The Changing Character of the Mental Lexicon:

An Information-based Account of Early Word Learning

Jesse Snedeker*

Harvard University

Michael Brent

Washington University

Lila Gleitman

University of Pennsylvania

Running Head: INPUT FOR EARLY WORD LEARNING

* Address correspondence to: Jesse Snedeker, Department of Psychology, 1138 William James Hall, Harvard University, 33 Kirkland Street, Cambridge MA, 02138; snedeker@wjh.harvard.edu.

Abstract

This work explores two phenomena in early lexical development. First, children's initial vocabularies contain a disproportionate number of nouns. Second, during the first year of word learning the overrepresentation of nouns increases steadily and an object-kind bias appears in word learning tasks. In these experiments, we explore whether the information available to novice learners can account for this developmental pattern by having adult subjects attempt to identify words from the situations in which they occur in infant-directed speech. Experiment 1 demonstrates that the input favors noun targets, particularly object labels. Experiment 2 shows that this input supports accurate self-assessments, which could tip off learners to the advantage of noun responses. Experiment 3 verifies that an object-kind bias leads to more successful word learning. We conclude with an account of how an object-kind bias could be learned from input with these characteristics.

Vocabulary acquisition poses a classic "poverty of the stimulus" problem. Amongst all the things that a new word could mean, how is the learner to choose? After all, a variety of different things can be – and are – uttered under the same external circumstances; conversely, a single utterance can be applicable to many different extralinguistic circumstances (cf., Chomsky, 1959). Yet even the most ordinary child has acquired a vocabulary of about 10,000 words during the first five years of life, and this vocabulary continues to expand throughout life.

In the present paper, we investigate word learning by focusing experimental attention on a typical property of the novice lexicon; namely, its concreteness (Gentner, 1978; Smith, Jones, and Landau, 1996). Early vocabularies are concrete in the sense that they usually label the kinds of things – the dogs and the spoons and the scooters – that populate our perceptual world at its "middle-sized" grain. Crosslinguistically, labels for such things tend to be nouns and so, artifactually, early vocabularies are "noun dominant." As we will try to demonstrate, a concrete and noun-dominant vocabulary will arise from the information structure of the input that is available to novice language learners, above and beyond any further constraints that may be imposed by their conceptual biases or limitations. To distinguish between information-based accounts and accounts invoking conceptual change, we hold conceptual sophistication constant by studying word learning in adults, and then we manipulate the information made available to these adult learners (for an earlier experimental series using these methods, Gillette, Gleitman, Gleitman, and Lederer, 1999).

Specifically, the experiments are directed to the following questions: If adults are asked to guess the meanings of words from observing their extralinguistic contexts, will they succeed on items as a function of their concreteness (hence, mimicking the "noun dominance" effect in young children)? Our second question pertains to the source of the noun dominance effect: Does information that is available in the child's input support accurate self-assessments of leaning (implicit or explicit) and could these assessments support the growth of a bias to acquire nouns in preference to items from other lexical classes? Third, if – by experimental artifice – we implant this bias to acquire nouns in our adult subjects, will they improve at learning words from their real-world contingencies? Before presenting the experiments, we want to provide some further detail about early child vocabularies and the kinds of accounts that have been given to explain them.

The Categorial Structure of Novice Vocabularies

These experiments are designed to mirror, and at best explain, the following specific properties of early word learning (see Figure 1). Children's early (productive and receptive) vocabularies are dominated by nouns that refer to people, animals, and small moveable objects (like toys). Even though adults speak to young children in full sentences, complete with verbs and function words, these elements are massively underrepresented in children's early vocabularies. This is true not only in languages like English (Brown, 1973; Bates, Dale & Thal, 1995) in which subjects and objects are mandatory and verbs are often buried in the middle of an utterance, but also in languages such as Italian (Caselli et al., 1995) and Korean (Au, Dapretto & Song, 1994) where verbs often occur in more salient positions, and in Hebrew, where verbs have more reliable morphological marking than nouns (Maital et al., 2000).[1]

Obviously, then, early vocabularies are not just straightforward reflexes of input characteristics. For some reason, nouns are easiest. In fact, if all words and word types were equally difficult to learn, we would predict that many verbs would be learned before any nouns. This is because, both in adult-to-adult and infant directed speech, the most frequent verbs are higher in token-frequency than the most frequent nouns (Miller, 1951; Johansson & Hofland, 1989; Gentner, 1981; Sandhofer, Smith, & Luo, 2000).

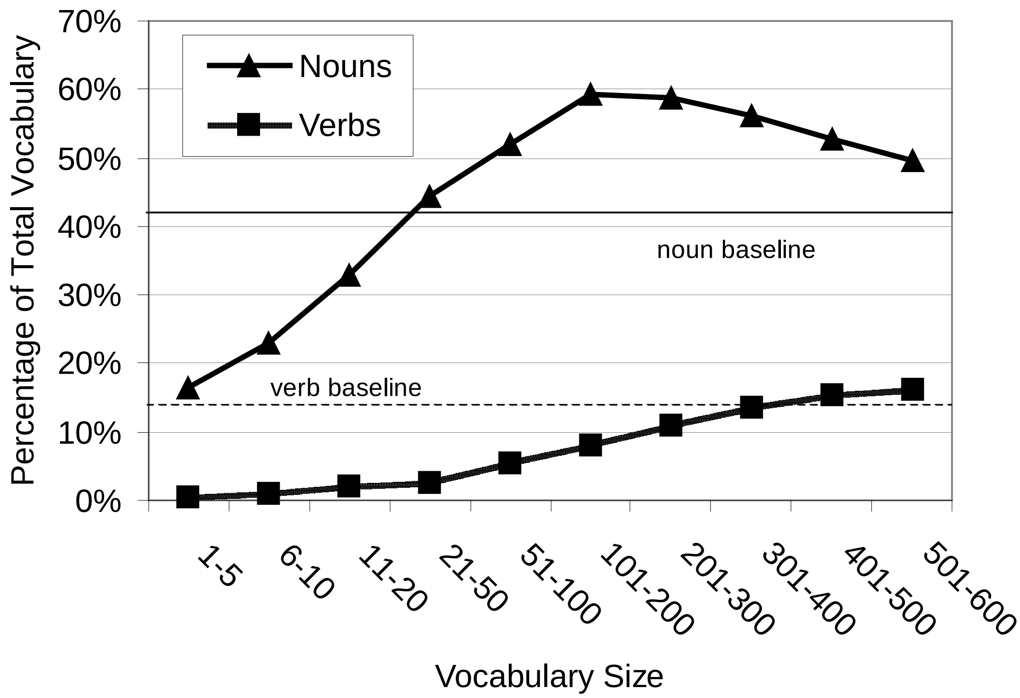

Figure 1. Composition of the production vocabulary of English learners as a function of vocabulary size (adapted from Caselli, Bates, Casadio, Fenson, Fenson, Sander & Weir, 1995 and Caselli, Casadio, & Bates, 1999). The noun and verb baselines indicate the percentage of words of each category on the MCDI toddler checklist.

To say "nouns are easiest" is, of course, a description in want of an explanation. But even as a description, the noun-first perspective fails to capture some of the detailed changes that occur in the young child's vocabulary composition. Earliest word learning (approximately, the first 20 words, usually in the period from 10-15 months) either does not show noun dominance, or shows it in only a weak form (Bates, Dale & Thal, 1995; Caselli et al., 1995). During this first period, word learning is slow and errorful. In fact, it is often hard to imagine just what the infant thinks his pathetic few words really mean (Dromi, 1987). Thereafter, through about the first 200 words (usually the period from approximately 16 – 24 months), noun dominance increases steadily. Nouns overpopulate the child vocabulary through approximately the next year of life (relative to noun frequency in parental-input speech). As Figure 1 also shows, the noun dominance effect then gradually tails off until the noun frequency in child vocabularies approximates that of input (adult) speech. As we will discuss further, the decline of noun dominance (or rather the rise of other word types in the child's vocabulary) is associated with the creation of more sophisticated linguistic representations of input speech. We will argue that these linguistic constraints act to refine and limit the conjectures made available by extralinguistic observation, and so support the acquisition of words from every lexical class.

Four Accounts of the Noun Dominance Effect

We now sketch four accounts of early noun dominance and its change over time. The first two accounts explain the changing categorial structure of the lexicon by alluding to properties of the learners themselves. The third and fourth accounts approach the same phenomena by pointing to properties of the input. The last of these is the one that will be put to experimental review in this paper.

Hypothesis 1: Nouns Encode Simple Concepts

The input-output disparity in early word learning is often attributed to the conceptual limitations of young children (Macnamara, 1972; Nelson, 1974; O'Grady, 1987). According to this view, the novice word learner can represent the world only in terms of "simple" concepts (such as those specifying individual persons and categories of perceptually similar objects). Able to entertain only such concepts, they can learn just the words that encode these simple concepts. And these are mostly nouns. As the child develops she becomes capable of representing more complex concepts (such as relations, events, and properties) and can now map these concepts to the words that express them (which are mostly verbs, adjectives, and prepositions). The view that conceptual change accounts for stages in lexical learning has been so influential that many researchers have used age-related changes in vocabulary composition as indices of conceptual change (e.g., Huttenlocher, Smiley, & Charney, 1983; Gopnik & Meltzoff, 1997; Smiley & Huttenlocher, 1995).

Hypothesis 2: Unlearned Biases in Lexical Learning

Other investigators posit that the noun dominance effect derives from learning mechanisms specific to the lexicon. It is not (or not relevantly) that learners are conceptually primitive in some way. Rather, they are biased to pair words with particular types of meanings and to avoid others, even through they are capable of using these lexically disfavored concepts in other domains. An important instance is the Whole-Object Constraint (Markman & Wachtel, 1988): when toddlers hear a novel word that could potentially refer either to a novel object or one of its properties or parts, they will generally extend that label to objects of the same kind. Simply hearing a linguistic label ("This is a blicket") can redirect a child's attention from a perceptually salient feature (such as motion) onto a novel object (Woodward, 1992). On the basis of such evidence, Markman proposed that children constrain their hypotheses about potential word meanings, assuming that a novel label will typically refer an entire object. Biases of this kind can provide an account for the noun-dominance effect in early vocabularies, even if children's conceptual repertoire is assumed to be quite similar to that of adults (but see also Markson & Bloom, 1997). Under this view, the origin of the object bias is independent of the child's experiences with word learning. It is either an innate mechanism specific to building the lexicon (e.g., Clark's principle of contrast (1987)) or a product of prior learning in other domains (e.g., in perception or prelinguistic cognition, Markman, 1992).

Hypothesis 3: The Information Structure of the Input

Gillette, Gleitman, Gleitman & Lederer (1999, henceforth GGGL) offered a different kind of explanation for children's limited lexical range, one that has little or nothing to do with early conceptual limitations. They proposed that the concreteness of early vocabulary and its close associate, noun dominance, arises from the fact that early word learning relies almost exclusively on what can be gleaned from observation of the real-world contingencies for word use. Words whose referential correlates are unobservable can't be learned by observation of the world. When only situational contexts are available word learning must rely on surface features of the ambient external world, even if the learner realizes that the meanings themselves are quite abstract. The outcome is a vocabulary heavily dominated by words with concrete real-world correlates; hence overloaded with nouns.

In contrast, according to GGGL, adults and even sophisticated three-year olds have multiple convergent sources of evidence that they use for determining word meanings. For instance, more experienced learners can make use of linguistic co-occurrence: knowledge that some word co-occurred with the other words "chef" and "bake" is a soft cue that it might mean "cake". But as a precondition for using such a word co-occurrence procedure, the learner must have built up a lexicon that stores probabilistic co-occurrence relations (Siskind, 1996). Thus this machinery cannot be used in the very earliest stages of word learning. Similarly, adults and preschoolers can use the morphology and syntactic structures in which word appear as cues to meaning (e.g., Bloom & Keleman, 1995; Brown, 1957; Gleitman, 1990). But again these kinds of cue are unavailable to the truly novice word learner, leaving him at the (often, not-so-tender) mercy of real-world observation.

In sum, GGGL's position is that noun dominance in early vocabulary is a consequence of the fact that only extralinguistic evidence ("word-to-world pairing") is available to youngest learners. It takes time for them to build up knowledge of the correlational structure of word use and of the morphology and structure of the exposure language so that these evidentiary sources can be added to the evidence available from observation of the outside world.

Hypothesis 4: Look Under the Lamppost

We propose here a view of lexical learning that turns on the learners' ability and inclination to focus on strategies that have worked in the past. This hypothesis presupposes and incorporates GGGL's explanation of noun dominance: when external observation provides the only or chief cue for word learning, the vocabulary will be concrete (thus noun-dominated) regardless of the conceptual repertoire of the learner. But this account without further amplification leaves two (related) properties of early word learning unexplained. First, the very earliest vocabulary learning appears to be categorially catholic (and almost uniformly unsuccessful, see again Figure 1). Second, the tendency to use surface shape as criterial for labeling is weak or nonexistent at the earliest stages of word learning, and emerges gradually (Landau, 1994; Samuelson & Smith, 1999; Samuelson, 2000). Why isn't the novice word learner biased toward concreteness from the very start of word learning? And what makes this bias grow?

We posit now (following Smith, 1995) that the bias to map words to object categories is absent in novice word learners and develops only through experience with word learning. We are not suggesting that novices are completely open-minded but merely that they are unbiased within the range of humanly natural concepts, whatever this range may turn out to be. Thus we exclude hypothetical learners who are inclined to entertain such concepts as grue and bleen or chimera composed of undetached rabbit-parts (cf., Goodman, 1966; Quine, 1960). We simply posit that novice word learners are open to the possibility that the word that they are hearing could come from any of the semantic categories available to natural languages.

As GGGL showed, even for sophisticated adults the extralinguistic environment provides strong cues only for the most concrete items. Yet this extralinguistic environment is the sole source of evidence early on. The learner rapidly discovers which kinds of word-meaning conjectures are sustained by the input data. Flushed with success on object labels, he progressively tunes his word learning in their direction. Thus by the present hypothesis concreteness effects feed upon themselves, accounting for the rise of noun dominance as the word-learning engine moves into high gear (Figure 1). Under hypothesis 3, as described by GGGL, the noun-dominance effect is the passive, adventitious outcome of a procedure that tries to map words to meanings of various types but succeeds only when the proposed meaning is concrete. Under the present hypothesis, the learner is much more active. He revises his learning procedure as a consequence of what has worked before. The novice looks for the word meanings under the lamppost, because that's where the light is.

Two more properties of this postulated feedback account are worth mentioning at the outset. The first is that the same mechanism that enhances the acquisition of concrete items should be inhibiting the acquisition of less concrete ones, because the learner will no longer search the abstract corners of his conceptual repertoire for conceptual correlates of word occurrence. The second is that the feedback procedure, just because it is a response to the input conditions, could have different outcomes in different environments —when languages, cultural practice, or even individual parents differ.

The Human Simulation Paradigm

The experiments that follow make use of a paradigm that was developed in GGGL to break the confound between cognitive limitations (as in Hypothesis 1) and insufficient knowledge of language (as in Hypothesis 3). Older children and adults can represent the linguistic input in more sophisticated ways, to be sure, but they also are more cognitively advanced than the usual novice word learner. GGGL disentangled these factors by presenting different information about words to cognitively mature (adult) learners, whose task was to guess the intended words. The target words were common nouns and verbs drawn from transcripts of videotapes of mothers playing with their children.

In the initial condition of GGGL, the subjects were to guess these "mystery words" based on two sources of evidence. First, they were told the syntactic category of the test word and then were shown six short video clips of scenes in which the word had been uttered by one of the mothers to her child. Because the audio had been turned off (and the mystery word itself replaced by a tone), the subject's only recourse was to inspect the videos themselves —what was going on —for their clues to the meanings. Under this forced word-to-world pairing procedure, college students, like toddlers, were able to learn more nouns than verbs. On the final trial, they identified 45% of the noun targets but only 15% of the verbs. In other experimental conditions of GGGL, populations of college students provided with alternative or additional information sources (e.g., co-occurrence information and/or syntactic information about the mystery words) were of course able to learn the verbs as well as the nouns. These experimental results provided support for an information-sensitive view of lexical learning, as in Hypothesis 3. The same cognitively mature subjects show the noun dominance in word learning if they are limited to extralinguistic information, but not if they are provided with co-occurrence and syntactic representations of the input.

Experimental Preview

In the experiments we now present, we again investigated the effects of various information sources on word learning using the GGGL methods but with two crucial differences:

Effects of Categorial Knowledge

In the studies by GGGL, subjects were always told the lexical category of each mystery word before they tried to guess it. In the present study, the subjects are always told to guess "the mystery word" rather than "the mystery verb" or "the mystery noun." The first benefit of this procedural change is to remove a glaring artificiality in the original GGGL experiments, for after all mothers do not ordinarily announce to their infants the lexical category of each word. GGGL had assumed that this prior knowledge would ease the task, but not interact with the experimental conditions. But that remains to be seen: Knowledge of category membership may be more informative for noun learning than for verb learning, or vice versa, and category membership may be more or less informative taken together with scene observation than with, say, syntactic information. Such interactions, as we will demonstrate, have consequences for how we should think about the mechanisms for acquisition of the lexicon.

Effects of Feedback

Could success in the learning task affect the learning procedure itself? We will try to demonstrate that the answer is yes. Specifically, we first test whether extralinguistic observation supports accurate confidence judgments; that is, whether learners can and do generate information about the quality of their hypotheses, information that could allow the word-learning system to monitor its success. Second, we explore whether these confidence judgments could support the development of an object bias by examining the relative confidence in object and non-object responses. As we will show, subjects whose information comes solely from observation of the extralinguistic world are not only better at noun than at verb learning — they know it. They are more confident about their noun guesses, right or wrong, than about their verb guesses. Finally, we test for the hypothesized effects of such a feedback scheme. As we will show experimentally, extralinguistically supported word learning is enhanced for learners in whom we induce the noun bias. In discussion, we will suggest that real child learners, like our experimental subjects, are best understood in much these same terms: Born into a community without specific knowledge of its language, they can succeed at first in acquiring only its most concrete lexical items. Therefore, until they acquire complex linguistic representations they increasingly focus their attention on learning concrete nouns. This learning-induced focus on a subset of the vocabulary efficiently partitions the language learning task: It creates an initial scaffold on which further linguistic developments are built.

Experiment 1: The Informational Structure of Word Learning

Here we ask whether adults, like child learners, identify nouns more accurately than verbs from observing the contexts for words' use, even when they are not given the syntactic category of the target word.

Method

Subjects

The subjects were 84 undergraduates at the University of Pennsylvania who received either partial class credit or a small payment for participation. All subjects were native English speakers.

Stimuli

The videotaped stimuli were adapted from GGGL. They were drawn from 4 hour-long videotaped "play sessions." In each session a mother and her 18-24 month old child explored a bag of toys that the experimenter had brought to the child's home. The toys were chosen with the goal of eliciting both object and action oriented utterances. For example, a cobbler's bench was selected because the experimenters expected that the mothers would both label the object and show the children how to use it. The experimenter kept the camera focused on the child and attempted to include in the shot both the child and everything that the child could see. The videotapes were transcribed and the 24 most common nouns and 24 most common verbs were chosen as targets. The target words were divided among three lists, each containing an equal number of nouns and verbs. The order of presentation of the targets in the first list was randomly generated, and this list was used as a template for the other two lists, so that nouns and verbs appeared in the same positions in all three lists. The order of presentation was reversed for half of the subjects who received each list.

For each target word, we chose 6 instances of its use by the mothers (Figure 2 depicts the video clips that were chosen for the example verb play). Candidate instances were excluded if the referent of the word was visible to the child but off camera, or if the mother's lips were visible and might provide information about the word form. When more than 6 instances of a word met these criteria, instances were selected at random. For each instance, we constructed a video clip that began 30 seconds before the word was uttered and ended 10 seconds after it was uttered. Often the mother said the target again during this period. In these cases, the clip was expanded to include 30 seconds before the first use and 10 seconds after the last and each instance in the clip was counted as one of the six stimuli.[2]

<u>Procedure</u>

Subjects were tested in small groups. Each group was assigned to one of the three stimulus lists and to one of the two presentation orders. They were told that they would be seeing short silent videos of mothers playing with their young children and that their task was to identify a word that the mother had said during that clip. The unknown word was represented by a tone that occurred at the exact point at which the mother had that uttered the word. For each word, subjects saw the set of scenes that included 6 instances of the word. After each tone they tried to identify word based on the scenes that they had seen thus far. After their sixth response, they were given 10 seconds to review all of their responses and make a final effort to identify the word. For each subject group, this procedure was repeated for 16 different words, 8 nouns and 8 verbs. Subjects also knew: that the children were 18-24 months of age; that the words had been chosen for their frequency in speech to young children; that roughly half of the words were nouns and roughly half were verbs[3]

At the end of the study the experimenter inspected the subjects' responses to find lexically ambiguous words and asked the subjects if they had intended the word to be a noun or a verb. If a subject had trouble making the distinction, the experiment presented the word in a noun frame ("a drum") and a verb frame ("drumming") and asked the subject to make a choice.

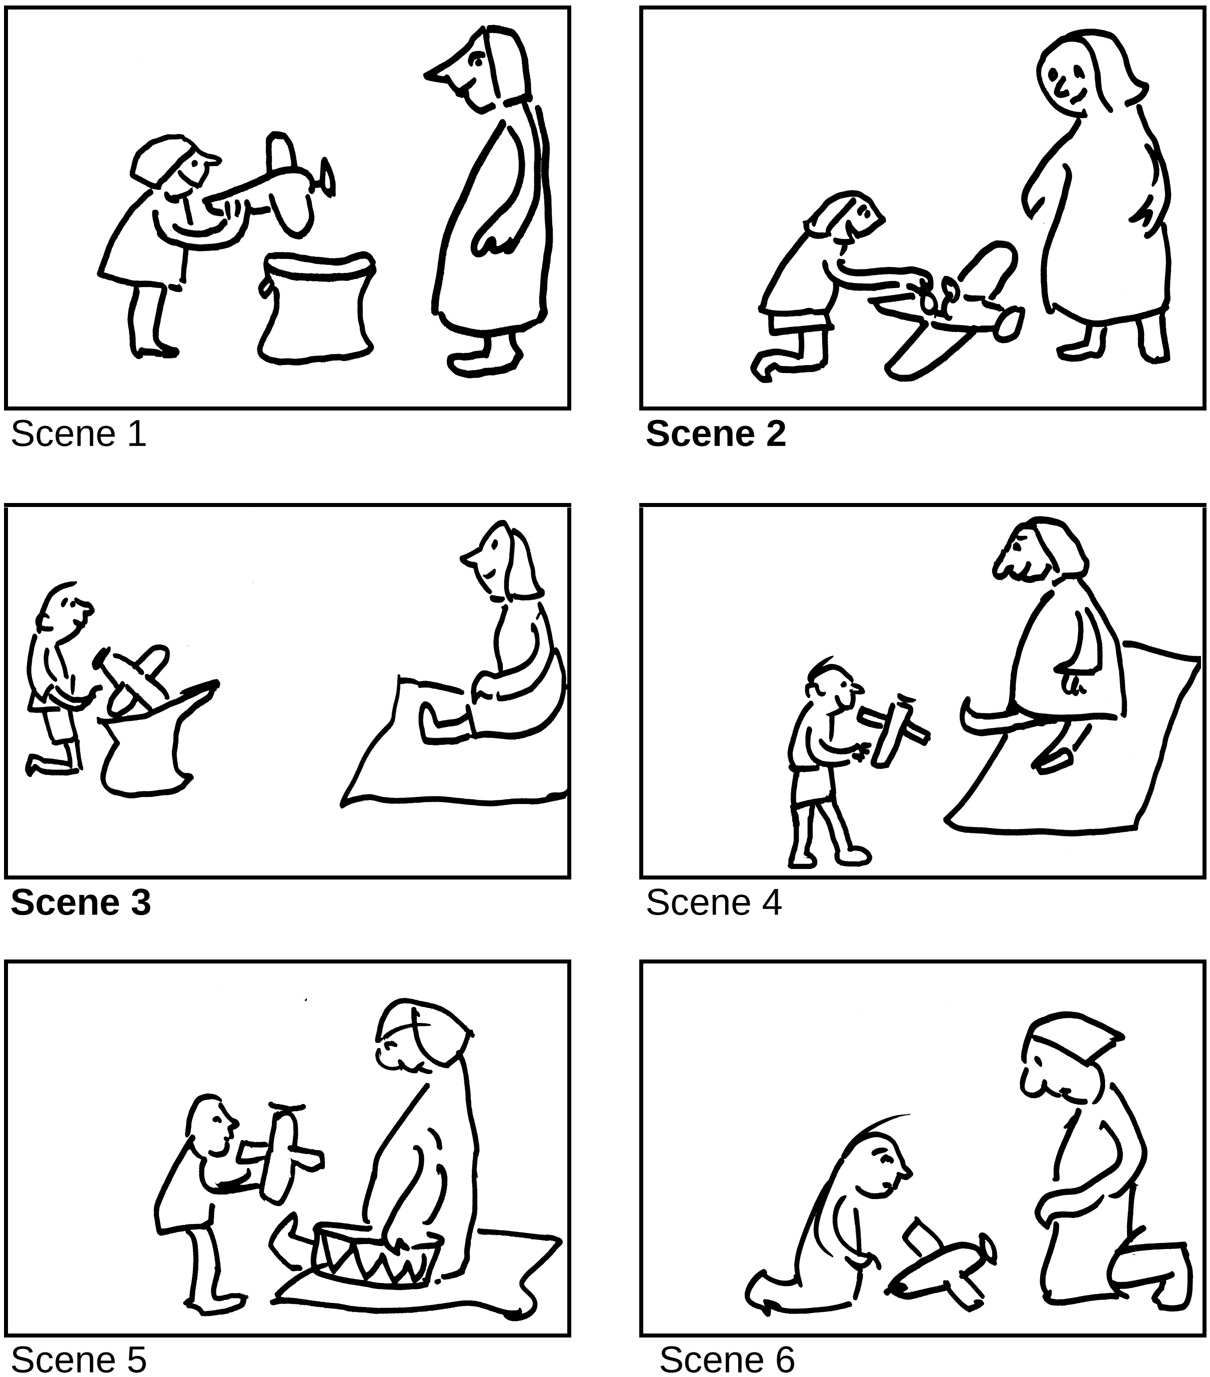

Figure 2. These cartoons approximate the visible information in the 6 videos for the item plane. Each video clip was about 40 seconds long.

Scoring

Responses were scored as correct if they contained the same base morpheme as the target word, even if they differed from the target in number, tense, or voice. If the subject provided the correct word form but indicated that their intended meaning came from a different lexical category than the target, the response was still scored correct.[4] If subjects responded with a phrase rather than a single word, the head of this phrase was taken as their response. No subject did this more than twice.

Results

The Effects of Word Type

Lexical category. Not all words are equally amenable to identification by observation of their real-world contingencies. Recall that GGGL's subjects correctly identified 3 times as many of the test nouns as the verbs (44% versus 15% mean success in these categories, on the final trial). The results in the present experiment, in which subjects are not told the category membership of the test words, are in the same direction but are not as dramatic. On the final trial, nouns were correctly identified 26.0% of the time and verbs only 11.8% of the time. This effect can be seen in Table 1, which shows the percentage of subjects who correctly identified each target item on the final trial.

To evaluate this apparent difference, we first summed the number of correct responses on the final trial for each subject and each item. These totals were entered into separate analyses of variance (ANOVAs) for items and subjects, each of which had 3 factors: Syntactic Category; presentation List; and Order of presentation. The results indicated a significant effect of Syntactic Category ($F1(1,78) = 51.05$, $p < .001$; $F2(1,42) = 4.84$, $p < .05$).

Insert Table 1 about here

There were also substantial differences in performance across the three lists. The effect of Syntactic Category was carried by List 1 (nouns $\underline{M}$ = 38.0%, verbs $\underline{M}$ = 2.7%) and List 2 (nouns $\underline{M}$ = 31.3%, verbs $\underline{M}$ = 18.8%). In List 3, nouns were identified less often than verbs ($\underline{M}$ = 8.1% and $\underline{M}$ = 13.8%, respectively). This resulted in a significant List by Syntactic Category interaction ($\underline{F1}(2,78) = 36.82$, $\underline{p} < .001$; $\underline{F2}(2,42) = 3.49$, $\underline{p} < .05$) and a main effect of List which was reliable only in the subjects analysis ($\underline{F1}(2,78) = 18.38$, $\underline{p} < .001$; $\underline{F2}(2,42) = 1.69$, $\underline{p} > .10$). Similar effects and interactions of control variables appear in the analyses and experiments that follow. Since we believe that they have no bearing on the proposals we will make, we will not be mentioning them again. We now turn to the factors that underlie this list effect.

What kinds of target words were easy to acquire by observing the scenes? Inspection of Table 1 reveals that there is considerable variability within as well as across syntactic category. The advantage for the nouns is largely attributable to a subset of them: Subjects did better on items that are more concrete. GGGL explored this phenomenon by asking (new) subjects to rate the imageability of the target items. They found a strong correlation between imageability and identification. In the present experiment we replicated this finding: item performance was correlated with ratings of imageability both for analyses of all targets ($\underline{r} = .42$; $\underline{p} < .005$) and for separate analyses of nouns ($\underline{r} = .30$; $\underline{p} < .01$) and verbs ($\underline{r} = .44$; $\underline{p} < .01$).

So far we have treated the target items in terms of the syntactic categories noun and verb. But a closer look suggests that the effects of syntactic category are an artifact of a semantic categorization that, at least in the maternal speech tested here, is highly correlated with syntactic category. In English and in all other known languages, words denoting persons or classes of objects are grouped together and classified as nouns (Pinker, 1984; Maratsos, 1991). Child-directed speech is heavily populated with these nouns that label objects and persons. And since the selection criteria for our targets reflected usage (we chose the most frequent nouns and verbs in the maternal corporus), 75% of the test nouns in fact were labels for classes of objects or people. Subjects were twice as accurate with these words as they were with the nouns that denoted abstractions, actions, or parts.

Even within the class of object labels there was substantial variation in performance. Half of the target nouns picked out categories at the so-called basic level of hierarchical classification (BLOCs; Rosch, 1988; Rosch, Mervis, Gray, Johnson & Boyes-Braem, 1976). These words were rated as highly imageable and were correctly identified on almost half of the final trials. BLOCs account in full for the noun advantage; when they are excluded from the analysis, performance on nouns is no higher than performance on verbs. For the verb targets, performance was highest for those items that refer to physical actions and lowest for items referring to mental states (see also Snedeker and Gleitman, in press).

How closely do these results mirror child word learning? The analyses just reported suggest that the effect of syntactic category was an artifact of concreteness. "Imageable" words were the easiest to learn from the extralinguistic information. In this regard too, the results are consistent with infant production vocabularies. We assessed the extent of this correspondence by comparing our results with data from the MacArthur CDI norming study (Fenson et al., 1994). Thirty-four of the 48 target words appeared on the toddler form of the MacArthur CDI, a parental report measuring early language development (Fenson et al., 1993). For these words, we compared the percentage of children (17-30 months) in the norming study who were reported to use that word spontaneously to the percentage of our subjects who identified it. We found a substantial correlation between the two measures ($r = .36$; $p < .05$).

Experience and Learning

Many studies of child learners, and common sense, inform us that learning rarely happens the first time that a child hears a word. Instead it is likely that information about the meaning of a word accrues over many opportunities to observe the range of its applicability to objects, events, and so forth. Just so with our adult subjects who, as a group, improved as they observed successive instances of the target words. As Figures 3 and 4 illustrate, the mean percentage of correct responses for nouns increased from 17.6% in Trial 1 to 25.7% in the final trial. For verbs the mean percentage correct rose from 5.1% to 11.8%. Thus although, the learning curve is steeper for noun targets than for verb targets, there was substantial improvement for words from both lexical categories.

To further examine this pattern, correct responses to each of the trials were summed by subjects and by items and submitted to separate ANOVAs with 4 factors: Syntactic Category, List, Order and Trial (Trial 1-6 and Final Trial). The effect of Syntactic Category, which we saw in the ANOVAs of the final trials, was reliable in these analyses as well. In addition, there was a reliable effect of Trial ($F1(6,468) = 21.60$, $p < .001$; $F2(6,252) = 6.88$, $p < .001$).[5] Critically, there was a significant effect of Trial in separate analyses of both noun targets ($F1(6,468) = 9.57$, $p < .001$; $F2(6,126) = 3.41$, $p < .005$) and verb targets ($F1(6,468) = 17.33$, $p < .001$; $F2(6,126) = 4.48$, $p < .001$), indicating that performance improved as the subjects were given more word-scene pairs.

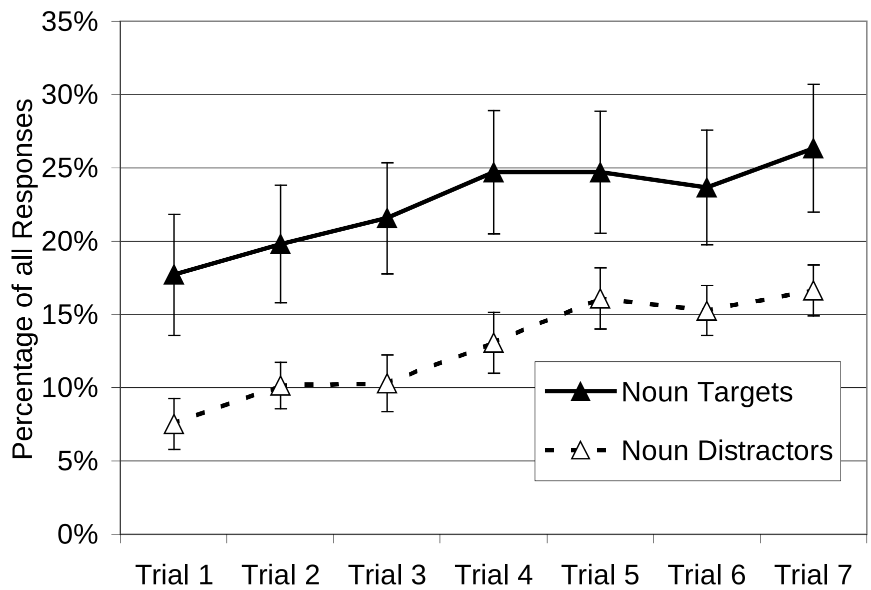

Figure 3. Percentage of target responses and distractor responses on the final trial for nouns.

These effects taken at face value seem to portend a successful career for cross-situational learning. That is, if additional instances – beyond the six that we allowed our subjects – continued to lead to similar improvement, then performance on both nouns and verbs would gradually approach ceiling. But another pattern in the data set suggests that this rosy forecast is highly implausible. While subjects are more likely to identify the correct verb as information accumulates, they also tend to converge on the same incorrect responses (see Figures 3 and 4). The ability to identify a word requires both that the input contains cues pointing to the correct meaning and that these cues are more systematic or salient than the cues pointing toward other possible meanings. If the potency of misleading cues increases faster than that of accurate cues, then additional contexts will not improve performance. Quite the contrary.

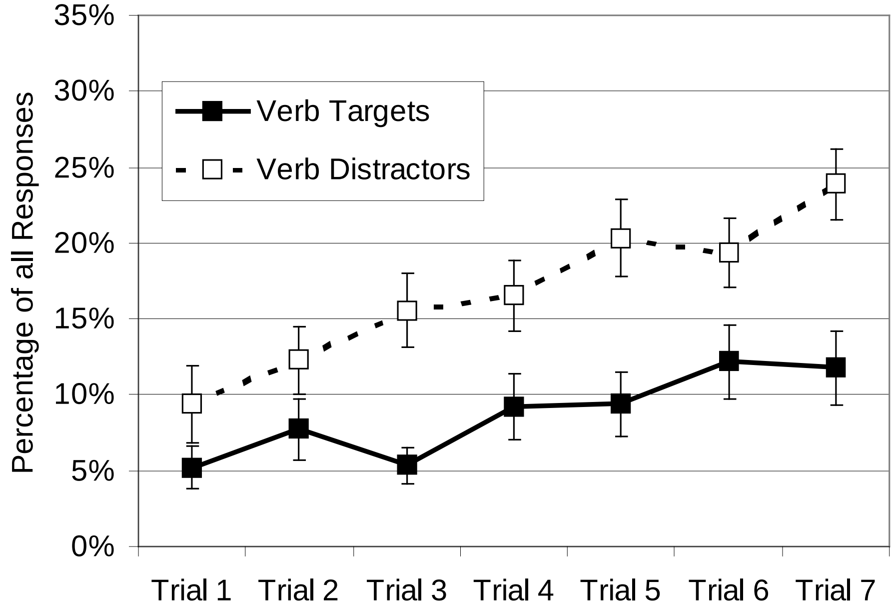

Figure 4. Percentage of target responses and distractor responses on the final trial for verbs

To explore this issue in more detail, we compared the strength of these two competing tendencies by: 1) identifying the most common incorrect answer for each target on the final trial (the "distractor"); 2) calculating the number of subjects who gave that response for each trial of that target; 3) subtracting the number of distractor responses from the number of target responses for that item; and 4) entering these difference scores into an items ANOVA. For the full set of nouns, the mean number of target responses was higher than the mean number of distractors from the first trial to the last (Figure 3) but this difference was not reliably greater than zero ($F(1,21) = 2.6$, $p = .12$). However, for the subset of nouns that picked out BLOCs, target responses were substantially and reliably more common than distractors ($F(1,9) = 7.53$, $p < .05$). For verbs, this pattern was reversed (Figure 4). Distractors began with an advantage that never waned and the difference between target responses and distractors was reliably below zero ($F(1,21) = 8.5$, $p < .001$). This means that word-scene pairs not only provided less evidence about verb meanings than about nouns, the scenes actually contained systematic misinformation, red herrings which led our subjects to converge on false hypotheses.[6]

How Category Labeling Affects Performance

As we have just shown, there is a difference in the informativeness of the scenes for identifying nouns and verbs. This was true both in GGGL, when subjects were informed about the syntactic category of each target word, and in the present experiment, where subjects received no such category labeling. Moreover, the percentage of subjects who correctly identified an item in the present experiment was very highly correlated ($r = .9$) with the number of subjects who had identified it in GGGL, $t(1,46) = 192.38$, $p < .0001$). Thus, the present experiment removed an unrealistic property of the GGGL procedure (children are not explicitly told whether a new word is a noun or a verb) but replicated their findings.

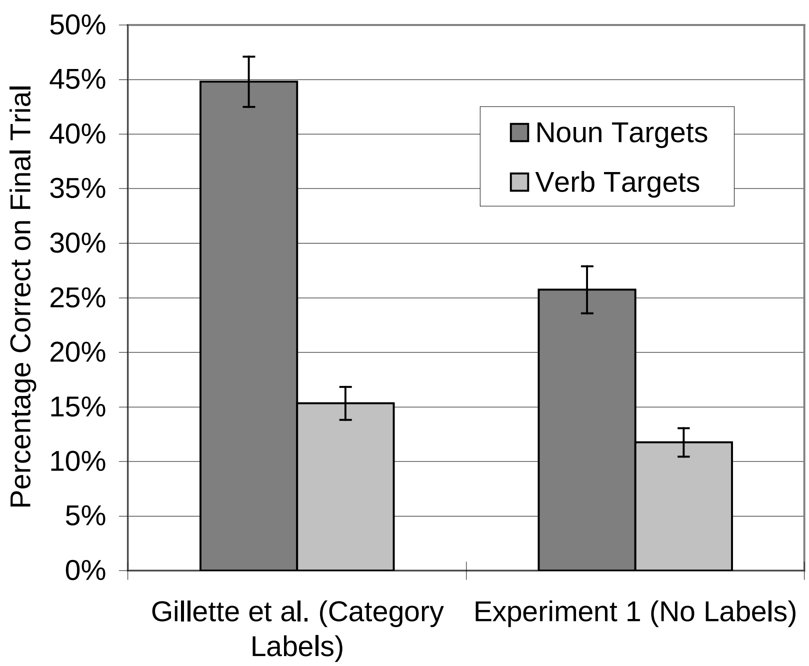

Figure 5. Percentage correct in GGGL (category labels) and Experiment 1 (no labels) as a function of the syntactic category of the target.

More interestingly, however, if we directly compare the two studies we find that the removal of category labels did not simply reduce performance accuracy across the board (see Figure 5). Instead it had a greater (negative) effect on noun identification than on verb identification. This effect was statistically evaluated as follows. Correct responses to the final trial for both experiments were submitted to subject and item ANOVAs with 4 factors each: Syntactic Category, List, Order and Experiment. Subjects generally performed better when they were given syntactic category labels and this improvement was much greater for nouns that for verbs. This resulted in a reliable main effect of Experiment ($F1(1,156) = 60.70$, $p < .001$; $F2(1,42) = 31.47$, $p < .001$) which was superseded by a significant interaction between Experiment and Syntactic Category ($F1(1,156) = 39.10$, $p < .001$; $F2(1,42) = 14.82$, $p < .001$). [7]

What Kinds of Responses Are Successful?

So far we have considered the types of targets that were relatively easy for subjects to identify; namely, BLOCs. Now we consider the types of responses that were successful. Subjects were more likely to correctly identify the targets that were nouns than the targets that were verbs, but this does not mean that they were always making a noun-like guess to noun targets and a verb-like guess to verb targets. Not surprisingly, since there were always objects in the scenes for verbs and events in the scenes for nouns, subjects often identified the wrong aspect of the scenario. In fact, incorrect guesses were no more likely to fall into the target lexical category than would be predicted based on chance alone (see Snedeker, 2000).

Our question here is whether any advantage in identification accrues to the subject when, in fact, his response is a noun rather than a verb. As Figure 6 illustrates, the answer is a clear yes. When subjects made a final response that was a noun, they were right 34% of the time but when they made a verb response, they were right only 13% of the time.[8] The higher success rate for noun responses was solely accounted for by the subset of noun responses that picked out BLOCs. When subjects made BLOC responses they were correct 44% of the time; other types of responses (verbs and non-BLOC nouns) were accurate on only 12% of the trials.[9]

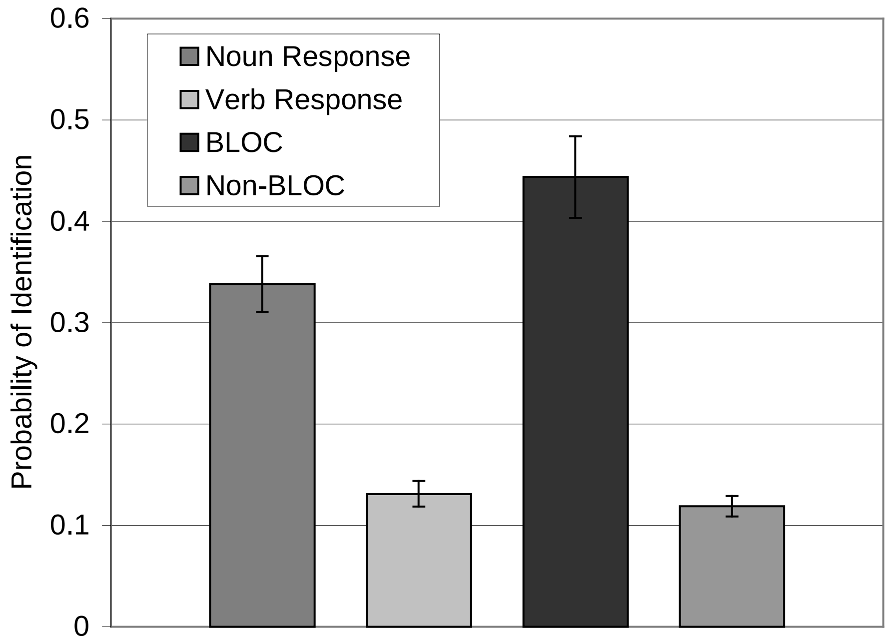

Figure 6. Probability of identification given type of response (Experiment 1).

Discussion

The present experiment had four general findings. The first was that our subjects identify nouns far more accurately than verbs when the input is limited to samples of the extralinguistic contingencies for word use. This result replicates the earlier study by GGGL whose outcomes were contaminated by the inclusion of category-label information.

The second finding was that the tested syntactic category distinction was not the real predictor of subject performance. Rather, the apparent effect of syntactic category on accuracy of identification was an artifact of a semantic classification: Basic level object terms are easier to identify from this information source than other items, including other semantic categories labeled by nouns (such as partitives and superordinates) as well as those those labeled by verbs (acts and states). This outcome is reminiscent of the dominance of object labels in the early vocabularies of children learning English.

The third finding was that withholding the category label resulted in a dramatic decrement of performance for noun but not for verb targets: Noun targets dropped from 44% correct in GGGL to only 26% correct in the present “labelless” version; while verb targets dropped only slightly, from 15% to 12% correct. That is, the explicit instruction to “be on the lookout for a noun” – which, as we showed, in the best case comes down to "be on the lookout for a BLOC” – enhanced the subjects’ chances of identifying that noun correctly. Knowing that one was looking for a verb hardly helped at all.

The fourth, and related, finding was that when subjects made a guess that was a BLOC they were almost four times as likely to be correct (44%) as when they made another type of guess (12%). This raises an interesting possibility: Given the limitations of word-to-world pairing, learners will primarily succeed when they attempt to map a new word form to a BLOC. Perhaps the dominance of object labels in early child vocabularies (and in the acquired “vocabularies” of our adult subjects) is the outcome of an adaptive learning procedure leads the learner to favor the kinds of meanings that have successful in the past (BLOCs)? If so, it would begin to account for why the noun-dominance effect, essentially invisible as learning begins, grows during the earliest period of vocabulary growth (Figure 1). With this motivation in mind, we now asked if these adult learners had any knowledge of their own performance. If so, the idea that the vocabulary learning adapts to its own successes and failures would gain some plausibility.

Experiment 2: The Presence of Evaluative Information

The comparison of GGGL and Experiment 1 suggested that learners might benefit by adopting a bias to map words to objects. But even to discover the utility of such a strategy, learners would need to implicitly recognize that success rate is linked to the semantic category of the hypothesis. The present experiment explored whether the word-scene pairs contain information that could support this kind of learning. Here again we asked adult subjects to identify words from extralinguistic contexts but this time they also were required to rate their confidence in their final responses. If word-scene pairs support accurate self-assessments, confidence ratings should reflect accuracy and they should be higher for object and noun responses.

Methods

Subjects

The subjects were 24 undergraduates at the University of Pennsylvania who received either partial class credit in a psychology course or a small payment for participation.

Stimuli and Procedure

This study used the same stimuli as in Experiment 1: the 6 video clips of each of the 24 verbs and 24 nouns that were most frequently used in the maternal speech samples, again divided into three lists and presented in two orders. The procedure was as in Experiment 1, except that after they made their final response for each word the subjects rated how certain they were that this was the correct answer. They were told to make these rating on a scale of 1 to 7, where 1 meant “guessing at random,” 7 meant “positive that [their] guess is correct,” and the other numbers stood for intermediate judgments.

Results

Identification and Comparisons with Experiment 1

The patterns of target identification in this experiment replicated those found in Experiment 1. The percentage of correct on the final trial was entered into a subject ANOVA with three factors: Syntactic Category, List and Order. There was a reliable effect of Syntactic Category ($F(1,18) = 22.55$, $p < .001$). To determine whether the additional task affected the participants’ performance or the type of responses that they made we conducted two comparisons of the data from Experiments 1 and 2. First, we entered the percentage of correct responses on the final trial for both studies into a subject ANOVA with four factors, those listed above plus Experiment. There was neither a significant main effect of Experiment (F (1,96) = 2.63, $p > .1$) nor any interaction involving this factor. Second, we examined the types of responses that participants made in each of the studies. For each subject the percentage of responses that were nouns was entered into an ANOVA with same four factors as before. Here again, no effect of Experiment was found (F (1,96) < 1, $p > .3$). The accuracy and the syntactic breakdown of the responses were apparently quite similar in the two studies despite the changes in the task.

Do Certainty Ratings Reflect Accuracy?

To determine whether participants were accurate in their appraisals, we calculated the mean certainty ratings for the correct and incorrect items for each participant and entered these ratings into an ANOVA with Correctness, List and Order and as factors. Critically, there was a strong effect of Correctness ($F(1,18) = 71.99$, $p < .001$). As shown in Figure 7, the difference between certainty ratings for correct and incorrect responses averaged 1.8 points. Twenty-two of the twenty-four participants gave higher ratings to their correct than to their incorrect responses and for twenty of those the mean difference was at least 1 rating point.

The strong effect of Correctness suggests that sets of world-scene pairs provide subjects with the information they need to evaluate the likelihood that their response is correct. But there is another plausible explanation. Subjects may have rated responses based on properties of the word, independent of the video contexts, that happened to correlate with performance. For example, subjects may have been highly confident in words that they believed were typical of maternal speech and this type of response may have been correct more often.

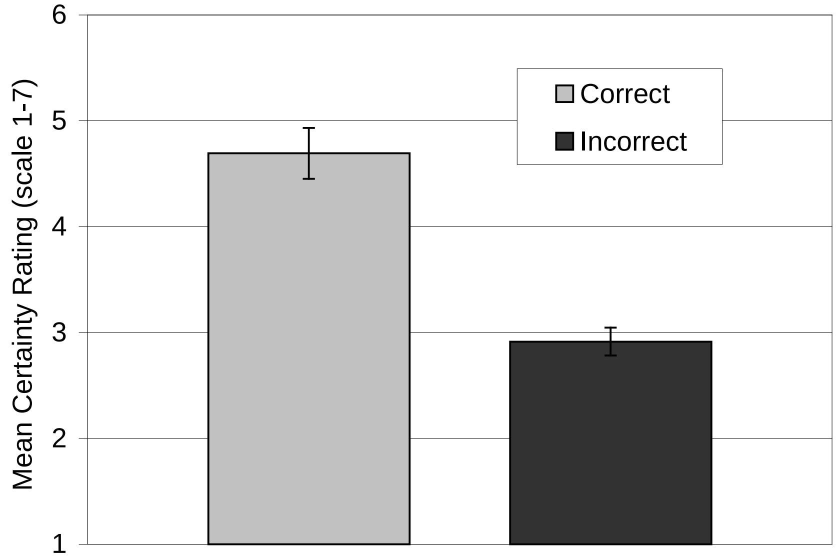

Figure 7. Mean certainty ratings for correct and incorrect responses (Experiment 2).

To determine whether properties of the response-words could account entirely for the effect of Correctness, we directly compared our subjects' confidence in a given word when that word was the target (hit trials) and when it was not (false alarms). This analysis was limited to just the 19 target words which were identified by at least one subject and appeared as incorrect responses for another target. The mean certainty ratings were entered into an item ANOVA which revealed an effect of Correctness ($F(1,18) = 8.80$, $p < .01$). Subjects gave the same response a higher rating when it was the actual target word than when it was not (M=4.23 and M=3.18, respectively). We conclude that our participants had some knowledge of their own performance and that this knowledge was based, in part, on information in the word-scene pairings.

Are Certainty Ratings Higher for Noun Responses?

In motivating the present study, we suggested that subjects who could evaluate their performance would discover that they succeeded more often when they made noun responses (and so a bias to guess more nouns should grow). To determine whether these ratings reflect the knowledge necessary for this discovery, we submitted the mean certainty ratings for noun and verb responses to a subjects ANOVA with Syntactic Category of Response, List and Order as factors.[10] As Figure 8 shows, there was a main effect of Syntactic Category of Response ($F(1,18) = 19.08$, $p < .001$).

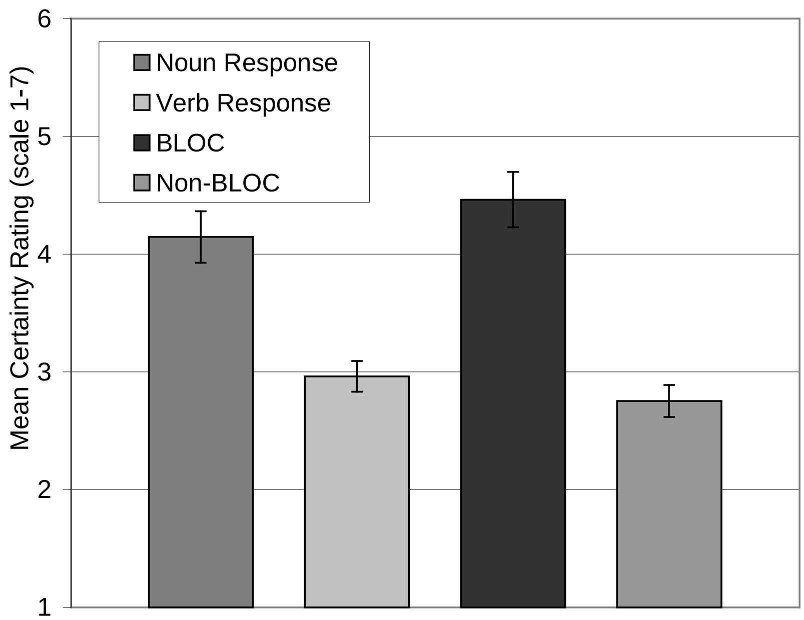

Figure 8. Mean certainty ratings given type of response (Experiment 2).

Recall that the noun advantage in Experiment 1 was attributable to high performance on the subset of noun targets that picked out basic level object categories. To determine whether BLOC responses were driving the effect of syntactic category found in the analysis above, we divided the responses into BLOCs and non-BLOCs and submitted the mean certainty ratings of each class to a subjects ANOVA with Response Type, List and Order as factors. Subjects gave substantially higher certainty ratings to the BLOC responses (see Figure 8), resulting in a main effect of Response Type ($F(1,18) = 43.43$, $p < .001$).

Why are Subjects more Certain about BLOC Responses?

There are two potential reasons why participants are more confident in the noun and BLOC response categories. First, these outcomes could simply be necessary side effects of their ability to gauge their level of knowledge. If noun and BLOC responses are more likely to be correct and subjects are able to accurately gauge their success, then it is analytic that the average rating for these types of responses would be higher. Second, confidence in the noun and BLOC responses could be higher than would be predicted on the basis of accuracy alone. Perhaps, our participants were less certain of the more abstract verbs and non-BLOC responses because they were unsure what type of evidence could be used to confirm or confute these kinds of hypothesis. If so they might have tended to rate these responses lower even when they were correct.

We explored the first explanation by conducting probability analyses identical to those described for Experiment 1. We separately calculated the mean percent correct for both noun responses and verb responses and submitted these percentages to a subjects ANOVA with the same three factors given above.[11] There was a main effect of Syntactic Category ($F(1,18) = 32.33$, $p < .001$). The probability of being correct was 3.5 times greater for noun responses than for verb responses. A parallel subjects ANOVA comparing BLOC and non-BLOC responses also produced a significant effect of Response Type ($F(1,18) = 60.19$, $p < .001$).

To determine whether the type of response had an effect on confidence in the noun and BLOC responses is directly related to the type of response itself, we needed to hold accuracy constant and test for an effect of response type. Ideally we'd do this by dividing each subject's responses along both dimensions, but the contingency between response and correctness made this problematic. Nine subjects did not have any correct verb responses. To overcome this obstacle we randomly paired one subject from each of the three presentation lists to create a single supersubject. The results of each supersubject consisted of one response to each of the 48 words. However, those responses had been contributed by three different individuals. Each supersubject had responses in all cells. The mean certainty ratings were evaluated in a subject ANOVA with Syntactic Category of Response and Correctness as factors. There was a main effect of Correctness ($F(1,6) = 21.29$, $p < .005$) but only a marginal effect of Syntactic Category of Response ($F(1,6) = 5.57$, $p = .056$). As Figure 9 illustrates, noun responses are rated somewhat higher than verbs both when they are correct and when they are incorrect.

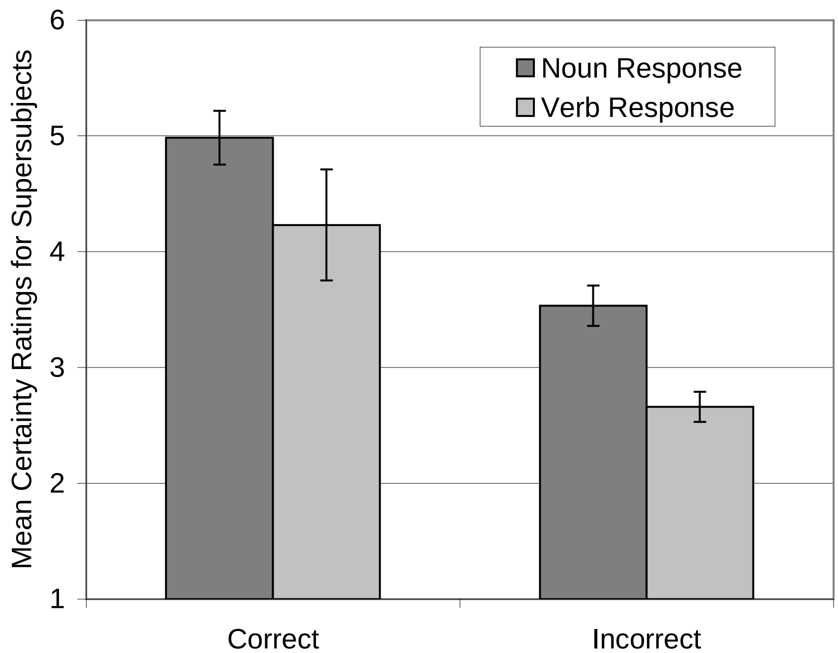

Figure 9. Certainty ratings as a function of accuracy for noun and verb responses (Experiment 2).

The same supersubjects were used to calculate mean certainty ratings for a subject ANOVA with Response Type and Correctness as factors. As Figure 10 illustrates, Correctness and Response Type had effects which were roughly equal in both reliability and magnitude ($F(1,6) = 21.66$, $p < .005$; $F(1,6) = 23.78$, $p < .005$, respectively) and there was no interaction between the two (F(1,6) = .000, p = .99). Participants were more certain of BLOC responses both when those responses were correct and when they were wrong.

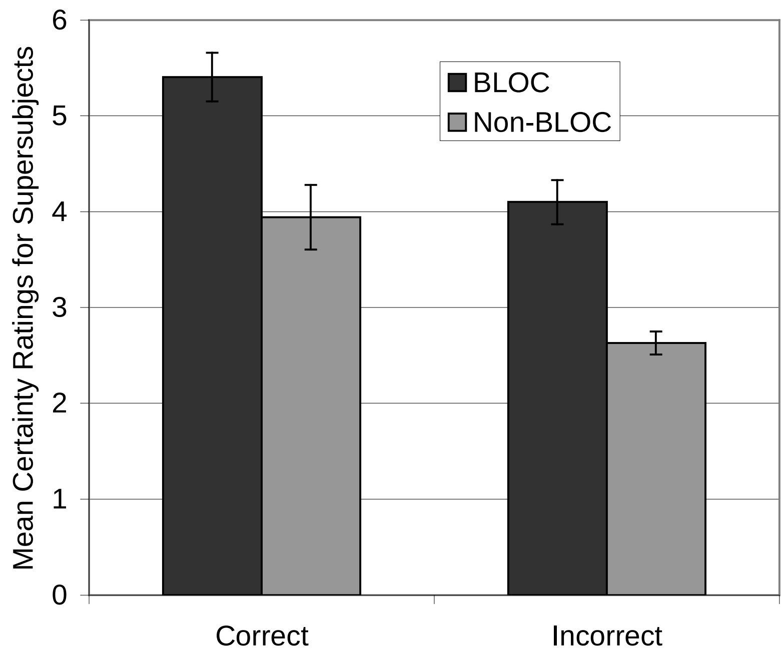

Figure 10. Certainty ratings as a function of accuracy for BLOC and non-BLOC responses (Experiment 2).

Discussion

The present results indicate that subjects were implicitly aware of when they did – and when they didn't – have adequate information to identify some of the words. There was a substantial and reliable difference between their confidence in their correct and incorrect responses. Specifically, subjects were most confident, and most accurate, for the nouns, particularly for the BLOC's. This finding bolsters the suggestion that word learners who must rely solely on extralinguistic context could profit from an examination of the categorial contingencies of correctness. The input supports accurate assessments of performance and these assessments correctly indicate that object hypotheses are more likely to be correct. Learners could increasingly gravitate toward BLOC learning because they learn that this is what works.

Experiment 3: The Effect of a Noun-Bias

The comparison of GGGL and Experiment 1 suggested that the strategy of searching for a nominal meaning for each word would improve actual performance. The present experiment tested this claim. We induced a noun bias by the simple expedient of lying to our subjects about the structure of the stimulus set. Then we compared their performance in this false belief state to that of the "unbiased" subjects in Experiment 1.

Methods

Subjects

The subjects were 42 undergraduates at the University of Pennsylvania who received either partial class credit in an introductory psychology course or a small payment for their participation.

Stimuli and Procedure

This experiment used the same stimuli, lists, and orders as the two previous experiments and the procedure was identical to that of Experiment 1. Only the subjects' instructions were changed. To induce a bias towards noun responses, and towards object categories in particular, we added the following sentences: "80% to 90% of the words on this tape are nouns; the rest are verbs. This is because mothers tend to emphasize nouns with children of this age. Most nouns pick out a category (or kind) of person or thing." This information was repeated by the experimenter when she reviewed the instructions with the subjects. In contrast, the subjects in Experiment 1 had been told that "roughly half of the words are nouns and roughly half are verbs".

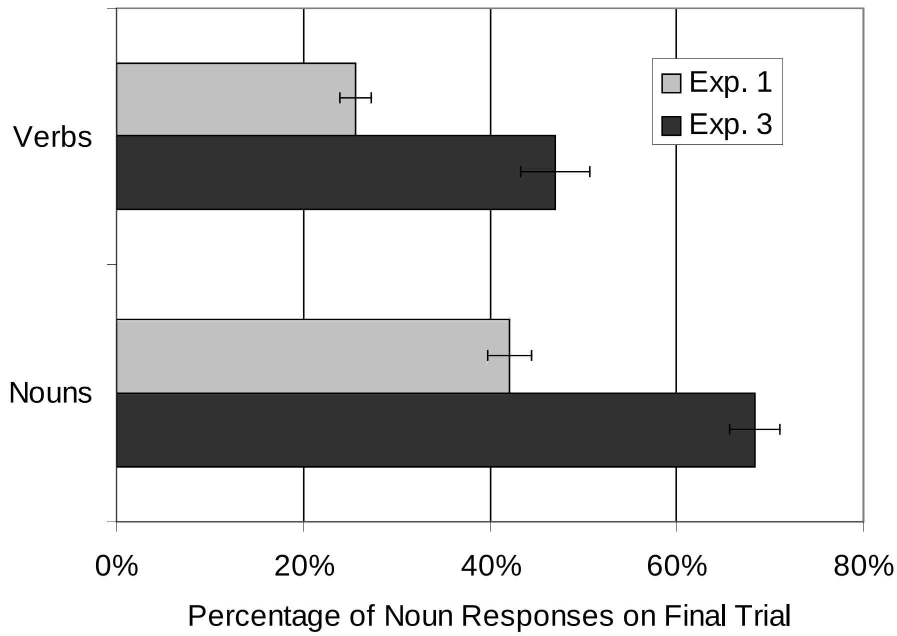

Figure 11. Percentage of noun responses in Experiment 1 (no bias) and Experiment 3 (induced bias) as a function of syntactic category of target

Results & Discussion

The Success of the Bias Manipulation

We checked the success of the bias manipulation by comparing the proportion of noun responses in Experiment 3 with the proportion in Experiment 1. All responses were classified as *noun*, *verb*, or *other* based on the subjects' post-experimental reports. For both studies, the total number of noun responses on the final trial was separately summed for noun targets and verb targets and these totals were submitted to a subject ANOVA with Syntactic Category of Target, Experiment, List, and Order as factors. There was a main effect of Syntactic Category ($F(1,114) = 69.70$, $p < .001$), reflecting the fact that in both studies noun responses were more common for noun targets. Crucially, there was also a main effect of Experiment ($F(1,114) = 80.66$, $p < .001$); nouns made up only 34% of the responses in Experiment 1 but comprised 58% in Experiment 3. There was no reliable interaction of Syntactic Category of Target and Experiment ($F(1,114) = 1.13$, $p > .25$). As Figure 11 indicates the increase in noun responses was spread equally across the two target types, suggesting that the change in instructions merely changed subjects response bias, not their sensitivity to syntactic category of the target.

Does the Induced Noun Bias Improve Performance?

Next we compared the relative success of participants in the two studies to find out whether this artificially created bias had led to the predicted improvement in word identification. For these analyses, homonymous words from a different syntactic category than the target were scored as incorrect.[12] Number of correct responses to verb targets and noun targets was separately summed for each trial and submitted to a subject ANOVA with Syntactic Category of Target, Trial, Experiment, List, and Order as factors. The effects discussed in the trial-by-trial analysis of Experiment 1 (main effects and interactions of List, Syntactic Category of Target, and Trial) were also found in this analysis. Here we will report only effects involving the factor of Experiment. Critically, there was a main effect of Experiment ($F(1,114) = 6.54$, $p < .05$). Across the entire set of trials, the subjects in Experiment 3 correctly identified more of the target words than those in Experiment 1. However, there was also an interaction between Experiment and Syntactic Category of Target ($F(1,114) = 32.99$, $p < .001$) and a three-way interaction between Experiment, Syntactic Category of Target, and Trial ($F(6,684) = 4.72$, $p < .001$). As Figure 12 illustrates, these interactions are due to differences in how the bias manipulation affected subjects' ability to identify noun targets and verb targets.

To characterize these differences more accurately we performed two additional analyses. First, since we were primarily interested in what subjects ultimately learned, we ran an additional ANOVA on the totals from the final trial. There was a reliable main effect of Experiment ($\underline{F}(1,114) = 6.91$, $p = .01$). The induced bias lead to an overall improvement in performance. However, there was also an interaction between Experiment and Syntactic Category of Target

(F(1,114) = 34.90, $p < .001$). Subjects in Experiment 3 performed much better on the nouns than those in Experiment 1 (M = 36.6%, M = 23.1%) and this difference was highly significant (F(1,114) = 26.3, $p < .001$). However, performance on verbs suffered in Experiment 3 (M = 6.5% down from M = 11.5%) and this difference, while smaller in magnitude, was reliable as well (F(1,114) = 7.07, $p < .01$).

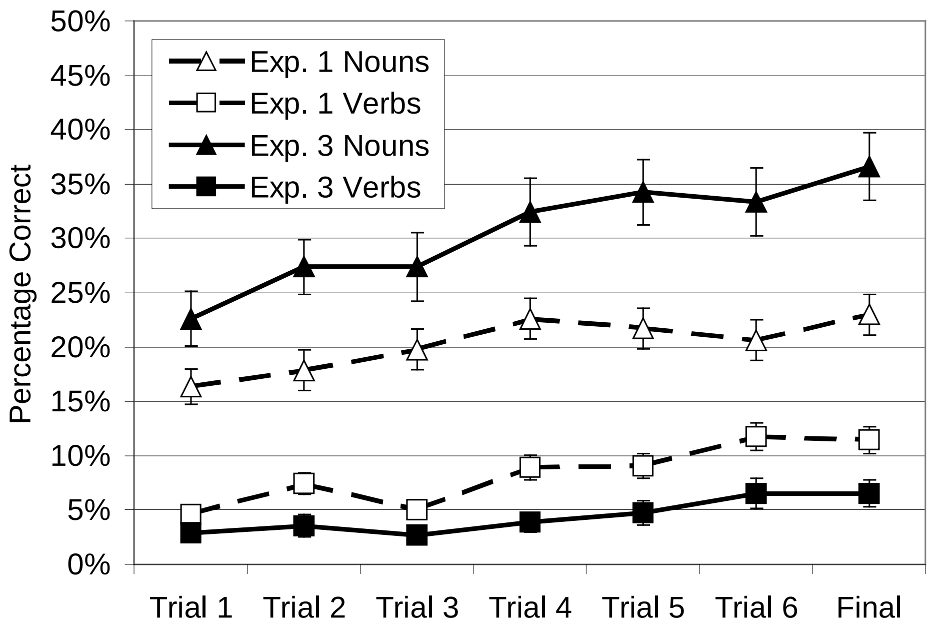

Figure 12. Percentage correct in Experiment 1 (no bias) and Experiment 3 (induced bias) as a function of trial

Next, we attempted to track down the source of the Syntactic Category of Target, Experiment and Trial interaction by conducting separate cross-trial analyses for noun targets and the verb targets. In the ANOVA for the noun targets there was a significant effect of Experiment (F(1,114) = 21.82, $p < .001$), attributable to consistently higher performance in Experiment 3 across trials, and a significant interaction between Trial and Experiment (F(6,684) = 3.14, $p < .005$). Figure 12 points to the source of this interaction: While the identification of noun targets in Experiment 1 leveled off by the fourth learning trial, identification of noun targets in

Experiment 3 continued to improve. In the ANOVA for the verb targets, we also found an effect of Experiment ($F(1,114) = 11.74$, $p < .001$); verbs were identified more often in Experiment 1. However, there was no significant interaction between Experiment and Trial ($F(6,684) = 1.59$, $p = .15$). As Figure 12 illustrates subjects in Experiment 3 showed a constant decrement in performance for verb targets.

General Discussion and Conclusions

The experiments reported in this paper were designed to uncover those aspects of vocabulary acquisition that are independent of conceptual growth and attributable, instead, to variation in the information base available to any human learner. This is the rationale behind the Human Simulation Paradigm, as investigated by GGGL: Their findings were that word-learning by adults as a function of extralinguistic observation – and extralinguistic observation only—reproduced robust properties of earliest learning in toddlers; namely, slow and errorful learning, with success largely limited to concrete nominal terms.

The experiments in the present paper had as their most general aim to understand the mechanisms that account for this patterning, as it relates to child language growth. We will suggest, based on the present experimental outcomes, that the rise and fall of "noun dominance" and "object dominance" (in the life of toddlers and their adult simulators) provides important clues to the machinery of word learning.

A Primitive Model of the Object Bias

The experiments just reported, taken together with collateral evidence from the literature of child learning, provide suggestive evidence for a learning model in which an object bias grows in learners as a consequence of how they represent and encode the situations that accompany word use. In introductory remarks, we sketched out an account of how this bias could be learned through interaction with the input. This primitive model consisted of six propositions, which we now reproduce, and reconsider in light of our findings:

Proposition 1: An object-kind bias is not present in the prelinguistic period, but appears and grows during the earliest stages of lexical learning.

Because our experimental "learners" were adults and not young children who are just beginning to acquire words, clearly their word-identification performance cannot tell us about the starting state of novice word learners. However, documentation of this first proposition of our model comes from a large literature showing that the object bias is nonexistent or very weak in the infant as word learning begins, and may not emerge at all under certain input conditions (a topic to which we will return presently). Vocabulary counts of learners of several languages suggest that noun dominance rises over the life of the infant (for English see Fig. 1 from Bates, Dale & Thal, 1995; Benedict, 1979 and Nelson, 1974; for Italian Caseli et al., 1995; for Hebrew Maital et al., 2000). During the same period of time children's ability to learn object labels improves (Leonard, Schwartz, Morris & Chapman, 1981). In addition, experimental studies of children's generalizations about the meanings of novel words suggest increasing potency of a shape or object-kind bias over developmental time (Samuelson & Smith, 1999). This rise in the object bias is a feature that will be critical to modeling early word learning.

Proposition 2: Because of the properties of word-to-world mapping, novices will tend to succeed when they map a word to an object or person but to fail otherwise.

In Experiments 1 and 2, we demonstrated that adults identifying very simple words by observing the contingencies for their use in maternal speech succeeded far more often when these words were nouns rather than verbs (Table 1, Figures 3 & 4). Thus we extended the findings of GGGL by finding a noun-verb difference even when subjects were not told about the syntactic category of the target. Pursuing this finding to its source, analysis showed that the advantage for nouns was an artifact of semantic typology; specifically, the ease of identifying BLOCs when the information base consists only of extralinguistic evidence (Figures 6 & 9). This is despite the fact that, in real life, our adult subjects know partitives like nose and superordinates like thing as well as they knew BLOCs like ball. Yet when their task was to recognize such items with only the scene information to guide them, the BLOCs were overwhelmingly easier. The same account pacifies the relevant data of early child vocabulary. Youngest learners cannot bring collateral language-internal evidence to bear on constraining their word-meaning conjectures. They have not yet built up the sizeable databases necessary to know (statistically speaking) which words go with which other words and under which structural descriptions, i.e., they cannot exploit clues that arise from the language-specific grammars of the exposure language. What these novices can do – right from the beginning of learning—is to observe the circumstances in which particular words are uttered. They can pair the words to the world. These context-based usage properties provide strong and reliable cues to the meanings of BLOCs, but not to the meanings of other words, be these abstract nouns or verbs. In short, the learner who proceeds by word-to-world pairing will succeed with BLOCs far more often than with other words. He needs no built-in object-kind bias for this to happen. The constraint on his evidentiary base guarantees this outcome.

Proposition 3: The word learning system has access to information that allows it to monitor its success.

To discover that the input favors BLOCs the learner must be able to accurately judge when she has correctly identified the meaning of a word. Experiment 2 showed both that the input contains information which could support accurate self-assessments and that our adult subjects were able to monitor these properties of the input (Figure 7).

But can children evaluate their own word learning performance? Obviously any learner must have some capacity and inclination to store past information and evaluate it against newly arriving input in the process of word learning. Otherwise all word-learning would be one-trial learning (and most of it would be wrong). The best experimental evidence that children can quite accurately determine when they have discovered the meaning of a word comes from the work of Clark (1987) and Markman and her colleagues (e.g., Markman & Wachtel, 1988) on children's ability to use the principle of contrast or mutual exclusivity in word learning. If a child can use the fact that a given word has a particular meaning to determine that a second word cannot have that same meaning, then she must be relatively confident in her knowledge of the meaning of the first word (or at least more certain of it than she is of the second).

Proposition 4: The system discovers the correlation between object responses and high performance (or confidence) and exploits this correlation by adopting a bias to map each word to an object category.

To discover that an object-kind bias would be advantageous, learners need to monitor more than their confidence in individual word meanings. They must also keep track of their relative confidence in meaning hypotheses from different semantic classes. Experiment 2 indicated both that the input contained information that could inform learners of the greater accuracy of object responses, and that our adult subjects had access to this information. The subjects were more confident of noun and BLOC responses, even when those responses were incorrect (Figures 8, 9 & 10).

Again, we have no direct evidence of whether and to what degree the child, can or does monitor such general properties of the input, though recent evidence makes a strong plausibility argument to this effect. For one thing, experimental studies of infant segmentation performance make clear that from at least the age of 8 months the human rapidly and efficiently extracts and encodes statistical properties of input speech (Saffran, Aslin & Newport, 1996; Newport, 1999). Suggestive findings more specifically related to our proposal have been put forward by authors studying crosslinguistic differences in word meanings and word learning. Very young children rapidly acquire linguistic-classificatory biases that are consistent with the vocabulary stock of their own language but inconsistent with other languages (e.g., linguistic spatial categories in Dutch vs Korean-speaking children, Choi & Bowerman, 1991; linguistic object-substance categories in English vs Japanese-speaking children, Imai & Gentner, 1997; linguistic path/manner verb categories in Spanish and Greek versus English speaking children, Papafragou, Massey, & Gleitman, 2000; Ozcaliskan & Slobin, 1999). These findings taken together with the results of Experiment 2 render the growth of an object bias understandable: Infant learners, like our subjects, have noticed that their success rate goes up when the meaning conjecture is a BLOC. Feeding upon success, then, the learner increasingly conjectures BLOCs even when such guesses are wrong. For a good stretch of time, it makes increasing sense to look under the lamppost.

Proposition 5: Adopting a bias for object kind categories leads to more successful and efficient word learning and is therefore maintained, at least until information about linguistic structure enters into the learning procedure.

The effect of an object bias on performance of the word learning system was explored in Experiment 3, where we compared the performance of subjects who were given an object bias with those who were not. We found an overall advantage in the bias condition that was superseded by an interaction between bias and syntactic category. The induced noun-bias led to a significant improvement for the noun targets and a smaller but reliable decrement for the verb targets. Furthermore, the improvement in performance for the noun targets increased across trials while the decrement for verb targets remained steady, suggesting that the biased learners would gain more from additional tokens than the unbiased learners (Figure 12).[13]

Proposition 6. The object bias is overcome when further types of information (such as co-occurrence and structural information) become available to the learning machinery.

Why should the object bias ever go away? Note in Figure 1 that by the time vocabulary size reached about 500 words (typically around 2 ½ or 3 years of age) noun dominance has subsided. How is it that older toddlers and young children go on to acquire words from every lexical class rapidly and efficiently? Again, there seems little reason to assign these new achievements to changes in the learner's mentality. It was noted at least as early as Lenneberg (1967) that rate of learning as well as the character of the lexicon (i.e., its categorial breadth) is so closely correlated with evidence of rudimentary syntactic knowledge as to suggest a causal relationship (for later commentary to this same effect: Gleitman, Gleitman, Landau, & Wanner, 1988; Bickerton, 1991; and Bates, Dale & Thal, 1995). GGGL and Snedeker and Gleitman (in press) studied this relationship with adults again using the Human Simulation Paradigm. The method was to supply additional, or alternate, sources of information to these subjects as they tried to guess the meanings of the verb vocabulary items that were so intractable to learning from extralinguistic context. Subjects efficiently identified the words as soon as co-occurrence and structural information were made available. Indeed, this language-internal information was most informative for learning least concrete terms: Where extralinguistic observation was the only source of evidence in these experiments, the most concrete words had the learning advantage; where linguistic observation (of the structural contexts for verb use) was the only source of evidence, the most abstract words had the learning advantage.[14]

The Place of the Object Bias in Vocabulary Learning

We have just described the efficacy of an object bias as it predicts the behavior of our simulated learners, and as it comports with – and to some degree explains – the natural history of this same bias in child language learners. The object bias enables novices to acquire a first stock of items in the exposure language. This small vocabulary of concrete items forms a scaffolding on which discoveries about the statistical and structural organization of the native language can be built (for more extensive discussion of later stages of vocabulary growth, see GGGL). The object bias (and its output, a concrete, noun-dominated vocabulary) critically grounds the first steps in language learning but later this bias must be supplemented or conditioned on linguistic context to make use of new information sources and enable further linguistic growth.

We now turn to two further issues whose resolution affects the viability of the object-bias model and its postulated role in language acquisition. Both these issues again pertain to the constitution of novice vocabularies. The first has to do with why abstract terms are not just wrong but almost absent from children's early vocabularies. The second has to do with the variability of the object bias in children learning different languages.

Identifiability and the Novice Lexicon

Throughout, we have approached the problem of vocabulary acquisition as a problem of information availability rather than as a problem of conceptual capacity. Children learn concrete items first because these are the items that submit to the word-to-world pairing procedure, the only word-learning device in the earliest armamentarium of the child learner. The opposing position is that young children learn concrete items first because their little brains and their deficient experience do not render them capable of appreciating abstractions. To compare these positions has been a basic purpose of the Human Simulation Paradigm which reproduces in sophisticated adults the same information basis presumed to be available to infants acquiring a first vocabulary. To the extent that these populations in possession of the same information solve the mapping problem for vocabulary in the same ways, the information-based account receives some support.

But there is, at least to first inspection, one huge difference in the vocabulary feats of young children versus our experimental subjects which, if left unresolved, points to a disanalogy in their learning performance, as follows: It is true that our adults, like infant learners, come to map many concrete nouns onto their correct meanings based on observation of their contexts. But our adults, observing such contingencies, come up with false mappings of most of the verbs (Experiment 1, see also GGGL, p. 148). These errors of commission are not characteristic of young children. The children don't seem to use verbs to refer falsely to states in events, e.g., saying "eat" when they mean 'want.' Instead, for some relatively lengthy period, they hardly utter verbs (Figure 1). So the question is: Given that verbs are hard to learn from extralinguistic observation alone, why don't young children utter verbs under the wrong semantic interpretations?

When we examine this question in light of the confidence-rating data, the child and adult outcomes are no longer inconsistent. The difference is largely that we forced our subjects to respond with their "best guess" as to the meaning of each word whether or not they felt they had adequate evidence for this guess. In contrast, young children have the luxury of not uttering words for whose meanings they have little or no confidence.

In sum, the confidence-rating findings indicate that our adult subjects knew that they didn't have adequate information to identify some of the words. Furthermore, they were particularly unsure of the identity of verb targets (the mean confidence rating for BLOCs was 4.5 while for verbs it was 3). While the input supports false hypotheses for verbs and for nouns that are not BLOCs, there is a clear difference between the weak extralinguistic evidence for these false responses and the strong extralinguistic evidence that allows them to identify BLOC terms. The suggestion is that littlest children are in much the same state. They say and understand many nouns and rarely utter verbs at all because they have the luxury of keeping quiet until they have some confidence that they are saying what they mean.

Variability in the Object Bias

As we have stated it, the object bias is an output of an adaptive strategy. It arises as the effect of the child's attempt to relate what she hears to its observable correlates in the conversational context rather than from, say, some innate or natural bias toward things rather than acts, properties, or abstractions. We have claimed, based on the confidence-judgments of Experiment 2 (Figures 7 & 8), that this strategy will feed upon its own success thus accounting for the rise in the object bias in toddlers. But if the object bias is an ascending function of its own success, this implies that the bias will vary under differing learning circumstances, and indeed this seems to be the case.

There may be a sense that can be made precise in which objects are more straightforwardly available to perception ("more easily observed") than actions (for useful discussions, see e.g., Baillargeon, 1993, Gentner, 1982; Landau & Jackendoff, 1993; Pylyshyn, 2000). If so, then the object bias would be expected to occur in most natural environments. However, just because the forces supporting this bias are adaptive and success-driven, the degree to which it manifests itself will depend upon how transparently a given cultural context or language evidences object mappings, and on how strongly other factors intrude in opposing directions. It is reasonable, then, to expect crosslinguistic differences in both the onset of the object bias and its potency.

Crosslinguistic work on vocabulary composition indeed shows the predicted variability. In a number of languages, and after a brief and categorially disorganized initial state (cf., Dromi, 1987), children's early vocabulary is dominated by nouns (Maital et al., 2000 for Hebrew; Casselli et al., 1995 for Italian; Goldin-Meadow, Seligman & Gelman, 1976, Benedict, 1979, and Fenson et al., 1994 for English; Jackson-Maldonado, Thal, Marchman, Bates & Gutierrez-Clellen, 1993 for Spanish; Gentner, 1982 for others). However, the degree of noun dominance appears to reflect the salience, concreteness, and frequency of different types of words in the input (Gentner & Boroditsky, 2001). For example, the noun advantage is smaller, or perhaps even non-existent, in languages like Korean or Mandarin which allow rampant omission of noun phrases (Tardif, 1996 and Tardif, Gelman & Xu, 1999 for Mandarin; Choi & Gopnik, 1995 for Korean). This difference is captured quite naturally by our adaptive model. In English, nouns in general and object labels in particular will be favored by their high imagability and the relative frequency with which they appear in the input. But in languages which allow "argument dropping," the nouns themselves are often to be omitted from the linguistic description of an event. Furthermore, the easier it is to recover the intended nominals pragmatically from observation, the greater the likelihood of their omission from speech. This means that object concepts will often be present in the child's representation of the event when the words encoding these are absent. This necessarily decreases the strength of the association between the object concepts and the object labels. Object labels may still be favored over verbs and abstract terms because of the countervailing influence of concreteness, but their dominance in early acquisition should clearly be weaker.

Symmetrically, to the extent that a language includes among its most frequent stock highly imageable verbs, both the degree of noun dominance and the strength of the object bias should be diminished. Recent work on two Mayan languages provides support for this idea. Both Brown (1998) and deLeon (1999) have presented acquisition data suggesting that children learning a language with object and subject omission and, critically, many extremely imageable verbs have early lexicons dominated by verbs.

In any case, findings based on counts of child vocabulary, even when taken from numerous languages, offer only a murky view of the learning function that is generating them. As Markman and Woodward (1991) point out, the existence of a bias may be masked by a variety of crosscutting factors influencing which words children know, and which of these they care to utter. That is, they may learn many words that are disfavored by a bias, either because they are favored by another bias (usually mutual exclusivity) or because a persistent and salient signal in the input overpowers the child's tendency to prefer meanings of another type. Testing the adaptive model with any precision will ultimately require controlled experiments on the development of word learning biases in children learning a variety of languages.

More generally, there are clear functional advantages to allowing biases for particular semantic types to be determined by the input. Doing so allows learners to maximize their success despite variation in the target language or the input they receive. In addition it allows for bias to be recalculated and reconfigured as the child's representation of the input becomes more elaborate. Waxman (1999) has found that the general bias to map all words to objects is rapidly reshaped into a bias to map nouns to objects. An account that relies on innate biases does not easily capture this contingency or the subsequent appearance of "biases" to map verbs to actions (e.g., Pinker, 1984; Landau & Gleitman, 1985) or adjectives to properties. (e.g., Klibanoff & Waxman, 2000; Mintz & Gleitman, 1998).

Final Comments

Throughout this paper we have argued that the early quirks and predominances in children's lexicons and learning procedures are materially the effect of the circumstances under which children solve the mapping problem, rather than the product of a pre-existing bias or limitation on what concepts they can or do represent. We are not suggesting that by enriching and deepening the postulated learning procedure for the lexicon, theories concerning natural ("innate") human categories can be weakened or dispensed with. We do not even consider such a tradeoff position to be coherent. What we have suggested is that the natural conceptual repertoires of children and adults are likely more similar than some authors have been disposed to believe. The lexical acquisition problem is often framed as one whose chief limiting factor is the cognitive status of the young learner. The evidence adduced is the concreteness and banality of children's early words (see, e.g., Gopnik & Meltzoff, 1997). Our findings suggest that the chief limiting factor may not be so much in what the young child can conceive, but in the information for deciding how some heard word (qua phonetic sequence) lines up with the discourse environment, and thus what it might mean.

References

Au, T., Dapretto, M. & Song, Y. (1994). Input versus constraints: Early word acquisition in Korean and English. Journal of Memory and Language, 33, 567-582.

Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. Granrud (Ed), Visual perception and cognition in infancy (pp. 265-315). Hillsdale, NJ: Erlbaum.

Bates, E., Dale, P.S. & Thal, D. (1995). Individual Differences and Their Implications for Theories of Language Development. In Fletcher, P. & MacWhinney, B. (Eds.), The handbook of child language (pp. 96-151), Oxford, England: Blackwell Publishers.

Benedict, H. (1979). Early lexical development: Comprehension and production. Journal of Child Language, 6, 183-200.

Bickerton, D. (1991). Haunted by the specter of creole genesis. Behavioral & Brain Sciences, 14, 364-366.

Bloom, P. & Kelemen, D. (1995). Syntactic cues in the acquisition of collective nouns. Cognition, 56, 1-30.

Brown, P. (1998). Children's first verbs in Tzeltal: Evidence for an early verb category. Linguistics, 36, 713-753.

Brown, R. (1957). Linguistic determinism and the part of speech. Journal of Abnormal & Social Psychology, 55, 1-5.

Brown, R. (1973). A First Language. Cambridge, MA: Harvard University Press.

Caselli, M.C., Casadio, P., & Bates, E. (1999). A comparison of the transition from first words to grammar in English and Italian, Journal of Child Language, 26, 69-111.

Caselli, M.C., Bates, E., Casadio, P., Fenson, J., Fenson, L., Sanderl, L., & Weir, J. (1995). A cross-linguistic study of early lexical development. Cognitive Development, 10, 159-199.

Choi, S. & Bowerman, M. (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. Cognition, 41, 83-121.

Choi, S. & Gopnik, A. (1995). Early acquisition of verbs in Korean: A cross-linguistic study. Journal of Child Language, 22, 497-529.

Chomsky, N. (1959). Verbal Behavior, By B.F. Skinner, Language, 35, 26-58.

Clark, E.V. (1987). The principle of contrast: a constraint in language acquisition. In B. MacWhinney (Ed.) Mechanisms of language acquisition (pp.1-33), Hillsdale NJ: Erlbaum.

DeLeon, L. (1999). Why do verb roots predominate in Tzotzil (Mayan) early lexicons?: An evaluation of factors of input and interaction. Paper presented at the Eighth Congress of the International Association for the Study of Child Language, San Sebastian, Spain, July 1999.

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

Fenson, L., Dale, P., Reznick, S., Bates, E., Thal, D., & Pethick, S.(1994). Variability in early communicative development. Monographs of the Society for Research in Child Development, 59.

Fenson, L., Dale, P., Reznick, S.,Thal, D., Bates, E., Harting, J., Pethick, S., & Reilly, J. (1993). The MacArthur Communicative Development Inventories: User's Guide and technical manual. San Diego, CA: Singular Publishing Group.

Gentner, D. (1978) On relational meaning: The acquisition of verb meaning. Child Development, 49, 988-998.

Gentner, D. (1981). Some interesting differences between nouns and verbs. Cognition and Brain Theory, 4, 161-178.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S.A. Kuczaj, II (Ed.), Language development: Vol 2. language, thought, and culture (pp. 31-53). Hillsdale, NJ: Erlbaum.

Gentner, D., & Boroditsky, L. (2001). Individuation, relativity and early word learning. In M. Bowerman & S. Levinson (Eds.), Language acquisition and conceptual development. England: Cambridge University Press.

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulation of vocabulary learning. Cognition, 73, 135-176.

Gleitman, L. (1990). The structural sources of word meaning. Language Acquisition, 1, 3-55.

Gleitman, L., Gleitman, H., Landau, B., & Wanner, E. (1988). Where learning begins: Initial representations for language learning. F. Newmeyer (Ed) Language: Psychological and biological aspects. Linguistics: The Cambridge survey, Vol. 3, (pp. 150-193). Cambridge: Cambridge University Press.

Goldin-Meadow, S., Seligman, M.E. & Gelman, R. (1976) Language in the two-year-old. Cognition, 4, 189-202.

Goodman, N. (1966). The structure of experience. 2nd ed. Indianapolis, IN: Bobbs-Merrill.

Gopnik, A. & Meltzoff, A.N. (1997) Words, thoughts, and theories. Cambridge, MA: MIT Press.

Huttenlocher, J., Smiley, P. & Charney, R. (1983). Emergence of action categories in the child: Evidence from verb meanings. Psychological Review, 90, 72-93.

Imai, M. & Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. Cognition, 62, 169-200.

Jackson-Maldonado, D., Thal, D., Marchman, V., Bates, E. & Gutierrez-Clellen, V. (1993). Early lexical development in Spanish-speaking infants and toddlers, Journal of Child Language, 20, 523-549.

Johansson, S. & Hofland, K. (1989). A frequency analysis of English vocabulary and grammar: Based on the LOB corpus. New York: Oxford University Press.

Klibanoff, R. & Waxman, S. (2000). Basic level object categories support the acquisition of novel adjectives: evidence from preschool-aged children. Child Development,71, 649-659.

Landau, B. (1994). Where's What and What's Where: The Language of Objects in Space. Lingua, 92, 259-296.

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

Landau, B. & Jackendoff, R. (1993). "What" and "where" in spatial language and spatial cognition. Behavioral & Brain Sciences, 16, 217-265.

Leonard, L., Schwartz, R., Morris, B., & Chapman, K. (1981). Factors influencing early lexical acquisition: lexical orientation and phonological composition. Child Development, 52, 882-887.

Lenneberg, E.H. (1967). Biological foundations of language. New York: Wiley.

Maital, S., Dromi, E., Sagi, A. & Bornstein, M. (2000) The Hebrew Communicative Development Inventory: Language specific properties and cross-linguistic generalizations. Journal of Child Language, 27, 43-67.

Macnamara, J. (1972). Cognitive basis of language learning in infants. Psychological Review, 79, 1-13.

Markman, E. (1989). Categorization and naming in children: Problems of induction. Cambridge, MA: MIT Press, Bradford Books.

Markman, E. (1992). Constraints on word learning: Speculations about their nature, origins, and domain specificity. The Minnesota symposia on child psychology, 25, Modularity and constraints in language and cognition (59-101). Hillsdale, NJ: Erlbaum

Markman, E. & Wachtel, G. (1988). Children's use of mutual exclusivity to constrain the meaning of words. Cognitive Psychology, 20, 121-157.

Markson, L. & Bloom, P. (1997). Evidence against a dedicated system for word learning in children. Nature, 385, 813-815.

Maratsos, M. (1991). How the acquisition of nouns may be different from that of verbs. In A. Krasnegor, D. Rumbaugh, R. Schiefelbusch, & M. Studdert-Kennedy (Eds.), Biological and Behavioral Determinants of Language Development (67-88). Hillsdale, NJ: Erlbaum.

Miller, G. (1951). Language and Communication. New York: McGraw-Hill.

Mintz, T. & Gleitman, L. (1998). Incremental language learning: two and three year olds' acquisition of adjectives. Proceedings of the 20th Annual Conference of the Cognitive Science Society, Hillsdale, NJ: Earlbaum.

Nelson, K. (1974). Concept, word and sentence: Interrelations n development. Psychological Review, 81, 267-285.

Nelson, K., Hampson, J., & Shaw, L.K. (1993). Nouns in early lexicons: evidence, explanations and implications, Journal of Child Language, 20, 61-84.

Newport, E. (1999). Innately constrained learning: blending old and new approaches to language acquisition. Presented at the Twenty-fourth Boston University Conference on Language Development, Boston, MA, November 1999.

O'Grady, W. (1987). Principles of grammar & learning. Chicago, IL: University of Chicago Press.

Ozcaliskan, S. & Slobin, D. (1999). Learning How to Search for the Frog: Expression of Manner of Motion in English, Spanish, and Turkish. Proceedings of the Twenty-third Boston University Conference on Language Development. Somerville, MA: Cascadilla Press.

Papafragou, A., Massey, C. & Gleitman, L. (2000). Shake, rattle & roll. Paper to be presented at the Twenty-fifth Boston University Conference on Language Development, Boston, MA, November 2000.

Pinker, S. (1984). Language learnability and language development. Cambridge, MA: Harvard University Press.

Pylyshyn, Z. (2000). Situating vision in the world. Trends in Cognitive Sciences, 4, 197-207.

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

Rosch, E. (1988). Principles of categorization. In A. Collins & E. Smith (Eds), Readings in cognitive science: A perspective from psychology and artificial intelligence (pp. 312-322), San Mateo, CA, USA: Morgan Kaufmann, Inc.

Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., & Boyes-Braem, P. (1976) Basic objects in natural categories. Cognitive Psychology, 8, 382-439.

Saffran, J., Aslin, R. & Newport, E. (1996). Statistical learning by 8-month-old infants. Science, 274, 1926-1928.

Samuelson, L. (2000). Attentional Biases in Artificial Noun Learning Tasks: Generalizations Across the Structure of Already-Learned Nouns. Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society, Mahwah, NJ, USA: Lawrence Erlbaum Associates, Inc.

Samuelson, L. & Smith, L. (1999). Early noun vocabularies: Do ontology, category structure and syntax correspond? Cognition, 73, 1-33.

Sandhofer, C., Smith, L., & Luo, J. (2000) Counting nouns and verbs in the input: Differential frequencies, different kinds of learning? Journal of Child Language, 27, 561-585.

Siskind, J. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. Cognition, 61, 39-91.

Smiley, P. & Huttenlocher, J. (1995). Conceptual development and the child's early words for events, objects, and persons. In M. Tomasello & W.E. Merriman (Eds.), Beyond names for things: Young children's acquisition of verbs (pp. 21-62). Hillsdale, NJ: Lawrence Erlbaum.

Smith, L. (1995). Self-organizing processes in learning to learn words: Development is not induction. The Minnesota Symposia on Child Psychology, 28, Basic and applied perspectives on learning, cognition, and development (pp. 1-32). Mahwah, NJ: Erlbaum

Smith, L., Jones, S. & Landau, B. (1996). Naming in young children: A dumb attentional mechanism? Cognition 60, 143-171.

Snedeker, J. (2000). Cross-situational observation and the semantic bootstrapping hypothesis. To appear in E. Clark (ed.), Proceedings of the Thirtieth Annual Child Language Research Forum. Stanford, CA: Center for the Study of Language and Information.

Snedeker, J. & Gleitman, L. (in press). Why it is so hard to label our concepts. To appear in Hall, G. & Waxman, S. (Eds).

Tardif, T. (1996). Nouns are not always learned before verbs: Evidence from Mandarin speakers' early vocabularies. Developmental Psychology, 32, 492-504.

Tardif, T., Gelman, S. & Xu, F. (1999). Putting the "noun bias" in context: A comparison of English and Mandarin. Child Development, 70, 620-635.

Waxman, S. (1999). Specifying the scope of 13-month-olds' expectations for novel words. Cognition, 70, 1999, 35-50.

Woodward, A.L. (1992). The role of the whole object assumption in early word learning. Dissertation, Stanford University.

Authors' Note

We had the advantage of discussions with Jerry Fodor, Dedre Gentner, Henry Gleitman, Ed Kako and Michael Kelly in developing the ideas in this work. We also thank the members of the Cheese Seminar who served as a sounding board throughout and made important experimental and theoretical contributions. We gratefully acknowledge the practical assistance offered by Henry Gleitman, who drew the explanatory cartoon (Figure 2), and Jane Gillette, who collected and edited the videotaped experimental stimuli. We were aided and abetted by Jamie Schuller, Bjoern Hartmann, Le'eat Sharoni, Deborah Small, Jelena Voijinovic, Michel Walker, and Dan Cunningham. This work was supported by an National Science Foundation Center Grant to the Institute for Research in Cognitive Science (SBR-89-20230), by an National Institute for Health Grant to John Trueswell and Lila Gleitman (1-R01-HD3750707-01), and a Predoctoral Fellowship from the Institute for Research in Cognitive Science to Jesse Snedeker.

Table 1

Percentage of Subjects Identifying each Target Word in Experiment 1

Nouns		Verbs	
Word	% correct	Word	% correct
plane	78.60%	throw	53.60%
swing	75.00%	come	46.40%
drum	71.40%	push	42.90%
elephant	71.40%	play	28.60%
pig	71.40%	hammer	21.40%
ball	67.90%	look	21.40%
bag	46.40%	see	17.90%
hammer	25.00%	wait	17.90%
hole	25.00%	put	10.70%
nose	17.90%	love	7.10%
hand	14.30%	fell	3.60%
tail	14.30%	go	3.60%
music	10.70%	like	3.60%
people	10.70%	stand	3.60%
peg	7.10%	do	0.00%
toy	7.10%	get	0.00%
hat	3.60%	have	0.00%

kiss	3.60%	know	0.00%
pilot	3.60%	make	0.00%
camera	0.00%	pop	0.00%
daddy	0.00%	say	0.00%
mommy	0.00%	think	0.00%
shoes	0.00%	turn	0.00%
things	0.00%	want	0.00%
Mean	26.00%	Mean	11.80%

Footnotes

[1] Though noun dominance has been documented in the early vocabularies of children learning a variety of languages, countervailing variables in language typology, maternal usage, and culture also clearly affect the contents of children's vocabularies (Gentner & Boroditsky, 2001). In addition, different data collection methods can lead to substantial differences in estimates of vocabulary composition (Tardif, Gelman & Xu, 1999).

[2] The number of clips per target was 3.1 for nouns and 3.6 for verbs.

[3] By telling our subjects that half the targets were nouns and half were verbs, we provided them with more information than is available to the novice word learner. This was done to ensure that any differences between our results and those of GGGL could be attributed to the subjects' knowledge of the syntactic category of target rather than to differences in the subjects' beliefs about the population of target words. We conducted a pilot study in which subjects were simply told that the words were nouns and verbs. The results of this pilot were quite similar to those of Experiment 1.

[4] This coding criterion guaranteed that any differences between the results of Experiment 1 and GGGL would not be attributable to the additional demand we placed on our subjects by asking them to identify the category of the target. By coding these responses as correct, we raised performance on noun targets by 5% and performance on verb targets by 2%.

[5] While there was a small but reliable interaction between Trial and Syntactic Category in the subjects analysis, it was not reliable in the items analysis suggesting that the improvement across trials varied substantially for items from the same syntactic category $F1(6,468) = 2.66$, $p < .05$; $F2(6,252) = .83$, $p > .5$.

[6] Given input with this structure you might expect child language learners to mistakenly map verbs to incorrect meanings, rather than simply failing to learn them at all. This issue is addressed in the General Discussion.

[7] Given the absolute difference in performance between nouns and verbs, one might wonder whether this apparent interaction could be attributed a main effect of Experiment that was multiplicative rather than additive. We discarded this hypothesis for two reasons. First, the difference in improvement between nouns and verbs is much greater than even a multiplicative model would predict (the proportional increase for nouns is 74% while for verbs it is only 30%). Second, there is no reason to believe that the effect of Experiment would depend upon the number of words that were identified before the additional information was added. In fact you might just as easily suppose that this improvement would be proportional to the number of words that had not already been learned on the basis of scenes alone.

[8] The effect of Syntactic Category of Response was highly reliable ($F(1,78) = 68.77$, $p < .001$) is a subject ANOVA.

[9] Responses were classified as BLOCs and NonBLOCs by the author (guided by the examples given in the literature (Rosch et al., 1976; Nelson, Hampson & Shaw, 1993). The proportion correct for each subject was separately summed for BLOC and Non-BLOC responses. The nine empty cells (all BLOCs) were replaced with the mean of the other cell for that subject. The effect of Response Type was highly reliable ($F(1,78) = 109.44$, $p < .001$).

[10] One subject made no verb responses, the empty cell was replaced with the average certainty rating for the subject’s noun responses.

[11] For the single subject who made no verb responses, we again substituted in the mean percent correct for the noun responses.

[12] This was done because the vast majority of these cross-category homonyms were verb responses to noun targets. Since the bias manipulation decreased the number of verb responses, we expected that these homonymous verbs would occur less frequently in Experiment 3. Thus if we included these words we would be systematically inflating the relative performance of subjects in Experiment 1. Reanalyzing the data from Experiment 1 in this way does not alter the pattern of effects discussed earlier.

[13] There are two additional advantages of an object bias that were not tested in Experiment 3. First, the presence of a bias should decrease the difficulty of the learning task. Even if an object bias did nothing to help the learner identify the correct meaning of the word, it still might be less taxing than open minded consideration of all the possibilities (Markman, 1989). Next, the input contains more types of nouns than types of verbs (Gentner, 1981). By presenting equal numbers of targets from each category we may have underestimated the effect of enhanced noun performance on overall success

[14] Other statistical properties of the input database, beyond the co-occurrence and syntactic properties just mentioned, could contribute to the fall of the object bias over time. For example, it could be that the learner has essentially used up the frequent vocabulary of object nouns to label the familiar things in his everyday world, in which case the success rate for object responses would begin to decline.