Developmental Psychology
1992, Vol. 28, No. 5, 862-873

0012-1649/92/$3.00

Relations Among Word Meanings in Early Lexical Development

Sandra R. Waxman
Harvard University

Ann Senghas
Massachusetts Institute of Technology

When deriving the meanings of individual words, children must also discern the relations among words. To ascertain how children interpret such relations, we taught twelve 2-year-olds novel count nouns for related but unfamiliar objects. Although we never specified the relations among words, children's interpretation of the relations was mediated by the similarity of the objects. For dissimilar objects, children interpreted the words as mutually exclusive. For more similar objects, children's performance was consistent with a hierarchical interpretation of meaning. Thus, by 2 years of age, children have the conceptual and lexical abilities necessary for the establishment of hierarchical inclusion relations. The significance of this finding for theories of lexical and conceptual development is discussed.

In the arena of lexical acquisition, toddlers demonstrate exceptional expertise. Typically, they produce their first words close to their first birthday and acquire additional words gradually over the next several months. Sometime between approximately 17 and 24 months, toddlers experience a rapid vocabulary burst, during which time they learn to produce an average of six new words per day (Carey & Bartlett, 1978; Dromi, 1987). This rate of acquisition is nothing short of phenomenal. Yet there is more to the process of lexical acquisition than the sheer accumulation of words: In addition to deriving the meaning of individual words, the young child must simultaneously work out the relations among words in the lexicon.

Many different relations among words are possible. For example, two words may mark mutually exclusive classes of objects (e.g., *dog* and *cat*), or they may refer to hierarchically related classes of objects (e.g., *dog* and *terrier*). Typically, the precise relation among words is not made explicit to children, particularly in the earliest phases of lexical development. Instead, adults tend to introduce new words simply and rely heavily on ostensive definitions or point-and-label strategies (Callanan, 1985, 1989; Mervis, 1986). For example, a parent may point to a collie and ask, "See the collie?" Later, the parent may point to a terrier and announce, "Look at the dog." Because ostensive definitions like these underdetermine the scope and meaning of novel words (Ninio, 1980; Quine, 1960; Wittgenstein, 1953), it is left largely to the child to work out how the new words *dog* and *collie* are related.

A considerable amount of attention has been devoted to ascertaining how young word learners so successfully derive meaning from so little explicit guidance (Landau & Gleitman, 1985; Markman, 1989; Pinker, 1979; Waxman, 1990, 1991). Many researchers have suggested that children's discovery of meaning is guided by implicit biases or expectations that lead them to favor some possible meanings over others. It has been argued that these expectations reduce the ambiguity inherent in the word learning scenario by narrowing the range of possible meanings a child will consider when ascribing meaning to a novel word. For example, experimental research has shown that young children have a very strong expectation that the first unfamiliar word applied to an unfamiliar object will refer to the whole object (as opposed to any salient part or attribute; Markman, 1989) and may be extended to other members of its basic-level kind (Callanan, 1989; Markman & Hutchinson, 1984; Mervis & Canada, 1982). The child's powerful expectation is no doubt augmented by adults' reciprocal tendency to label objects initially at the basic level and to use this basic-level term to anchor subsequent meanings (Callanan, 1985; Mervis, 1986; Ninio, 1980; Shipley, Kuhn, & Madden, 1983).

Children are also able to take advantage of syntactic form class and other contextual clues to interpret the meaning of novel words. For example, once they have learned a basic-level term for an object, children may interpret subsequent words as referring to higher or lower order classes (Clark, Gelman, & Lane, 1985; Gelman, Wilcox, & Clark, 1989; Taylor & Gelman, 1988, 1989; Waxman, 1990); to particular attributes of an object, such as its shape, color, or texture (Au, 1990; Hall, Waxman, & Hurwitz, 1991; Landau, Smith, & Jones, 1988; Soja, Carey, & Spelke, 1991); to a salient part (Markman & Wachtel, 1988); to the particular named individual (Gelman & Taylor, 1984; Hall, 1991a; Katz, Baker, & Macnamara, 1974); or to individuals within a restricted life phase (e.g., baby and teen) or restricted context (e.g., passenger and pedestrian; Hall, 1991b; Hall & Waxman, 1991).

This line of research on early lexical development has yielded

Parts of this research were presented at the International Conference for Infancy Studies, April 1990, Montreal, Canada, and at the Boston University Child Language Conference, October 1990, Boston, Massachusetts. This research was supported by grants to Sandra R. Waxman from the Spencer Foundation and the National Academy of Education and from Harvard University.

We are grateful to the children, teachers, and parents of the Arlington Infant and Toddler Center, Harvard Yard Child Care Center, Radcliffe Child Development Center, and Soldier's Field Park Children's Center. Special thanks to A. Densmore for her contribution.

Correspondence concerning this article should be addressed to Sandra R. Waxman, who is now at Department of Psychology, Northwestern University, 2029 Sheridan Road, Evanston, Illinois 60208-2710.

important insights into the processes young children use to derive the meanings of and relations among novel words. However, the empirical investigations conducted thus far have focused almost exclusively on one aspect of this complex process. Most have examined how children work out the relations among multiple terms (e.g., basic, superordinate, subordinate, and property terms) applied to a single object or class of objects. For example, researchers have examined how children learn that the terms *dog, animal, collie, puppy, fluffy,* and *friendly* may all correctly apply to a given object or class of objects.

In the present experiments, we addressed the problem of word learning from a different vantage point. We suspect that children must also encounter interpretive difficulties when it comes to learning the scope and meaning of different words applied to different objects or classes of objects. Consider again the parent who points to a collie and later to a terrier and labels them a *collie* and a *dog,* respectively. Another parent, under the same external circumstances, might label these objects as a *collie* and a *terrier,* respectively. In both cases, the context in which the novel words are introduced is identical. However, the words used and the intended relations among them differ. In the former case, the terms are related by inclusion; in the latter, they are mutually exclusive. How, then, do children come to determine the meanings of and relations among novel words?

To explore this question, we used a method that weds naturalistic observation with controlled experimental procedure and is therefore particularly well suited to exploring lexical acquisition in toddlers (Tomasello, Mannle, & Werdenschlag, 1988). We introduced 2-year-old children to novel words in the context of four informal play sessions. Each child learned three novel words. Two of these (e.g., *horn* and *flute*) referred to a pair of related target objects (e.g., a small plastic horn and a small plastic flute). The third word (e.g., *whisk*) referred to a third unfamiliar target object (e.g., a wire whisk) that was unrelated to the original pair. For clarity of exposition, we have adopted the following terminology: We refer to the members of the related pair as A1 and A2 and refer to the unrelated object as B. In a preliminary study (Experiment 1), we confirmed that A1 and A2 were perceptually more similar to one another then either was to B.

In Experiment 2, we observed toddlers' interpretations of the novel words and the relations among them, using both production and comprehension measures. On the first visit, we introduced only Object A1 (e.g., the horn) in conjunction with Word A1 (e.g., *horn*). On two subsequent visits, we introduced Object A2 (e.g., the flute) in conjunction with Word A2 (e.g., *flute*) and Object B (e.g., the whisk) in conjunction with Word B (e.g., *whisk*). The fourth visit was a follow-up, conducted 1 week later.

This design yielded detailed information concerning toddlers' interpretation of novel words and the relations among them. In addition, because the evolution of word meaning cannot be captured in a single visit (Vygotsky, 1962), we charted each child's pattern of interpretation at each visit. In this way, we were able to observe any changes in interpretation over the course of the experiment.

The design also gave us an opportunity to ascertain whether children regarded Objects A1 and A2 as more similar to one another than they were to Object B. We expected children's extensions of the novel terms to reflect this in the following manner: We predicted that children might extend a label applied to one member of the related pair (e.g., A1) to include the other member (e.g., A2) but not to include the unrelated object (B). Similarly, we predicted that children would not generalize the term for the unrelated object to include the related targets.

Also incorporated into our design was the assumption that children are capable of fast-mapping (Carey, 1978; Carey & Bartlett, 1978; Heibeck & Markman, 1987) and therefore would apply the target words systematically to target objects. To check this assumption, we included several nontarget objects whose labels were unfamiliar to the children. We predicted that children would map the target words onto the target objects and would not generalize them indiscriminantly to the unfamiliar distractor objects.

Our primary question concerned the children's interpretations of the words for Objects A1 and A2 and the relation between them. As can be seen in Table 1, four distinct relations between the words are consistent with the simple conditions under which we introduced the words. Perhaps the most straightforward interpretation is Strategy 1. In this interpretation, the child would assume that the two words are mutually exclusive with respect to the target objects. To make this interpretation, the child would simply map each novel word conservatively onto the object on which it was taught. This interpretation would not require the child to go beyond the input offered by the experimenter.

The next three interpretations are also consistent with the experimenter's input, but each would involve extending the novel terms beyond the input in a systematic way. For instance, in Strategy 2, the child would interpret the two words as interchangeable with respect to the related target objects (e.g., *rabbit* and *bunny*). To make this interpretation, a child would apply both of the novel words to both members of the related pair.[1]

In Strategies 3 and 4, the child spontaneously would extend one of the terms to include both of the related target objects, but would reserve the other term exclusively for the object on which it was taught. Children who adopted Strategy 3 would extend the first word (e.g., *dog*) to include both members of the related pair, but restrict the second word to the object on which it was taught (e.g., *terrier*). Conversely, children who adopted Strategy 4, would restrict the first word (e.g., *terrier*) to the object on which it was taught, but extend the second word to include both members of the related pair (e.g., *dog*).

Although Strategies 3 and 4 are consistent with a hierarchical pattern of interpretation, it is important to note that they are also consistent with other pairs of meanings that overlap. For example, consider the classes *former president* and *born in Massachusetts.* Although both John Kennedy and Richard Nixon

[1] Note that there are several independent ways in which the child could arrive at this pattern of extension. First, the words may be interpreted as synonyms (but see Clark, 1987, for the argument that there are no true synonyms). Second, like the words *cop* and *policeman,* the words may pick out the same set of referents, differing only in connotation (see Clark, 1987). Third, the words may refer to two different categories with overlapping extensions, such that each category includes both of the target objects among its members. For example, the words *mammal* and *animal* may be used interchangeably to refer to dogs and cats (but not fish); the words *mammal* and *pet* may be used interchangeably when referring to dogs and cats (but not parakeets or grizzly bears).

Table 1
Four Possible Interpretations of the Relation Between Words A1 and A2

Interpretation	Behavior
1. Words A1 and A2 are mutually exclusive (e.g., A1 = *dog*, A2 = *cat*).	Word A1 is applied exclusively to Object A1. Word A2 is applied exclusively to Object A2.
2. Words A1 and A2 are interchangeable (e.g., A1 = *rabbit*, A2 = *bunny*).	Word A1 is applied to both Objects A1 and A2. Word A2 is applied to both Objects A1 and A2.
3. Word A1 is inclusive; Word A2 is restrictive (e.g., A1 = *dog*, A2 = *terrier*).	Word A1 is applied to both Objects A1 and A2. Word A2 is restricted to Object A2.
4. Word A2 is inclusive; Word A1 is restrictive (e.g., A1 = *terrier*, A2 = *dog*).	Word A1 is restricted to Object A1. Word A2 is applied to both Objects A1 and A2.

are members of the class *former president,* and although only Kennedy is also a member of the class of individuals *born in Massachusetts,* it does not follow that the class *born in Massachusetts* constitutes a proper subset of the class *former presidents.* Therefore, Strategies 3 and 4 are consistent with, but not exclusive to, a hierarchical interpretation of meaning.

Each of the four interpretations outlined in Table 1 is consistent with Clark's (1987) principle of contrast, which states that no two words in a given language are perfectly synonymous. However, only Strategy 1 is also consistent with the more stringent principle of mutual exclusivity (Markman, 1989), which states that children assume that words will refer to categories that do not overlap. Strategies 2–4 represent violations of this principle because each involves overlap in the extensions of the target labels.

Strategies 2–4 may be distinguished in the following manner. We assume that Strategy 2, in which a child would systematically extend each target word to include both related objects, rests on an appreciation of the similarity of the two target objects. We assume that Strategies 3 and 4 go beyond this appreciation in one crucial respect: By extending one word while simultaneously restricting the other, a child would signal an implicit appreciation of both the similarities as well as the differences between the target objects. The ability to appreciate simultaneously the similarities and differences among objects is crucial to the development of class–subclass relations. The ability to mark these relations lexically (as in Strategies 3 and 4) is crucial to the establishment of hierarchically related meanings. In view of the extensive literature concerning the late emergence of hierarchical inclusion relations, one might expect Strategies 3 and 4 to occur infrequently in children as young as 2 years of age.

Experiment 1: Ensuring Perceptual Similarity Between Related Targets

The purpose of Experiment 1 was to gather information concerning the perceptual similarity of the target objects to be included in the subsequent toddler study (Experiment 2). In particular, we sought to ensure that A1 and A2 were, in fact, more similar to one another than either was to B.

Method

Subjects

Twenty-one undergraduate students at Harvard University served as subjects. There were approximately equal numbers of men and women.

Materials

Materials consisted of three pairs of related objects. These included (a) a toy horn and a toy flute, (b) a hook and a clip, and (c) a whisk and a tong. The horn and flute were made by the same manufacturer, and each measured approximately 21 cm in length. The horn, which resembled a trumpet, was pink and green; the flute, which resembled a recorder, was pink and yellow. Both were made of plastic. The clip and the hook were each red, were each mounted on a red plastic disk with a magnetic backing, and were each approximately 7 cm long. Both the whisk and the tongs were made of stainless steel, had red rubber-coated plastic handles, and were approximately 30 cm long.

Procedure

To begin, the experimenter explained that she was interested in gathering adults' ratings of the perceptual similarity between pairs of objects in an effort to select materials for a subsequent study to be conducted with 2-year-old children. She further explained that because the 2-year-olds would not know the words for the various objects, they should strive to complete their ratings on the basis of the perceptual similarity of the objects themselves, rather than on their labels. She then presented the subjects with all possible pairs of those targets in random order.

Scoring

Subjects rated the similarity of each pair on a 7-point scale from *extremely similar* (1) to *extremely dissimilar* (7).

Results

The similarity ratings were distributed discontinuously, with mean scores for all three pairs of related objects falling at or below 3.10 and mean scores for all unrelated objects (e.g., hook–flute) falling uniformly above 5.30. The mean similarity rating and standard deviation for each related pair were as follows: For

horn–flute, $M = 2.38$, $SD = 1.02$; for hook–clip, $M = 1.76$, $SD = .77$; and for whisk–tong, $M = 3.10$; $SD = .10$. The mean similarity ratings for the unrelated pairs (e.g., hook–flute) ranged from 5.30 to 6.33. A contrast analysis conducted on the similarity ratings obtained for the three pairs of related objects revealed small, but significant, differences among them (all $ps < .05$). This suggests that our related target pairs actually covered a range of perceptual similarity.

To determine whether A1 and A2 were indeed more closely related to one another than either was to any of the other objects in the set, we compared the similarity rating obtained for each pair of related objects with the similarity rating obtained when each member of the pair was presented in conjunction with each of the other, unrelated objects from the set. For example, we compared the mean similarity rating for hook–clip with the mean similarity rating for each of the following pairs: hook–flute, hook–tong, clip–flute, and clip–tong. In all cases, the pair of related objects were rated as significantly more similar to one another than either was to the other, unrelated objects (all $ps < .0001$).

In sum, Experiment 1 ensured that A1 and A2 were more similar to each other than either was to B.

Experiment 2: Word-Learning Strategies in Toddlers

In Experiment 2, we turned our attention to the toddlers and asked how they interpreted the meanings of and relations among novel words applied to these objects.

Method

Subjects

Twelve toddlers, with a mean age of 25 months (range = 22–32 months), served as subjects. All were enrolled in preschools serving middle- and upper middle-class populations around Boston and Cambridge, Massachusetts. All were native speakers of English. One boy, who produced no intelligible speech in the first session, was replaced.

Materials

A complete list of objects used in this experiment may be found in Table 2.

Target objects. The three pairs of target objects were those described in Experiment 1. Each child learned one of the pairs as well as one unrelated object selected from one of the remaining two pairs. The unrelated object for each child was balanced in the following manner.

Table 2
Complete List of Stimuli

Nontarget objects		Target objects
Familiar	Unfamiliar	
Box	Stopwatch	Horn–flute
Pig	Pinch-purse (rubber)	Hook–clip
Truck	Egg separator	Whisk–tong
Doll shirt		
Two plastic cups		
Horse		
Rooster		

For the children learning the horn–flute pair, half learned the tongs as their B, and half learned the hook as their B. For the children learning the hook–clip pair, half learned the flute as their B, and half learned the tongs as their B. For children learning the whisk–tong pair, half learned the flute as their B, and half learned the hook as their B. The mean ages for children learning the horn–flute, hook–clip, and whisk–tong pairs were 25.5 months, 28.5 months, and 25.5 months, respectively.

We made every effort to use the novel terms as naturally as possible. However, for the purposes of experimental design, there were some minor differences between our usage and that typical of everyday discourse. For instance, although *tongs* (like *scissors*) is pluralized in the adult lexicon, we introduced it as a singular count noun to ensure that all target words were presented in identical syntactic contexts. That is, we introduced the tongs as the (or a) *tong*. In addition, we labeled the whisk with a novel word, *jire*. This was because in pilot work, we found ourselves unable to unambiguously determine how the children were answering the experimenter's query, "What's this?" The toddlers' articulation was not clear enough to distinguish their production of *whisk* from their production of *this*.

Nontarget objects. The remaining objects used in the experiment were selected on the basis of a control study conducted in advance of the experiment proper with an independent group of eight 2-year-old children, each of whom was asked to label a series of objects. From these, we selected 11 nontarget objects, which were seen by all children in the current experiment. Of these 11, 8 were familiar objects that all control subjects readily labeled at the basic level. The remaining 3 were unfamiliar objects. None of the control subjects were able to provide labels for these.

Procedure

To simulate the natural word-learning process, we introduced toddlers to their designated target objects and target labels in the context of four informal play sessions. At each visit, the child and one of the experimenters sat either on the floor in a quiet corner of the child's preschool room or in an adjoining room. The other experimenter videotaped the visit and assisted the first experimenter.

At the beginning of each visit, the experimenter emptied a bag of toys onto the floor within easy reach of the child. The child was encouraged to play freely with the toys. During the first three visits, the experimenter modeled the novel labels for the target items. For a labeling episode to count as a model, the following criteria had to be met: The experimenter had to produce the target label, with stress and in phrase-final position, while the child's attention was focused on the target object. If the child looked away from the object during this time, the modeling episode was repeated. Twice during each visit, the experimenter tested the child's production and comprehension of the target words.

The first three visits were all completed within a 2-week period, and each lasted approximately 20 min. The fourth visit, which took place between 7 and 10 days after the third visit, lasted approximately 10–15 min. All visits were videotaped for later transcription. The details of the procedure adopted for each visit are described in detail below.

Visit 1. The goal of the first session was to introduce children to Object A1 and teach them Word A1. Eight toys were included in the first visit: one target item (a horn, a clip, or a whisk), four familiar items, and three unfamiliar items. The experimenter picked up each object, one at a time, and asked the child to name it. She selected the objects in a different random order for each child, with the following two constraints: (a) A familiar object was always selected first, and (b) Object A1 was never selected first or last.

We took the following steps to ensure that although the toddlers may have seen some of the target objects around their homes or schools, the target words they were about to learn were indeed unfamiliar to them.

If a child was able to label the target object, a different target was selected. Even if a child provided an incorrect label for the target from the appropriate semantic domain (e.g., if a child labeled the flute a *whistle*), another target was selected for that child. When an appropriate target object was determined, the experimenter labeled it clearly for the child. This initial labeling episode counted as the first modeling of the target word.

If the child offered a correct label for the familiar nontarget objects, the experimenter simply repeated it. In most cases, the experimenter adopted the child's label for the familiar toys. For example, some children labeled the plastic rooster *a rooster;* others labeled it *a chicken.* Occasionally, a child neglected to label one of the familiar toys. In such cases, the experimenter offered the correct label. For the unfamiliar nontarget objects, the experimenter accepted any label offered by the child but did not supply a label if the child failed to do so.

For the remainder of the session, the experimenter and child continued to play with the objects. During this time, the experimenter modeled Word A1 on 10 different occasions. After the 5th and the 10th modeling, she tested the child's production and comprehension of the target word in succession. To test production, she drew the child's attention to the target item and asked, "What is this?" To test comprehension, she placed all toys within the reach of the child and asked, "Where is the _____?" or "Can you put the _____ on the box?" Children were given no corrective feedback.

Visits 2 and 3. The goals of Visits 2 and 3 were to (a) introduce children to Objects A2 and B and the target words and (b) examine the children's interpretation of all three target words (A1, A2, and B). A total of 13 objects were included in Visits 2 and 3. These included the 8 items used in Visit 1, 3 additional familiar objects, and 2 additional target objects (A2 and B; see Table 1). For example, a child who had learned *horn* as Word A1 in Visit 1 learned *flute* as Word A2 and either *whisk* or *hook* as Word B. The procedures used in Visits 2 and 3 were identical.

As in Visit 1, the experimenter picked up each of the toys one at a time and asked the child to name it. The toys were picked up in a different random order for each child, with the following three constraints: (a) A familiar toy was always selected first, (b) target toys were never selected first or last, and (c) Object A1 was selected before either Object A2 or Object B.

If the child offered a correct label for Object A1 and the familiar nontargets, the experimenter simply repeated it; if the child failed to label the toy, the experimenter offered the correct label. For the unfamiliar nontarget objects, the experimenter acknowledged any label offered by the child but did not supply a label if the child failed to do so. For the new target objects (A2 and B), none of the children provided correct labels. The experimenter introduced these labels. These initial labeling episodes constituted the first modeling of Words A2 and B.

As the experimenter and child continued to play with the toys, the experimenter modeled Words A2 and B on 10 different occasions each. After the 5th and the 10th modeling, she tested the child's production of all three target words, in random order. This was immediately followed by assessment of the child's comprehension of all three target words (A1, A2, and B). No corrective feedback was provided.

Visit 4. The fourth session constituted a follow-up visit and was conducted between 7 and 10 days after the third visit. Because this session included no modeling, it lasted only 10–15 min. The experimenter began by taking the toys out of the bag one at a time in random order, with the following two constraints: (a) The target items (A1, A2, and B) were never the first or last items presented, and (b) presentation of Object A1 preceded presentation of the A2 and B objects. Twice during the session, the experimenter elicited production of all three target words in random order. She then tested comprehension of all three target words in random order. The experimenter did not repeat or correct the child's labels.

Coding

The videotaped interaction of each visit was transcribed for all episodes related to the target objects, including the utterances that accompanied each episode, the context in which the utterance occurred, and the producer of the utterance (experimenter or child). The transcripts included all productions, both spontaneous and elicited.

Results

Children participated enthusiastically and quickly and accurately mapped the target words onto appropriate target objects. During Visit 1, children produced Word A1 an average of 2.9 times. On Visits 2–4, children produced each target term an average of 2.5, 3.4, and 3.6 times, respectively. Further, children did not apply the target words indiscriminantly to the other unfamiliar objects within their reach. In the first visit, they applied Word A1 correctly at a rate of 97.1% and 91.6% in production and comprehension, respectively.[2] Percentages reflect the number of correct usages of a word in either production or comprehension expressed as a function of the total number of usages of that word in either production or comprehension. For example, if a child produced Word A1 six times in a given session, and if five of these six productions were in reference to Object A1, then that child would be credited with producing Word A1 accurately at a rate of 83.5% for that session. By the end of the first visit, fully 100% of the toddlers had successfully learned to map Word A1 onto Object A1.

Performance on subsequent visits revealed the systematicity with which the toddlers mapped out the relations among the three target terms. A preliminary analysis revealed no reliable differences among the patterns of extension obtained in Visits 2–4. Therefore, we pooled the data from these visits for subsequent analysis. The pooled data are displayed in Tables 3 and 4 for production and comprehension, respectively.

Consider first the acquisition of Word B. As predicted, children applied Word B to Object B almost exclusively. Interestingly, they achieved a greater degree of accuracy in production (97%) than in comprehension (79%), paired $t(11) = 2.91$, $p < .01$, a point to which we will return later in the discussion. Their successful acquisition of Word B makes it clear that the toddlers were not overwhelmed by the process of learning three novel words for three unfamiliar objects in this experiment. This finding is important because it sets the stage for addressing our principal research question: How do children interpret the relation between the two related target words?

To address this question thoroughly, we examined children's use of Words A1 and A2 in two different but complementary ways. We first analyzed the group data that are summarized in Tables 3 and 4 and then examined the particular interpretation adopted by each individual toddler concerning the relation between the A1 and A2 terms.

[2] In Visit 1, only 1 child misapplied Word A1, and she did so on only one of eight productions. In comprehension, 2 children made errors on the first of their two comprehension probes. All other children responded correctly to both of the comprehension probes.

Table 3
Percentage of Extensions of the Three Target Words to the Three Target Objects During Production in Visits 2–4

	Referent		
Word (used by child)	A1	A2	B
A1	63	32	2
A2	8	91	0
B	0	3	97

Table 4
Percentage of Extensions of the Three Target Words to the Three Target Objects During Comprehension in Visits 2–4

	Referent (selected by child)		
Word (used by experimenter)	A1	A2	B
A1	77	21	1
A2	10	82	3
B	4	11	79

Group Data

Inspection of Tables 3 and 4 reveals that in both production and comprehension, children's most frequent application of each novel label was to the particular target object on which it was taught. For example, children applied Word A1 to Object A1 on 63% and 77% of their production and comprehension trials, respectively. The difference between production and comprehension approached significance, paired $t(11) = 2.12$, $p < .06$. They applied Word A2 to Object A2 on 91% and 82% of their production and comprehension trials, respectively. This difference did not reach significance.

Although children primarily mapped the target words onto the objects on which they were taught, they did not restrict themselves to this application alone. They also went beyond the information provided in the experimenter's modeling to extend the novel words to the other related target objects. That is, they extended Word A1 to Object A2 and extended Word A2 to Object A1. Although in principle these two types of extension are equally likely, children extended Word A1 to Object A2 more often than they extended Word A2 to Object A1. This asymmetry obtained in both production, paired $t(11) = 3.24$, $p < .01$, and comprehension, paired $t(11) = 2.28$, $p < .05$.

The preceding analyses support two conclusions. First, the fact that children's extensions of Words A1 and A2 rarely included Object B illustrates that their extension of novel terms is principled and guided by the similarity of the novel objects under consideration. Second, the fact that children extended a label applied to one member of the related pair to include the other member reveals that they readily go beyond the input they receive when ascribing meaning to a novel word.

In the next analysis, we asked whether the patterns of extension of the target terms varied as a function of the particular set of objects to which the children were introduced. Recall that in Experiment 1, we found that the three related pairs (A1 and A2) fell along a continuum of similarity; the similarity rating for each pair differed significantly from the ratings for all other pairs. The hook and clip were rated as most similar, followed by the flute and horn, and finally the whisk and tong. We therefore sought to ascertain whether this variation in degree of similarity of Objects A1 and A2 would be reflected in the children's interpretations of the target words. In particular, we suspected that children's tendency to extend a label (e.g., Word A1) applied to one member of a related pair to include the other member (e.g., Object A2) might vary as a function of the similarity of the particular pair. Performance on Visit 2, when Object A2 was first shown, provided a hint that this might be the case. Four children (2 in the horn–flute condition and 2 in the hook–clip condition) initially extended Word A1 to include Object A2. Following strict experimental protocol, the experimenter acknowledged the children's initial labels and went on to offer the A2 label, as she did with all children. It is interesting to note that the subsequent performance of these 4 children was indistinguishable from that of the remaining 4 children (2 in the horn–flute and 2 in the hook–clip condition) who did not offer the A1 term initially. Table 5 displays the relevant production and comprehension data for all children in each of the three A1–A2 pairs, pooled for Visits 2–4.

An analysis of variance (ANOVA) based on children's tendency to produce Word A1 in reference to Object A2 revealed a main effect for stimulus set, $F(2, 10) = 15.263$, $p < .002$. Post hoc analyses (Fisher's least significant difference) indicated that children who learned the hook–clip and the horn–flute pairs were more likely to extend Word A1 to Object A2 than were children learning the whisk–tong pair. There was no difference between the former two pairs.

Consider next the tendency to produce Word A2 to refer to Object A1. An ANOVA based on this measure also revealed a main effect for stimulus set, $F(2, 10) = 4.64$, $p < .04$. As can be seen in Table 5, there were no instances of this type of extension in either horn–flute or whisk–tong pairs. Only children who learned the hook–clip pair produced Word A2 in reference to Object A1. A careful inspection of the data revealed that across all sessions, there were only four such occurrences (accounting

Table 5
Percentage of Extensions of the Two Related Target Words to the Related Target Pairs for Each Target Pair During Production and Comprehension in Visits 2–4

	Referent					
	Horn–flute		Hook–clip		Whisk–tong	
Word	A1	A2	A1	A2	A1	A2
	Production					
A1	50	47	54	38	93	4
A2	0	99	23	77	0	1
	Comprehension					
A1	69	31	70	27	89	8
A2	23	64	13	79	0	94

for 23% of the productions of Word A2 in the hook–clip condition). One child never produced Word A2 in reference to Object A1; 2 children did so on only one occasion each; the remaining child produced Word A2 to refer to Object A1 twice. Interestingly, on both occasions, this child corrected herself immediately by spontaneously producing Word A1.

Although the general pattern of results observed in the comprehension data parallels that obtained in production, ANOVAs based on the comprehension data revealed no significant main effects or interactions.

Taken together, these analyses support the hypothesis that children's tendency to extend one of the A terms to include the other A object varied as a function of the particular A1–A2 pair to which they had been introduced. Our analyses were limited, however, in one very essential respect. Because they were based on group data, they did not permit us to ascertain precisely how each individual child interpreted the relation between the two related words. We addressed this issue directly by characterizing each child according to the particular strategy adopted in interpreting the relation between the A1 and A2 words.

Individual Children's Interpretations

To determine individual strategies, we characterized each child's pattern of extension on both the production and comprehension measures. For each measure, a child was characterized as making a mutually exclusive interpretation (Strategy 1) if that child restricted his or her application of Words A1 and A2 to the particular objects on which the words were taught. A child was characterized as making an interchangeable interpretation (Strategy 2) if that child extended each of the target words to include both of the related target objects. A child was characterized as making an inclusion interpretation (Strategy 3 or 4) if that child extended one of the target terms to include both Objects A1 and A2 but restricted application of the other target term to the particular item on which it was taught. For a child to be credited with extending one word to include the other related object, the extension had to occur on more than one occasion during a given session. We adopted this criterion to distinguish systematic extensions of the novel words from slips of tongue or occasional accidental misapplications.

Before we turn to the strategies adopted by the individual subjects, one final note bears mention. A child who adopts an inclusion strategy in production may, at the same time, reveal a mutually exclusive pattern in comprehension. Consider, for example, an individual who has at his or her command only two terms to refer to cows: *cow* and *guernsey*. This individual, like a child in our experiment, may well use the term *cow* productively to refer to both a hereford and a guernsey, while restricting the word *guernsey* to the latter. Nonetheless, when asked to select a cow on a comprehension trial, she should prefer to pick the hereford, because (according to the Gricean maxims of conversation) if the speaker had intended to refer to the guernsey, he or she would have said so. Therefore, a mutually exclusive pattern in comprehension is, in fact, entirely consistent with either a mutually exclusive or an inclusion strategy.

A close inspection of the children's individual strategies revealed two important findings. First, what strategy a child used depended on the specific object pair to which the child had been introduced. Second, children as young as 2 years of age spontaneously extended one term to include both related target objects but restricted the other term to the object on which it was taught.

All 4 of the children who learned the whisk–tong pair interpreted the words as mutually exclusive (Strategy 1). In both production and comprehension, these children restricted their application of Word A1 to Object A1 and restricted their application of Word A2 to Object A2. Furthermore, these children adopted a mutually exclusive interpretation from the start (Visit 2) and maintained this interpretation on each subsequent visit.

In contrast, all 8 children who learned the hook–clip or horn–flute pair adopted Strategy 3. In production, all of these children consistently applied Word A1 to both Objects A1 and A2 but reserved Word A2 for Object A2. They adopted this production strategy from the start (Visit 2) and consistently exhibited it on all subsequent visits. In comprehension, 6 of the 8 children treated the words as mutually exclusive with respect to the targets on all visits. (The 2 remaining children, 1 who learned the clip–hook pair and 1 who learned the horn–flute pair, treated the two target words interchangeably in comprehension in one session, but abandoned this pattern in favor of a mutually exclusive comprehension pattern in the remaining two sessions.)

In sum, children who learned the whisk–tong pair interpreted the words as mutually exclusive whereas those who learned the other target pairs adopted an inclusion strategy. Recall that in Experiment 1, adults judged that the three related target pairs fell along a continuum of perceptual similarity. It was therefore somewhat surprising that the children's word-learning strategies did not mirror this continuum but instead fell into two distinct groups.

To examine more closely the correlation between the strategies adopted by individual children and the target pair to which they had been introduced, we reviewed our videotaped sessions to further inspect the circumstances under which the novel words and objects had been introduced to the children. We found that when the experimenter introduced the flute and the horn, she used identical actions. For example, she would blow on each instrument and move her fingers to vary their tone. Similarly, when she introduced the hook and the clip, she used identical actions. For instance, she would attach these to other objects and would stick each onto a magnetized section of the toy box. However, the experimenter used different actions when she introduced the whisk and the tongs. She rubbed the whisk lightly over the box (or over the child's shoe) to create a soft rhythm. She also placed small objects in the wire end and shook it like a rattle. She performed none of these actions with the tongs. Instead, she used the tongs to pick up small objects and to tap on the toy box or the floor.

Thus, although we had taken care to select object pairs (e.g., Objects A1 and A2) that bore clear static *perceptual* resemblances, perhaps the accompanying functional information provided when the experimenter introduced the objects rendered some object pairs (the horn–flute and hook–clip pairs) more closely related than others (the whisk–tong pair).

Follow-Up Study: The Contribution of Functional and Perceptual Similarity

To assess the possibility that the functional information the experimenter provided when she introduced objects made

some object pairs more closely related, we conducted a follow-up study with 21 adult subjects. All subjects were students at Harvard University and were unaware of the purpose of the study or the outcome of the previous study. The procedure was identical to that used in Experiment 1 with one exception: Subjects in this experiment were instructed to rate the objects on the basis of both perceptual and functional similarity. The experimenter explained that she was gathering similarity ratings from adults to pursue a finding previously obtained with 2-year-old children. She further explained that the toddlers did not know the words for the various objects and therefore they should strive to complete their ratings on the basis of perceptual and functional aspects of the objects themselves, rather than on their labels. Next, the experimenter held up each of the six target objects, one at a time in random order, and briefly demonstrated the functions that had been used in the experiment proper. She then presented the subjects with all possible pairs of those targets, in random order, and subjects rated the similarity of each pair on a 7-point scale from *extremely similar* (1) to *extremely dissimilar* (7).

Like the toddlers in the word-learning study, adults who were provided with both perceptual and functional information treated the whisk–tong pair differently than they did either of the other two pairs: The whisk–tong pair was singled out. Contrast analyses revealed significant differences between the whisk–tong pair ($M = 3.24$, $SD = 1.22$) and each of the other two related pairs (both $ps < .001$). There was no difference in the adults' ratings of the hook–clip pair ($M = 2.10$, $SD = 1.38$) and horn–flute pair ($M = 1.91$, $SD = .77$). Thus, when adults were asked to consider both perceptual and functional information, their intuitions corresponded well with the strategies adopted by the toddlers in word learning.

These data are consistent with a very rich literature indicating that children, like adults, take advantage of both perceptual and functional similarity, as well as correlations between them, in establishing lexical and conceptual categories (Gentner, 1978; Gentner & Rattermann, 1991; Medin & Shoben, 1988; Mervis, 1986; Nelson, 1974; Tomikawa & Dodd, 1980; Tversky, 1985).

General Discussion

These data provide a rich and detailed portrayal of 12 toddlers' performance in the exacting task of interpretating the meaning of novel words and the relations among them. In most of the current experimental literature concerning early word learning, a principal focus has been on the process by which children work out the relations among multiple labels applied to a single object or class of objects (e.g., *dog, puppy, animal, collie,* and *furry*). In the present study, we examined the means by which children work out the relations among different labels (e.g., *horn* and *flute*) that are applied to different, but related, objects (e.g., a horn and a flute). We also examined the stability of their interpretations over a series of visits.

Summary of Findings

Twenty-five-month-old children learned three novel words over the course of four visits. Words A1 and A2 referred to a pair of related objects; Word B referred to an object that was unrelated to the pair. (See Experiment 1 for evidence concerning the perceptual similarity between the related and the unrelated objects.) The toddlers interpreted the novel words and the relations among them in a remarkably systematic fashion. Their success in this demanding task is consistent with the view that children are capable of rapidly mapping novel words to appropriate referents. They did not apply the target words to other novel objects or apply the term for Object B to Objects A1 and A2.

Moreover, the toddlers adopted clear strategies for working out the relations between the two related target words. These strategies were mediated by the functional and perceptual similarity among the target objects. (See the follow-up study for evidence pertaining to the functional and perceptual similarity among the objects.) When Object A2 was sufficiently different from Object A1 (as was the case for the whisk–tong pair), children assumed that the two labels were mutually exclusive, with each presumably referring to contrastive basic-level kinds. But when Objects A1 and A2 were more similar (as was the case for the horn–flute and hook–clip pairs), the children's interpretation was quite different. In this case, the toddlers assumed that both objects were members of a common class. They interpreted one word as referring to both objects (presumably at the basic level) and the other as referring to only one of the objects. In particular, they extended the first label to both Objects A1 and A2 but restricted the second label to the particular object on which it was taught.

Theoretical Issues in Children's Conceptual and Lexical Development

These results are important for theories of conceptual and lexical development. They reveal that the manner in which children's extensions of novel words applied to unfamiliar objects is guided by the similarity of the objects under consideration. Moreover, they reveal that when toddlers interpret the relation between novel words, they go considerably beyond the simple observation that some object pairs are more alike than others. If the children had taken note of only the similarity in the horn–flute and hook–clip pairs, they might have interpreted the labels as interchangeable with respect to the objects. Instead, these toddlers interpreted one word as being more inclusive than the other, despite the fact that there was nothing in the experimenter's presentation directing them to do so. In this way, they evidenced a natural and spontaneous ability to coordinate their appreciation of similarities as well as differences between the objects and to mark this appreciation lexically.

Hierarchical relations. The present findings are striking because they contrast sharply with the view that young children do not have at their command the conceptual and lexical abilities necessary to support the establishment of hierarchical inclusion relations (Inhelder & Piaget, 1964; Vygotsky, 1962). For example, our data reveal in toddlers a nascent ability to coordinate two different overlapping classes with two different extensions. The ability to coordinate overlapping classes in this manner is important for theories of cognitive development because it is a logical prerequisite to establishing truly hierarchical inclusion relations. To conceptualize a hierarchical system, one must be able to recognize that a given object or class of objects is simultaneously a member of multiple overlapping, hierarchi-

cally related classes. This in turn requires the ability to coordinate overlapping classes. Children in Experiment 2 spontaneously coordinated overlapping extensions of Words A1 and A2. Hence, their performance challenges one conceptually based explanation for children's presumed difficulty in establishing hierarchies.

These data also challenge a language-based explanation for children's presumed difficulty in establishing hierarchies. According to the principle of mutual exclusivity, children assume early in the process of lexical development that words mark mutually exclusive, nonoverlapping classes of objects (Markman, 1989; Markman & Wachtel, 1988). Markman has argued that although children's adherence to this principle facilitates word learning, it actually impedes the establishment of hierarchical systems. This is because in a hierarchical system, labels are not mutually exclusive; any given object or class of objects can necessarily be described by a series of overlapping, hierarchically related labels (e.g., *terrier, dog,* and *mammal*).

The present data are in direct conflict with the position that children use the principle of mutual exclusivity as a default assumption in ascribing meaning to new words. (The data are, however, consistent with Clark's, 1987, principle of contrast, which makes the more moderate claim that no two words in a given language share precisely the same meaning.) We see no evidence that children are biased to construe words as referring to nonoverlapping, mutually exclusive classes. Instead, they readily labeled Object A2 with both Words A1 and A2. Their performance, particularly in production, illustrates how freely and fluently children violate the principle of mutual exclusivity (for other violations of this principle, see Au & Glusman, 1990; Gathercole, 1987; Merriman & Bowman, 1989; Waxman & Hatch, 1992).

This finding also raises an important methodological issue. Most of the published work supporting the principle of mutual exclusivity has been based primarily on comprehension as a measure. However, as we noted earlier, a child who exhibits a mutually exclusive pattern in comprehension may indeed have interpreted the two target words as bearing an inclusion relation. To illustrate this point, we described the hypothetical case of an individual who knows only two terms to refer to cows: *cow* and *guernsey.* This individual, like a child in our experiment, may use the term *cow* productively to refer to both a hereford and a guernsey, while restricting the word *guernsey* to the latter. Yet in comprehension, when that individual is asked to select a cow, she may select only the hereford, on the basis of the (Gricean) assumption that if the speaker had intended to refer to the guernsey, he or she would have said so. Thus, what appears to be a mutually exclusive pattern in comprehension is, in fact, also consistent with an inclusion pattern of interpretation. Thus, the present data signal the importance of examining both comprehension and production when evaluating children's interpretation of the relations among novel words. (Other differences between children's performance in production and comprehension are addressed later in this discussion.)

Our findings advance the arguments made in other published work concerning early lexical and conceptual development. For example, Merriman and Bowman (1989) failed to uncover evidence of the principle of mutual exclusivity until later in the preschool years. Other researchers (Blewitt, 1989; Clark et al., 1985; Gelman, Wilcox, & Clark, 1989; Shipley & Kuhn, 1983; Smith, 1979; Taylor & Gelman, 1989; Waxman, 1990) have noted that children appear to have begun to establish hierarchical systems. For example, Banigan and Mervis (1988) reported that toddlers often interpret novel words (e.g., *goggles*) as referring to subcategories of existing familiar basic-level categories (*glasses*). Merriman (1986) has made the related claim that toddlers tend to interpret novel words as referring to subcategories of newly established basic-level classes.

Taylor and Gelman (1989) also provided preliminary evidence for lexical hierarchies in 2-year-old children. In a series of experiments using a forced-choice method, Taylor and Gelman introduced toddlers to novel nouns (e.g., *fep*) for familiar objects with familiar basic-level labels (e.g., *dog*). They reported that children revealed a strong inclination to interpret the novel word as referring to a subordinate of the familiar basic-level category. Our data are entirely consistent with this important finding and augment it in several ways. First, notice that in Taylor and Gelman's task, children were faced with the puzzle of understanding how the same object (e.g., a particular dog) could sensibly be described by two labels (e.g., *dog* and *fep*). However, in our experiment, children encountered no such puzzle, for they heard two different words applied to two different objects. Thus, children could easily have interpreted the words as being mutually exclusive. Indeed, this was the strategy adopted by children who were introduced to the whisk and the tongs. Although the children who were introduced to the flute and horn or hook and clip could have done the same, they instead interpreted the words as referring to overlapping, but distinct, sets.

Although the empirical foci and procedures adopted by the aforementioned researchers differ considerably from our own, the patterns adopted by the children are strikingly similar. These results are significant because they provide strong and independent support for the claim that children as young as 2 years of age have the requisite conceptual and lexical abilities to set up hierarchical relations. Taken together, these reports call into serious question the assumption that children have difficulty constructing hierarchies.

Nonetheless, some caution is in order in interpreting this body of research. First, it is important to note that although the patterns of data reported here and elsewhere are consistent with a hierarchical interpretation of meaning (e.g., *dog* and *terrier*), they are also consistent with other pairs of meaning that overlap, but in a nonhierarchical fashion (e.g., *former president* and *born in Massachusetts*). The assertion that children establish truly hierarchical relations requires one of two types of evidence: (a) logical evidence, which entails a commitment to logical inclusion statements such as "All flutes are horns," or (b) empirical evidence, which entails an exhaustive delineation of the full range of items that are considered to be members of a given class and subclass. Children's linguistic and metacognitive difficulties in providing the first type of evidence have been well documented (Smith, 1979). For empirical evidence of a hierarchical relation, it must be shown that all things that the child includes in a given subset are also necessarily members of the more inclusive set. For example, it must be shown that all goggles are also considered to be glasses (Banigan & Mervis, 1988), all terriers are necessarily thought to be dogs (Taylor & Gelman, 1989), or all flutes are also considered to be horns (Experiment 2). Such exhaustive empirical data are difficult to

obtain, particularly in children as young as 25 months. Indeed, some have argued that exhaustive evidence is, in principle, impossible to obtain (i.e., Hume, 1748/1974). Unfortunately, in the absence of such evidence, it is difficult to distinguish clearly between overlapping patterns of interpretation that are hierarchical and those that are not. Thus, the data published to date offer strong support for the claim that young children interpret novel words in a manner that is consistent with, but not necessarily exclusive to, a hierarchical interpretation.

Even if it were possible to show that children were indeed establishing truly hierarchical systems, it is important to bear in mind that the rudimentary hierarchical systems of very young children will be considerably less elaborate and less accessible for logical reasoning than are those assembled by older children and adults. Moreover, it will be important in future work to determine more precisely the structure, representation, and inductive power of these early classification systems and to chart more completely the children's interpretation of each novel word.

Extension and overextension. The present data also bear on other fundamental issues in language development, conceptual development, and the interface between them. For instance, children's systematic extension of Word A1 to refer to Object A2 calls to mind the phenomenon of overextension and may provide some insight into the controversy concerning the source of overextension early in language development. Our data are consistent with the observation that overextensions are more likely to occur in production than in comprehension. Children's tendency to generalize Word A1 to include Object A2 was much more prevalent in production than in comprehension. Several researchers have interpreted this type of data as evidence that overextensions are a consequence of children's limited productive vocabularies (e.g., Bloom, 1973; Clark, 1978; Fremgen & Fay, 1980; Kuczaj, 1982). However, in Experiment 2, children's extensions of Word A1 could *not* be attributed to a limitation in the lexicon, for our subjects clearly did have productive command of both Words A1 and A2 (see Tables 3 and 4). Instead, their generalization of Word A1 appears to reflect the power of their conceptual systems rather than any deficiencies in their lexicons.

Production versus comprehension. The present results are also relevant to comparisons of production and comprehension in early language development. Recall that although the patterns observed in comprehension were consistent with those observed in production, children were generally more precise on the latter measure. There are two possible explanations for this difference. First, the children enjoyed the informal aspects of the play sessions, during which, unbeknownst to them, the experimenter modeled each target word a prescribed number of times. They were less enthralled (but still compliant) with responding to the experimenter's probes for production and comprehension. They quickly learned to respond promptly to the probes so that she would cease probing and they could resume playing freely. To respond to a production probe, a child simply had to focus attention on the object in the experiment's hand and label it. The fact that children labeled with such accuracy is testimony to their proficiency in word learning. Interestingly enough, responding to the comprehension probes required more sustained attention. Recall that on comprehension probes, children were asked, "Where is the [target object]?" To respond, they had to search for the requested object among the others in the array. We noticed that children sometimes searched for a moment and then, if they could not quickly locate the appropriate object, would simply hold up any object.

Second, there were actually more opportunities to observe the children's production than their comprehension. Most children produced the target terms spontaneously during the visits. Yet there was no natural opportunity for them to spontaneously demonstrate their comprehension. Only during the prescribed comprehension probes (two probes per target word per visit) did the experimenter use the target word without simultaneously indicating the target object. As a result, our data set includes both spontaneous and elicited productions of the target words, but only elicited comprehension. This methodological difference leaves open the possibility that if children had been given more opportunities to demonstrate their comprehension of the target terms, they might have more frequently extended them. Their patterns in comprehension might then have mirrored more exactly the patterns in production.

Nonetheless, the differences obtained on our production and comprehension measures cannot be reduced entirely to methodological concerns, because similar discrepancies have been reported by other researchers working from several different vantage points. See, for example, Huttenlocher and Smiley's (1987) argument concerning the limitations of comprehension data compared with production data.

Accounting for the asymmetrical pattern of extension. Despite the relatively small number of children who participated in our experiment, the data are remarkably consistent in two important ways: (a) Each child's pattern of interpretation was consistent across the repeated sessions, and (b) each child's pattern of interpretation closely resembled that of other children learning words for the same target object pair. However, it is unclear why Word A1, and not Word A2, was extended by all 8 children who learned the hook–clip and horn–flute pairs. One possible explanation involves the interpretive biases of young word learners and the strategies adopted by their parents. Several researchers have demonstrated children's tendency to interpret the first word applied to a novel object as referring to the basic-level kind and parents' reciprocal tendency to label objects initially with basic-level terms (Callanan, 1989; Hall & Waxman, 1991; Mervis, 1986; Shipley et al., 1983). Findings like these suggest that children in the present experiment may have interpreted the A1 term at the basic level, and then interpreted the A2 term as referring to a more restricted set (Blewitt, 1989; Merriman, 1986; Smith, 1979; Taylor & Gelman, 1989; Waxman, 1990). Another possibility is that the circumstances under which we introduced the related target terms influenced the children's interpretations. Recall, for example, that Words A1 and A2 were not introduced in identical fashions. Word A1 was modeled on the first visit but not on the second and third; Words A2 and B were modeled on the second and third visits but not on the first. Although it is unclear precisely how these circumstances could have led to the observed asymmetry, it may be worthwhile in future work to pursue this possibility. For example, if Words A1 and A2 had been modeled on the same days, children might have been more inclined to interpret them as mutually exclusive. Finally, although we have extensive and detailed data from each child, it is possible that the asymmetry was a consequence of the small number of subjects involved.

Additional research will help to pinpoint the source of the consistent asymmetry in production observed in this study.

Evolution of word meaning. Finally, we were surprised to discover that children's interpretations of the target words remained stable over the 2-week period during which we repeatedly visited each child. This finding is reminiscent of the argument that many object words have proper extensions from the start (Bowerman, 1978; Gruendel, 1977; Huttenlocher & Smiley, 1987; Macnamara, 1982; Nelson & Bonvillian, 1973; Rescorla, 1980). However, it is important to recall that we gave children no corrective feedback regarding their extensions of the novel words. Therefore, we were essentially studying children's word-learning strategies when they were left to their own devices. This makes it all the more impressive that they went beyond the input to establish overlapping meanings. However, we suspect that if we had provided feedback, we would have seen more modification in meaning over the course of the visits. For example, if the experimenter had corrected children when they applied Word A1 (e.g., *hook*) to Object A2 (e.g., the clip), the children would have learned to apply the two words as mutually exclusive. Therefore, the stability observed in the present experiment does not in itself directly challenge the argument that word meanings evolve through a process of social convergence and coordinated feedback between adults and toddlers (Adams & Bullock, 1986; Callanan, 1985; Clark, 1973; Fischer & Bullock, 1984; Mervis & Mervis, 1982; Vygotsky, 1962).

Conclusion

In the past, there has been very little experimental research of this extensive nature with 2-year-olds, in part because they are difficult and perplexing subjects. Toddlers are too young to perform consistently on most classification and language tasks (but see Crain & Thornton, 1989; Golinkoff, Hirsh-Pasek, Cauley, & Gordon, 1987; Sugarman, 1981; and Waxman & Kosowski, 1991 for innovative and promising methods that have been developed for children at this age). Yet they are too old (and often too active) to perform reliably on most measures designed for infants. Perhaps our method, which weds naturalistic observations to controlled experimental methods, will continue to provide insight into convergences between conceptual and lexical acquisition.

The evidence thus far bolsters the argument that 2-year-old children have at their command the conceptual and lexical abilities that are fundamental to the establishment of hierarchical systems. Toddlers adopted clear strategies and went beyond the experimenter's input to work out the relations among novel labels. Strategies like these are an essential element in early lexical development and make possible the rapid acquisition of novel words and the relations among them.

References

Adams, A. K., & Bullock, D. (1986). Apprenticeship in word use: Social convergence processes in learning categorically related nouns. In S. A. Kuczaj II & M. D. Barrett (Eds.), *The development of word meaning: Progress in cognitive development research* (pp. 155–197). New York: Springer-Verlag.

Au, T. K. (1990). Children's use of information in word learning. *Journal of Child Language, 17,* 393–416.

Au, T. K., & Glusman, M. (1990). The principle of mutual exclusivity in word learning: To honor or not to honor? *Child Development, 61,* 1474–1490.

Banigan, R. L., & Mervis, C. B. (1988). Role of adult input in young children's category evolution: II. An experimental study. *Journal of Child Language, 15,* 493–504.

Blewitt, P. (1989). Categorical hierarchies: Levels of knowledge and skill. *Genetic Epistemologist, 17*(2), 21–27.

Bloom, L. (1973). *One word at a time: The use of single word utterances before syntax.* The Hague, The Netherlands: Mouton.

Bowerman, M. (1978). Systematizing semantic knowledge: Changes over time in the child's organization of word meaning. *Child Development, 49,* 977–987.

Callanan, M. A. (1985). How parents label objects for young children: The role of input in the acquisition of category hierarchies. *Child Development, 56,* 508–523.

Callanan, M. A. (1989). Development of object categories and inclusion relations: Preschoolers' hypotheses about word meanings. *Developmental Psychology, 25,* 207–216.

Carey, S. (1978). The child as word learner. In M. Halle, J. Bresnan, & G. Miller (Eds.), *Linguistic theory and psychological reality* (pp. 264–293). Cambridge, MA: MIT Press.

Carey, S., & Bartlett, E. (1978). Acquiring a single new word. *Papers and Reports on Child Language Development, 15,* 17–29.

Clark, E. V. (1973). What's in a word: On the child's acquisition of semantics in his first language. In T. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 65–110). San Diego, CA: Academic Press.

Clark, E. V. (1978). Strategies for communicating. *Child Development, 49,* 953–959.

Clark, E. V. (1987). The principle of contrast: A constraint on language acquisition. In B. McWhinney (Ed.), *Mechanisms of language acquisition: The 20th Annual Carnegie Symposium on Cognition* (pp. 1–33). Hillsdale, NJ: Erlbaum.

Clark, E. V., Gelman, S. A., & Lane, N. M. (1985). Compound nouns and category structure in young children. *Child Development, 56,* 84–94.

Crain, S., & Thornton, R. (1989, October). *Acquisition of *that-trace effects.* Paper presented at the Boston Child Language Conference, Boston, MA.

Dromi, E. (1987). *Early lexical development.* Cambridge, England: Cambridge University Press.

Fischer, K. W., & Bullock, D. (1984). Cognitive development in school-age children: Conclusions and new directions. In W. A. Collins (Ed.), *Development during middle childhood* (pp. 70–146). Washington, DC: National Academy Press.

Fremgen, A., & Fay, D. (1980). Overextensions in production and comprehension: A methodological clarification. *Journal of Child Language, 7,* 205–211.

Gathercole, V. C. (1987). The contrastive hypothesis for the acquisition of word meaning: A reconsideration of the theory. *Journal of Child Language, 14,* 493–531.

Gelman, S. A., & Taylor, M. (1984). How two-year-old children interpret proper and common names for unfamiliar objects. *Child Development, 55,* 1535–1540.

Gelman, S. A., Wilcox, S. A., & Clark, E. V. (1989). Conceptual and lexical hierarchies in young children. *Cognitive Development, 4,* 309–326.

Gentner, D. (1978). A study of early word meaning using artificial objects: What looks like a jiggy but acts like a zimbo? *Papers and Reports on Child Language Development, 15,* 1–6.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman, & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development* (pp. 225–277). Cambridge, England: Cambridge University Press.

Golinkoff, R. M., Hirsh-Pasek, K., Cauley, K., & Gordon, L. (1987). The eyes have it: Lexical and syntactic comprehension in a new paradigm. *Journal of Child Language, 14,* 23–45.

Gruendel, J. M. (1977). Referential extension in early language development. *Child Development, 48,* 1567–1576.

Hall, D. G. (1990, April). *The restricted–unrestricted distinction in lexical acquisition.* Poster presented at the Biennial Meeting of the Society for Research in Child Development, Seattle, WA.

Hall, D. G. (1991a). Acquiring proper nouns for familiar and unfamiliar animate objects: Two-year-olds' word-learning biases. *Child Development, 62,* 1142–1154.

Hall, D. G. (1991b). *Acquiring count nouns: The unrestricted-restricted distinction.* Unpublished doctoral dissertation, Harvard University, Cambridge, MA.

Hall, D. G., & Waxman, S. R. (1991). *On the acquisition of restricted and unrestricted nouns.* Manuscript submitted for publication.

Hall, D. G., Waxman, S. R., & Hurwitz, W. M. (1991). *How 2- and 4-year-old children interpret adjectives and count nouns.* Manuscript submitted for publication.

Heibeck, T. H., & Markman, E. M. (1987). Word learning in children: An examination of fast mapping. *Child Development, 58,* 1021–1034.

Hume, D. (1974). An enquiry concerning human understanding. In *The empiricists* (pp. 307–430). Garden City, NY: Anchor Books. (Original work published 1748)

Huttenlocher, J., & Smiley, P. (1987). Early word meanings: The case of object names. *Cognitive Psychology, 19,* 63–89.

Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child.* New York: Newton.

Katz, N., Baker, E., & Macnamara, J. (1974). What's in a name? A study of how children learn common and proper names. *Child Development, 45,* 469–473.

Kuczaj, S. A. II (1982). Young children's overextensions of object words in comprehension and/or production: Support for a prototype theory of early object word meaning. *First Language, 3,* 93–105.

Landau, B., & Gleitman, L. R. (1985). *Language and experience: Evidence from the blind child.* Cambridge, MA: Harvard University Press.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3,* 299–321.

Macnamara, J. (1982). *Names for things.* Cambridge, MA: MIT Press.

Markman, E. M. (1989). *Categorization and naming in children.* Cambridge, MA: MIT Press.

Markman, E. M., & Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology, 16,* 1–27.

Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology, 20,* 121–157.

Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20,* 158–190.

Merriman, W. E. (1986). Some reasons for the occurrence and eventual correction of children's naming errors. *Child Development, 57,* 942–952.

Merriman, W. E., & Bowman, L. L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development, 54*(3–4, Serial No. 220).

Mervis, C. B. (1986). Child-basic categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 201–233). Cambridge, England: Cambridge University Press.

Mervis, C. B., & Canada, K. (1982). On the existence of competence errors in early comprehension: A reply to Fremgen & Fay and Chapman & Thomson. *Journal of Child Language, 10,* 431–440.

Mervis, C. B., & Mervis, C. A. (1982). Leopards are kitty-cats: Object labeling by mothers for their thirteen-months olds. *Child Development, 53,* 267–273.

Nelson, K. (1974). Concept, word and sentence: Interrelations in acquisition and development. *Psychological Review, 81,* 267–285.

Nelson, K. E., & Bonvillian, J. D. (1973). Concepts and words in the 18-month-old: Acquiring concept names under controlled conditions. *Cognition, 2,* 435–450.

Ninio, A. (1980). Picture-book reading in mother-infant dyads belonging to two subgroups in Israel. *Child Development, 51,* 587–590.

Pinker, S. (1979). Formal models of language learning. *Cognition, 7,* 217–283.

Quine, W. V. (1960). *Word and object.* Cambridge, MA: MIT Press.

Rescorla, L. (1980). Overextension in early language development. *Journal of Child Language, 7,* 321–335.

Shipley, E. F., & Kuhn, I. F. (1983). A constraint on comparisons: Equally detailed alternatives. *Journal of Experimental Child Psychology, 35,* 195–222.

Shipley, E. F., Kuhn, I. F., & Madden, E. C. (1983). Mothers' use of superordinate category terms. *Journal of Child Language, 10,* 571–588.

Smith, C. (1979). Children's understanding of natural language hierarchies. *Journal of Experimental Child Psychology, 27,* 437–458.

Soja, N. N., Carey, S., & Spelke, E. S. (1991). Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition, 38,* 179–277.

Sugarman, S. (1981). The cognitive basis of classification in very young children: An analysis of object ordering trends. *Child Development, 52,* 1172–1178.

Taylor, M., & Gelman, S. A. (1988). Adjectives and nouns: Children's strategies for learning new words. *Child Development, 59,* 411–419.

Taylor, M., & Gelman, S. A. (1989). Incorporating new words into the lexicon: Preliminary evidence for language hierarchies in two-year-old children. *Child Development, 60,* 625–636.

Tomasello, M., Mannle, S., & Werdenschlag, L. (1988). The effect of previously learned words on the child's acquisition of words for similar referents. *Journal of Child Language, 15,* 505–516.

Tomikawa, S. A., & Dodd, D. H. (1980). Early word meaning: Perceptually based or functionally based. *Child Development, 51,* 1103–1109.

Tversky, B. (1985). Development of taxonomic organization of named and pictured categories. *Developmental Psychology, 21,* 1111–1119.

Vygotsky, L. (1962). *Thought and language.* Cambridge, MA: MIT Press.

Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development, 5,* 123–150.

Waxman, S. R. (1991). Convergences between semantic and conceptual organization in the preschool years. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on language and cognition: Interrelations in development* (pp. 107–145). Cambridge, England: Cambridge University Press.

Waxman, S. R., & Hatch, T. (1992). Beyond the basics: Preschool children label objects flexibly at multiple hierarchical levels. *Journal of Child Language, 19,* 153–166.

Waxman, S. R., & Kosowski, T. D. (1990). Nouns mark category relations: Toddlers' and preschoolers' word-learning bases. *Child Development, 61,* 1461–1473.

Wittgenstein, L. (1953). *Philosophical investigations.* New York: Macmillan.

Received March 26, 1991
Revision received February 14, 1992
Accepted February 24, 1992 ■