: (1) to clarify that infants could offer *either* pictures or objects when the experimenter asked for them by name, (2) to support infants' selection of pictures on trials in which an object and picture were pitted against one another (in *picture-object training*, below, the experimenter consistently

## LEARNING WORDS FROM PICTURES

requested the picture), and (3) to identify infants who were unable to select a familiar picture when asked for it by name. Infants who failed to respond correctly on all training questions were replaced (see Participants, above). As a result, we can be confident that all infants included in the experiment-proper understood the task and were able to select either pictures or objects as referents of familiar nouns.

*Single-object training.* To begin, the experimenter placed a single familiar object (e.g., toy dog) in front of the infant, touched the object, and named it (e.g., “This is a dog”). She then placed her hand on the table and pointed at her open palm, saying, e.g., “Can you give me the dog?” If the infant selected the object, the experimenter provided positive reinforcement (e.g., “Yay, that is a dog!”). On this, and on all subsequent training and test trials, if the infant failed to make a selection, the experimenter repeated the question twice, paraphrasing it slightly (e.g., “Can you show me / give me the dog?”) before moving on.

*Object-object training.* The experimenter presented two familiar objects (e.g., apple, car), saying, “Look! Here is an [apple] and here is a [car].” With her palm outstretched, she asked, “Can you give me the [apple]?”

*Picture-picture training.* This was identical to the object-object training (above), except that pictures of familiar objects (rather than the objects themselves) were presented.

*Picture-object training.* This was identical to the prior training trials (above), except that a familiar object was pitted against a familiar picture (counterbalanced for side). The experimenter introduced the picture (e.g., “Here is a [ball]”) and the object (e.g., “Here is a [cup]”) (counterbalanced for order). She then requested the picture (e.g.,

## LEARNING WORDS FROM PICTURES

“Can you give me the [ball]?”). This permitted us to assess whether an infant could select the named familiar item, even if it meant selecting a picture rather than an object.

*Picture–Object (novel) training.* This was identical to the prior trial, except that an unfamiliar object (one that had not been presented in familiarization) was pitted against a familiar picture. The experimenter introduced both items, saying, “Here is a [car] (pointing to the picture), and look at this (pointing to the rolling pin).” She then requested the picture, saying “Can you give me the [car]?” This permitted us to assess whether an infant could select the named item, even if it meant selecting a familiar picture, rather than a novel (and presumably attractive) three-dimensional toy object.

### 2.3.3 Book-Reading Phase

The experimenter read through the picture book with the infant. Whenever a picture of a familiar object appeared, the experimenter pointed to it and named it twice (e.g., “Look, this is a cup. Do you see the cup?”). When a picture of the novel target kind (whisk) appeared, she did the same, saying, “Look, this is a whisk. Do you see the whisk?” When a picture of the novel distractor kind (garlic press) appeared, she pointed and drew attention to it, but provided no name (“Ooh, look at this! Wow, do you see that?”). Because the book included six pictures of whisks (three orange, three purple) and because with each appearance the experimenter named it twice, this yielded a total of 12 naming episodes in which the novel noun was paired with a picture of a whisk.

## LEARNING WORDS FROM PICTURES

### 2.3.4 Test Phase

In this phase, the experimenter presented infants with three test questions, each involving a pair of items. All of the test items were silver in color and none had appeared earlier in either the picture book or training phases. The test questions, inspired by previous designs (Ganea et al., 2009; Preissler & Carey, 2004), were presented in counterbalanced order.

*Target Picture-Target Object test.* The experimenter presented a real silver whisk and a picture of a silver whisk, saying, “There is a whisk here. Can you show me (or give me) a whisk?” Both the picture and object were members of the same novel object category that had been named during the book-reading phase, but both differed in color from the pictures with which the novel noun had been introduced. We noted which choices infants made, as well as the order in which they made them.

At issue was whether infants would extend the novel word to a new member of the same target kind that differed in color from those depicted during book reading (e.g., to a silver whisk), and whether they would extend it to a picture, an object, or both. Notice that if infants are unable to generalize the novel word to either the new color or the new medium, then they should offer no choice on this test question or should select randomly (at chance levels) between the two silver choices. But if, as we predict, infants appreciate that the novel noun refers to the object kind, then they should extend “whisk” to other exemplars of that kind, including a picture of a whisk that differs in color (here, silver) and a real three-dimensional silver whisk.

Notice that on this trial (where both the object and the picture were members of the named object category), if infants appreciate the referential status of novel nouns,

## LEARNING WORDS FROM PICTURES

then both choices may be seen as referents of “whisk”. But recall that in the training phase, infants had responded to questions in which only *one* choice (either the picture or the object) was correct. Therefore, to insure that infants understood that on this trial, they could select both the picture and the object, if they wished, we adopted the following strategy: if an infant selected only the picture or only the object in response to this question, the experimenter was instructed to go on saying, “There is another whisk here. Can you show me another whisk?” As will be clear in the Results section, because most infants spontaneously selected both the object and the picture, most did not receive this prompt. Strikingly, even among the few who did receive it, not a single infant changed their answer in response.

*Target Object – Distractor Object.* The experimenter presented two objects (silver whisk, silver garlic press) saying, “Can you show me the whisk?” At issue was whether infants would extend “whisk” selectively to another member of the named target kind, and not to an object of a different kind.

*Distractor Object - Target Picture.* The experimenter presented a real silver garlic press and a picture of a silver whisk, asking, “Can you show me the whisk?” At issue was whether infants’ interpretation of “whisk” was sufficiently robust to permit them to select the picture, even if this required ignoring a real three-dimensional object.

### 2.4 Coding

Infants’ behavior was coded from their videotaped sessions, following established criteria (Ganea et al., 2009; Preissler & Carey, 2004). We first identified all responses that were intentional. To be credited as intentional, infants had to indicate their choice(s)

## LEARNING WORDS FROM PICTURES

explicitly, either by giving it to the experimenter, pointing to it, or picking it up while making eye contact with the experimenter. Responses were coded as unintentional if an infant simply touched, played with, or explored an item but did not clearly indicate it to the experimenter. (If an infant failed to make an intentional response, the experimenter presented the items again and repeated the question before moving on to the next set of items.) For example, if on the Picture-Object test, an infant pointed to the picture while making eye contact with the experimenter, and then picked up the object only to play with it, this was coded as a “picture alone” response. If the infant indicated the object to the experimenter and then subsequently played with the picture, this was coded as an “object alone” response. If the infant pointed to both the picture and object while making eye contact with the experimenter, or gave both to the experimenter, this was coded as a “both” response, and we identified the order in which infants made their selections (simultaneously, object first, or picture first). A second rater, blind to the hypotheses, coded all test trials. Inter-rater reliability was 90%; disagreements were resolved by a third rater.

### 3. Results

Infants at both 15 and 17 months understood their task: 91% of their responses were intentional and unambiguous. Moreover, infants’ interpretation of the novel word was not tethered tightly to the particular perceptual features of the instances with which it had been introduced. Instead, infants at both ages spontaneously extended “whisk” to the real three-dimensional silver whisk, even though it differed both in color and



## LEARNING WORDS FROM PICTURES

representational status from the pictures with which this novel noun had previously been paired. Infants' responses to each of the three test questions are displayed in Table 1.

Test Question	Age	
	15 months (n = 12)	17 months (n = 12)
<b>Target Object vs. Target Picture</b>		
Both	7	11
Selection order:		
Object and picture simultaneously	3	4
Object first	4	6
Picture first	0	1
Object only	2	1
Picture Only	2	0
No Choice	1	0
<b>Target Object vs. Distractor Object</b>		
Both	0	0
Target Object only	9	10
Distractor Object only	0	0
No Choice	3	2
<b>Target Picture vs. Distractor Object</b>		
Both	0	0
Target Picture only	7	8
Distractor Object only	3	3
No Choice	2	1

Table 1. For each test question, the number of infants at each age choosing both test alternatives, only one of the two test alternatives, or making no choice.

Consider first infants' responses on the Target Picture-Target Object test.

Although “whisk” had been introduced exclusively in conjunction with pictures, infants did not favor the picture of a whisk over a real object. On the contrary, infants overwhelmingly extended “whisk” to the real silver whisk. This indicates that their representation of the novel word was sufficiently abstract to permit them to extend it systematically to new members of the same object kind, despite a change in color and in representational medium (from 2-dimensional to a 3-dimensional).

As can be seen in Table 1, 21 of the 24 infants (88%) chose the real object, either selecting alone or in conjunction with the picture (binomial test,  $p = .0001$ ); this pattern

## LEARNING WORDS FROM PICTURES

of performance did not differ as a function of infants' age (Fisher's exact test,  $p = .22$ ).

Table 1 also reveals that infants' predominant response on this question was to select both the object and picture. There was a reliable difference between infants' likelihood of extending the word to both the picture and object, the picture alone, the object alone, or not making a choice,  $\chi^2(3, 24) = 32.333, p < .0001$ . Their tendency to select both the picture and the object exceeded chance levels (binomial test,  $p = .008$ ), and this did not differ as a function of infants' age (Fisher's exact test,  $p = .16$ ). Moreover, among infants selecting both alternatives, some made their selections simultaneously and others made them sequentially. Among this latter group, infants were significantly more likely to select the object – rather than the picture -- first, (binomial test,  $p < .05$ ); this pattern was evident at both 15-months (all 4 infants making sequential selections chose the object first) and 17-months (6 of the 7 infants making sequential selections chose the object first), Fisher's exact test,  $p = .10$ . This clear tendency to select the object – rather than the picture – first is not consistent with the predictions of the associationist account. On this view, infants should favor the picture over the object.

Finally, we considered whether infants' tendency to select the object on this critical test trial might somehow have been influenced by their participation in other test trials. To address this concern, we focused on infants who received the Target Picture-Target Object first, before they had an opportunity to see the target object on any other test trial. The pattern was clear: Every infant this group first chose the object (either selecting it alone or together with the picture). This rules out any possibility that infants' tendency to select the object might somehow have been influenced by its presence in other test trials.

## LEARNING WORDS FROM PICTURES

Infants' responses on the remaining test questions provide converging evidence that infants extended "whisk" systematically to other members of the same object kind, despite changes in color (the picture of silver whisk) and in representational medium (the three-dimensional silver whisk) (Table 1). In response to these questions, not a single infant selected both alternatives. Instead, they systematically singled out the target, whether it was a picture or an object. On the Target Object–Distractor Object test, infants overwhelmingly selected the Target Object (silver whisk), but not Distractor Object (silver garlic press), binomial test,  $p = .003$ ; this robust pattern did not differ as a function of infants' age, Fisher's exact test,  $p = 1.0$ . This provides assurances that when infants extend "whisk" from pictures to objects, they do so selectively, systematically identifying a new 3-dimensional object from the same object category (silver whisk), but not a distractor object (silver garlic press). On the Target Picture – Distractor Object test, infants showed no tendency to favor a three-dimensional object over a picture. On the contrary, infants favored the picture of the whisk over the real garlic press by a factor of greater than 2:1, binomial test,  $p = .04$ ; this did not differ as a function of age, Fisher's exact test,  $p = 1.0$ . This provides additional assurances that infants' overwhelming selection of the real silver whisk in the Target Picture-Target Object test cannot be attributed to an *a priori* bias favoring objects over pictures.

Finally, we tabulated each individual infant's responses across the three test questions. Our goal was to discover whether individual infants extended "whisk" in a principled fashion, flexibly and correctly extending it to either the picture and to the real object, but systematically rejecting the non-whisk distractors. The probability of guessing correctly on all three trials is .125; if infants were guessing on each question, then only 3

## LEARNING WORDS FROM PICTURES

of the 24 infants should guess correctly on all three. Instead, we found that 17 of the 24 infants responded correctly on all three test questions, 8 at 15-months and 9 at 17-months (both  $p$ 's  $< .001$ , binomial tests). Taken together, these results reveal that infants' interpretation of the novel noun, paired exclusively with pictures during training, was sufficiently broad to include other members of the same object category, despite a change in both color and representational modality.

### 4. Discussion

The goal of the current experiment was to clarify the representational status of infants' newly learned words. In particular, we asked whether infants could extend a novel word, applied exclusively to pictures of novel object, to real three-dimensional members of the same object category. We introduced 15- and 17-month-old infants to a novel noun, applying it to two distinct pictures representing a particular object kind, and then asked whether infants would extend this newly-learned noun to other members of the same kind (other whisks), one which differed in color only (a picture of a silver whisk) and another which differed in both color and representational modality (a real three-dimensional silver whisk). If a word is nothing more than a feature of the perceptual and sensory experience(s) with which it has co-occurred, then infants should be more likely to extend the novel word to the picture of the whisk than to the real three-dimensional object, because the picture is perceptually closer to the instances on which the word had previously been paired. But if infants interpret a novel noun as referring to an object category, then they should extend the word to both the real three-dimensional

## LEARNING WORDS FROM PICTURES

silver whisk as well as the picture, because both are representations of the same object category.

The results were straightforward: 15- and 17-month-old infants' responses reflected neither (a) a strict reliance on perceptual similarity (which would have led them to favor the picture over the object in the Target Picture-Target Object test) nor (b) an overarching preference to select objects over pictures (which would have led them to favor the real garlic-press over the picture of the whisk in the Target Picture-Distractor Object test). Instead, infants interpreted "whisk" as referring to an object concept, extending it in a principled and flexible way to other whisks. This reveals that infants' interpretation of "whisk" was not tethered tightly to the particular perceptual features with which the word had previously been paired. Instead, their interpretation was sufficiently abstract to include new (and as yet, unnamed) members of the same object category, although they differ in color and in representational medium (real, three-dimensional objects).

At the same time, it is important to bear in mind that many infants were excluded here, as in prior reports with 15- and 19-month-olds using this paradigm (Ganea et al., 2009). (See Footnote 1.) Like Ganea et al. (2009), we have pointed out the advantages of including infants who successfully complete the training and who reveal no *a priori* 'object bias' (a bias to select an object over a picture, independent of the experimenter's request). Nonetheless, in future work it will be important to assess the generalizability of these results to a broader sample of 15- and 17-month-olds.

The benefits of introducing variability in learning. These results also shed light on infants' sensitivity to variability in learning. When infants and young children are

## LEARNING WORDS FROM PICTURES

exposed to variability in learning contexts, they acquire more abstract representations in tasks including grammatical rule learning (Torkildsen, Dailey, Aguilar, Gomez, & Plante, in press), object categorization and word learning (Waxman, 1990; Waxman & Klibanoff, 2000; Xu & Tenenbaum, 2007), and category-based induction (Rhodes, Gelman, & Brickman, 2010). The current results provide additional evidence that infants benefit from variability in the input. In Ganea et al. (2009, Study 2), where the novel noun was applied to a picture of a whisk that did not vary (four pictures of a white whisk), infants adopted a conservative interpretation, rarely (if ever) extending the noun to a three-dimensional object of a different color. The results of the current experiment suggest that this conservative pattern likely reflects infants' sensitivity to the range of exemplars with which novel words were introduced, rather than an inability to appreciate their abstract referential status. After all, we provided variability during learning, introducing the novel noun to two different pictures of a whisk (one purple whisk, one orange whisk). Under these conditions, infants at both ages spontaneously and systematically extended the novel noun to a new color (a picture of a silver whisk) and to a new medium (a real 3-dimensional silver whisk). At 17 months, 11/12 infants extended the noun to include the three-dimensional silver whisk; at 15 months, 10/12 of the infants did so. We suspect that this enabled infants in to extend the novel label broadly to the category "whisk" rather than adopting a more conservative interpretation. We therefore propose that infants' spontaneous extension of the novel noun – across colors and across representational media – reflects not only their appreciation of a link between words and object categories, but also their sensitivity to the input conditions.

## LEARNING WORDS FROM PICTURES

However, because we also introduced other procedural modifications that may have improved infants' performance (e.g., providing a training phase early in the procedure; providing infants with an opportunity to play with the test objects before engaging them in the experimental protocol<sup>3</sup>), additional research will be required to ascertain precisely which modification(s), individually or collectively, led to infants' successful extension of the novel noun. Indeed, there may be developmental change in how readily infants interpret a naming event – in a picture book or in the world at large – as kind-referring. The current findings serve as a demonstration that young word learners are capable of extending a label that is initially tied only to pictorial representations to the more abstract kind that the pictures represent.

## 5. Conclusion

In sum, the current results reveal that for infants as young as 15 months of age, words refer to categories of objects whose meaning extends beyond the particular items with which they were introduced. Armed with this early appreciation that nouns refer to object categories, infants attend carefully to the range of application for a novel noun and weave this together with their concepts of individuals and kinds to discover the meaning of a word. This work also offers 'good news' about reading picture books to infants and toddlers: Infants glean information from two-dimensional representations and apply it systematically to real three-dimensional objects when they encounter them.

## LEARNING WORDS FROM PICTURES

### **Acknowledgements**

This research was supported by NSF BCS-0950376 to SRW. We are grateful to Ilana Clift for assistance with coding, and to the children and caregivers who participated.



## LEARNING WORDS FROM PICTURES

## References

- Baillargeon, R. (2008). Innate ideas revisited: for a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science*, 3, 2–13.
- Baldwin, D.A. (1995). Understanding the link between joint attention and language. In C. Moore and P.J. Dunham (Eds.), *Joint attention: its origins and role in development* (pp. 131–158). Mahwah, NJ: Erlbaum.
- Baron-Cohen, S., Baldwin, D.A., & Crowson, M. (1997). Do children with autism use the speaker's direction of gaze strategy to crack the code of language? *Child Development*, 68(1), 48-57.
- Booth, A.E., & Waxman, S.R. (2003). Bringing theories of word learning in line with the evidence. *Cognition*, 87(3), 215-218.
- Booth, A., Waxman, S. R., & Huang, Y.T. (2005). Conceptual information permeates word learning in infancy. *Developmental Psychology*, 41(3), 491-505.
- Booth, A.E. & Waxman, S.R. (2006). Déjà vu all over again: re-re-visiting the conceptual status of early word learning: Comment on Smith and Samuelson (2006). *Developmental Psychology*, 42(6), 1344-1346.
- Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.
- Fenson, L., Dale, P.S., Reznick, J.S., Thal, D., Bates, E., Hartung, J.P., Pethick, S., & Reilly, J.S. (1993). *The MacArthur Communicative Development Inventories: User's Guide and Technical Manual*. Baltimore: Paul H. Brookes Publishing Co.
- Frith, U., & Happé, F. (1994). Autism: beyond “theory of mind”. *Cognition*, 50(1–3), 115–132.
- Ganea, P.A., Allen, M.A., Butler, L. Carey, S., & DeLoache, J.S. (2009). Toddlers' referential understanding of pictures. *Journal of Child Experimental Psychology*, (104), 283-295.
- Gelman, R., & Williams, E.M. (1998). Enabling constraints for cognitive development and learning: domain specificity and epigenesis. In D. Kuhn and R. Siegler, (Eds.), *Handbook of child psychology: Cognition, perception, and language* (Vol. 2, 4th ed., pp. 575–630). Hoboken, NJ: John Wiley & Sons, Inc.
- Gelman, S.A., & Waxman, S.R. (2009). Taking development seriously: Theories cannot emerge from associations alone. *Trends in Cognitive Sciences*, 13(8), 332-333.
- Graham, S.A., Booth, A., & Waxman, S.R. (2012). Words are not features of objects:

## LEARNING WORDS FROM PICTURES

- Only consistently applied nouns guide 4-year-olds' inferences about object categories. *Language Learning and Development*, (8), 1-11.
- Graham, S.A., Kilbreath C.S., & Welder, A.N. (2004). Thirteen-month-olds rely on shared labels and shape similarity for inductive inferences. *Child Development*, 75, 409–427.
- Horst, J. S., & Samuelson, L. K. (2008). Fast mapping but poor retention by 24-month-old infants. *Infancy*, 13(2), 128-157.
- Keates, J., & Graham, S.A. (2008). Category markers or attributes: Why do labels guide infants' inductive inferences? *Psychological Science*, 19, 1287-1293.
- Plunkett, K. (1997). Theories of early language acquisition. *Trends in Cognitive Science*, 1(4), 146-153.
- Preissler, M. A. (2008). Associative learning of pictures and words by low-functioning children with autism. *Autism*, 12(3), 231-248.
- Preissler, M.A., & Carey, S. (2004). Do both pictures and words function as symbols for 18- and 24-month-old children? *Journal of Cognitive Development*, 5, 185–212.
- Putnam, H. (1973). Meaning and reference. *Journal of Philosophy*, 70, 699–711.
- Rakison, D. H., & Lupyan, G. (2008). Developing object concepts in infancy: An associative learning perspective. *SRCD Monographs*.
- Rhodes, M., Gelman, S.A., & Brickman, D. (2010). Children's attention to sample composition in learning, teaching, and discovery. *Developmental Science*, 12, 421-429.
- Samuelson, L. K., & Smith, L.B. (1999). Early noun vocabularies: Do ontology, category structure and syntax correspond? *Cognition*, 73(1), 1-33.
- Sloutsky, V.M. (2009). Theories about “theories”: where is the explanation? Comment on Waxman and Gelman. *Trends in Cognitive Science*, 13, 331–332.
- Sloutsky, V.M., & Fisher, A.V. (in press). Linguistic labels: Conceptual markers or object features? *Journal of Experimental Child Psychology*.
- Sloutsky, V.M., Kloss, H., & Fisher, A. (2007). When looks are everything: appearance similarity versus kind information in early induction. *Psychological Science*, 18, 179–185.
- Sloutsky, V.M., Lo, Y., & Fisher, A. (2001). How much does a shared name make things

## LEARNING WORDS FROM PICTURES

- similar? Linguistic labels, similarity and the development of inductive inference. *Child Development*, 72, 1695–1709.
- Smith, L.B., & Samuelson, L. (2006). An attentional learning account of the shape bias: reply to Cimpian and Markman (2005) and Booth, Waxman, and Huang (2005). *American Psychological Association* 42 (6), 1339.
- Smith, L.B., Jones, S.S., Yoshida, H., & Colunga, E. (2003). Whose DAM account? Attentional learning explains Booth and Waxman. *Cognition*, 87(3), 209-213.
- Spelke, E.S. (2000). Core knowledge. *American Psychology*, 55, 1233–1243.
- Torkildsen, J. V. K., Dailey, N. S., Aguilar, J. M., Gomez, R. L., & Plante, E. (in press). Exemplar variability facilitates rapid learning of an otherwise unlearnable grammar. *Journal of Speech, Language, and Hearing Research*.
- Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development*, 5(2), 123-150.
- Waxman, S. R. (1999). Specifying the scope of 13-month-olds' expectations for novel words. *Cognition*, 70, B35-B50.
- Waxman, S.R., & Gelman, S.A. (2009). Early word-learning entails reference, not merely associations. *Trends in Cognitive Sciences*, (13)6, 332-333.
- Waxman, S.R., & Klibanoff, R.S. (2000). The role of comparison in the extension of novel adjectives. *Developmental Psychology*, 36(5), 571-581.
- Xu, F., Cote, M., & Baker, A. (2005). Labeling guides object individuation in 12-month-old infants. *Psychological Science*, 16(5), 372-377.
- Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, 10(3), 288-297.

---

<sup>1</sup>This rate is roughly comparable to the attrition rate reported in previous work. For example, in Ganea et al.'s (2009) Study 2, in which 15-month-olds were included and in which the procedure most closely resembles the one we have adopted here, 20 15-month-olds were engaged in the procedure. Of these, 5 were excluded for fussiness or failure to complete the procedure. Of the remaining 15, 8 were subsequently excluded from analysis because they demonstrated an "object bias". Thus, in Ganea et al.'s study, 65% of the 15-month-olds (13/20) were excluded from analysis. In the current experiment, 50% of the 15-month-olds were excluded from analysis: 12 due to fussiness, parental interference, or failure to respond correctly on one or more training trials.

<sup>2</sup> Ganea et al. (2009) also provided a training phase, but it differed from the current training phase in several ways. First, it followed (rather than preceded) the book-reading phase. Second, during training infants were presented with pictures of two *familiar* objects (ball, cup) and asked, "Show me the ball/cup". Third, children were presented with pictures of the novel objects (target, distractor) and asked to indicate the target ("Show me the whisk"). Immediately after this training phase, infants participated in an object familiarization phase during which the experimenter presented the novel target and distractor objects (different in color from the ones to be presented in testing), in succession for a few seconds each, saying, "Look at this" and allowed infants to handle the objects.

<sup>3</sup> A potential concern is that permitting infants to play with the unfamiliar objects in the Familiarization Phase, before introducing the novel words in the Picture-book Reading Phase, may have led infants to connect the novel noun to the unfamiliar object that they had seen before. Notice however, that for this to be the case, infants must already possess the representational capacity that we are testing. That is, if infants do indeed connect the novel noun (applied to a picture) to a real object that they seen earlier but in the absence of any name, this would mean that infants (a) appreciate that both the picture and the object are members of the same object category, and (b) appreciate that a novel noun applied to the picture alone would also apply to other members of that category, even if they differ in color and representational format.