Published in final edited form as:

Cogn Psychol. 2011 September ; 63(2): 93–106. doi:10.1016/j.cogpsych.2011.06.002.

Linking sounds to meanings: Infant statistical learning in a natural language

Jessica F. Hay¹, Bruna Pelucchi^{2,3}, Katharine Graf Estes⁴, and Jenny R. Saffran²

¹University of Tennessee – Knoxville; Department of Psychology

²University of Wisconsin – Madison; Department of Psychology and Waisman Center

³University of Ferrara, Italy; Department of Evolutionary Biology

⁴University of California, Davis; Department of Psychology

Abstract

The processes of infant word segmentation and infant word learning have largely been studied separately. However, the ease with which potential word forms are segmented from fluent speech seems likely to influence subsequent mappings between words and their referents. To explore this process, we tested the link between the statistical coherence of sequences presented in fluent speech and infants' subsequent use of those sequences as labels for novel objects. Notably, the materials were drawn from a natural language unfamiliar to the infants (Italian). The results of three experiments suggest that there is a close relationship between the statistics of the speech stream and subsequent mapping of labels to referents. Mapping was facilitated when the labels contained high transitional probabilities in the forward and/or backward direction (Experiment 1). When no transitional probability information was available (Experiment 2), or when the internal transitional probabilities of the labels were low in both directions (Experiment 3), infants failed to link the labels to their referents. Word learning appears to be strongly influenced by infants' prior experience with the distribution of sounds that make up words in natural languages.

1. Introduction

Before infants utter their first words, they face the daunting task of finding those words in continuous speech. This is not a trivial challenge, as words are rarely produced in isolation [1, 2] and there are no fully reliable markers of word boundaries in fluent speech [3, 4]. However, a substantial literature has demonstrated that infants learn about many of the cues that indicate word boundaries very early in development (for an overview, see [5]). While this research has revealed impressive learning in young infants, much of it has employed highly simplified language stimuli that are quite different from the rich and complex language input infants typically receive. This literature is also largely based on measures of infants' listening preferences. While this method is very useful for assessing infants' ability to discriminate between multiple test items, it gives little indication of how infants might use their early knowledge of sound sequences during the course of language acquisition.

© 2011 Elsevier Inc. All rights reserved.

Corresponding author: Jessica F. Hay Department of Psychology University of Tennessee – Knoxville 1404 Circle Drive Knoxville, TN 37996 Office: 865-974-3329 Fax: 865-974-3330 jhay@tennessee.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

In the present experiments, we investigated infants' ability to detect word boundary cues in complex natural language input, and to use this information to solve a basic task in language acquisition: linking the sounds of words with their referents. The extensive literature on word segmentation indicates that infants are sensitive to distributional patterns in their native languages. By their first birthday, infants have discovered numerous language-specific regularities in the ambient input, including how sounds pattern together (i.e., phonotactic patterns; [6, 7]), how sounds change according to their context (i.e., allophonic variations; [8, 9]), and how they group into rhythmic units (i.e., prosodic patterns; [10, 11, 12, 13]). Language-specific knowledge begins to guide infants' ability to find words in continuous speech by 9 months of age, and possibly earlier (e.g., [14, 7, 15, 16]).

Infants are also sensitive to distributional regularities that function independently of specific language experience, such as the likelihood that two syllables will co-occur (e.g., [17, 18]). At word boundaries, the likelihood of syllable co-occurrence tends to be quite low, whereas the likelihood of syllable co-occurrence within words tends to be much higher (e.g., [19]). Co-occurrence likelihood signals statistical coherence or internal predictability, is frequently referred to as the transitional probability (TP) between syllables, and is expressed as the likelihood of Y given X as a function of the frequency of the co-occurrence of XY :

$$Y|X = \frac{[XY]}{[X]} \quad (1)$$

Eight-month-old infants are sensitive to differences in transitional probability, and discriminate sound sequences with high internal predictability (typically referred to as *words*) from sound sequences with low internal predictability (or *part-words*) and from unattested sound sequences (i.e., sequences with zero predictability, also called *non-words*) following very brief familiarization with an artificial language [17, 18].

Studies of statistical learning have primarily been done using artificial languages and have revealed powerful learning mechanisms in young infants. Artificial languages have allowed researchers to isolate variables of interest and control the precise experience infants receive. However, the cost of this control is that artificial languages provide greatly simplified linguistic input. They tend to contain few words (typically, just four to six), have a limited set of phonemes and syllables, and lack other sequential regularities, rhythmic patterning, pitch changes, and other acoustic variability associated with natural languages. For these reasons, the learning challenges presented by artificial languages are quite different than those presented by natural languages. In order to address this concern, researchers have systematically increased the complexity of artificial languages used in infant studies in several ways. For example, researchers have included naturally produced syllables instead of synthetically produced syllables [20, 21], varied word length [22, 23, 24] and stress cues [14, 16, 25], and provided multiple correlated cues to word boundaries [21]. Generally, these studies have found that infants are still able to track statistical regularities in the input despite added complexity (although see [22], for a counter-example), and that in some cases additional complexity facilitates statistical learning (e.g., [24]). However, while these miniature languages contain more structure than earlier artificial materials, they continue to lack the richness of natural languages.

Recently, however, Pelucchi and colleagues [26, 27] demonstrated that English-learning 8-month-old infants track transitional probabilities even when presented with complex language input that is naturally produced, grammatically correct, and semantically meaningful. Importantly, the natural language was not English (which English learning infants may be able to segment using known words [28]), but instead was an unfamiliar

natural language – Italian. Italian and English share a characteristic strong/weak stress patterning for disyllables [29, 30, 31], but differ in their allophonic and phonotactic regularities as well as in other rhythmic properties [32]. These Italian corpora maintained virtually all of the complexities found in natural language, but the internal transitional probabilities between syllable sequences were expressly manipulated in a subset of the words. Despite the complexity of the natural language input, 8-month-old infants discriminated high vs. low transitional probability words following a very brief familiarization period.

Pelucchi et al.'s [26, 27] demonstration of statistical learning in a natural language allowed for greater ecological validity than previous experiments using artificial languages. However, these results – successful discrimination between test items that differed in transitional probability – tell us little about the representations that infants formed while listening to the fluent speech. For example, does the output of the segmentation task provide infants with something akin to words that are useful for subsequent language development? Infants may be able to use word segmentation processes to identify sound sequences that are likely to be individual words and thus ready to be associated with meanings. Alternatively, the representations that emerge from this statistical learning process may fail to intersect with subsequent word learning. On this latter view, statistical learning mechanisms, while available to infant language learners faced with fluent speech, do not render representations that can serve as potential words, thereby limiting the explanatory power of statistical learning accounts.

To begin to address these issues, Graf Estes et al. [20] designed a method to study the connection between word segmentation and word learning. In this experiment, 17-month-old infants were first familiarized with an artificial language. The language consisted of four disyllabic words, recorded so that the transitional probabilities between syllables were the only reliable cues to word boundaries. Infants were then presented with a word-object association task. They were taught two novel label-object associations where the labels were words from the artificial language (TP = 1.0), nonwords (sequences that did not appear in the language; TP = 0.0), or part-words (sequences that crossed word boundaries in the artificial language; TP = 0.5). Following habituation, infants were tested using a modified Switch paradigm [33]. Half of the test trials were Same trials, in which the label-object pairings presented during habituation remained unchanged. The other half of the test trials were Switch trials, in which the object and label pairings were switched such that Object A was paired with Label B and vice versa. The results confirmed that the distributional information presented in the initial speech stream affected infants' subsequent word learning success. Infants who were taught labels that were words in the speech stream showed significantly longer looking to the Switch trials than to the Same trials, indicating that they had learned the label-object pairings. However, infants who were taught labels that were nonwords or part-words from the language did not show differential looking to the two types of test trials. Only the high transitional probability sequences appeared to function as good candidate words. This pattern of results suggests that not only are infants able to track statistical regularities in sound sequences, but also that the output of this process can function as the input to future word learning.

Word learning tasks, combined with word segmentation tasks, provide a window on what infants learn by tracking statistical regularities in speech streams, because infants can use the output of statistical learning to support subsequent label-object association learning. One learning task feeds the next. However, like most statistical learning experiments, Graf Estes et al. [20] used simple artificial language materials with native language phonology. It is not yet clear whether infants tracking regularities in natural languages are also able to form representations of sound sequences that are ready to link to meanings.

The present study examined whether infants can use the output of statistical learning in a natural language as the input to word learning. To that end, we combined the method developed by Graf Estes and colleagues [20] with the Italian materials developed by Pelucchi et al. [26]. In Experiment 1, 17-month-old infants were first familiarized with an Italian corpus of fluent speech, followed by a label-object association task. The crucial manipulation concerned the statistics of the labels relative to the Italian corpus. Half of the infants were trained on label-object pairings in which the labels were words from the corpus with a high internal transitional probability (**HTP condition**). The other half of the infants were trained on label-object pairings in which the labels had a low internal transitional probability (**LTP condition**). Importantly, both types of words occurred equally often in the Italian corpus. Based on the Graf Estes et al. [20] results with artificial language materials, we predicted that the infants for whom the labels were high transitional probability words would show more successful word learning than infants for whom the labels were low transitional probability words.

2. Experiment 1

2.1. Introduction

In Experiment 1, we investigated how differences in word-internal transitional probability patterns from a word segmentation task affected infants' ability to learn new object labels. Critically, all target words presented in the word segmentation task were matched for frequency of occurrence and stress pattern. They differed only in their internal transitional probabilities: high transitional probability (HTP) words had an internal transitional probability of 1.0 and low transitional probability (LTP) words had an internal transitional probability of 0.33. We expected infants to have difficulty associating LTP words with novel objects, because of their lower internal statistical consistency.

2.2. Method

2.2.1. Participants—Forty 17-month-old monolingual English-learning infants (mean age 17.6 months, range: 17.2–18.0) were randomly assigned to either the HTP or the LTP condition. All infants were born full-term, had fewer than five prior ear infections, no history of hearing or vision impairments, and no exposure to Italian or Spanish. Infants were recruited from a local birth announcement database maintained by the Waisman Center at the University of Wisconsin–Madison. Data from 15 additional infants were excluded from analysis because of fussiness (8), experimenter error (1), or not paying attention (5). One additional infant showed a looking time preference > 3 SD from the mean and was excluded from the analyses.

2.2.2. Stimuli

2.2.2.1. Word Segmentation Task: Materials were identical to those used in a previous word segmentation study [26]. They were recorded in a lively prosody by a native Italian-speaking female who was unaware of the purpose of the study. Four trochaic Italian target words (*bici*, *casa*, *fuga*, *melo*) were included in a set of 12 naturally produced, semantically meaningful and grammatically correct Italian sentences. Each target word occurred 6 times within the corpus. Two of the target words (HTP words) had high internal transitional probabilities [1.0; see Eq. (1)] because their component syllables never appeared elsewhere in the corpus. The transitional probabilities of the other two target words (LTP words) were lowered to 0.33 by including 12 further occurrences of their first syllable within the corpus. Thus, the first syllable of each LTP word appeared 18 times within the corpus, only 6 of which were in the target words (see Table 1). The average transitional probability at word boundaries was comparable for the two word types: HTP = 0.26; LTP = 0.27. The corpus was repeated 3 times during familiarization, for a total of 18 presentations of each target

word. To control for possible arbitrary listening preferences, we created two counterbalanced sets of sentences, Language A and Language B, in which the HTP and the LTP words were switched (see Appendix A for sentence lists). The total duration of each corpus was 2 min 15 sec.

2.2.2.2. Label-object association task: Infants received one of two sets of counterbalanced label-object pairs. In the HTP condition, the labels were the two HTP words from the segmentation task (e.g., *melo* and *fuga* in Language A). In the LTP condition, the labels were the two LTP words (e.g., *bici* and *casa* in Language A). The combination of the two counterbalanced sets of sentences (Language A and Language B) with the between-subjects manipulation of the transitional probability of the labels (HTP condition and LTP condition) led to four different testing conditions. The labels were recorded in citation form, with relatively monotone pitch contours. The novel objects were the same for all conditions and consisted of the two computerized 3-D images utilized by Graf Estes et al. [20].

2.2.3. Procedure—The procedure was identical to Graf Estes et al. [20]. Infants were seated on a parent's lap in a sound attenuated booth, approximately 1 meter from a flat screen monitor. Infants first listened to one of the two Italian corpora (Language A or Language B) for 2 min 15 sec while watching a silent cartoon. Immediately following familiarization, infants participated in the label-object association task. The stimuli were presented using Habit X 1.0 [34]. An observer viewed infants' responses on a monitor and indicated looking times by pressing a button on the computer running Habit. To avoid potential bias, the observer was blind to the identity of the materials being presented, and the parent listened to masking music over headphones.

Infants were first habituated to two label-object pairs, presented one at a time, in random order. On each trial, the object moved across the computer screen while its label played from a speaker located beneath the screen. This label-object combination played continuously while the infant looked at the screen, and terminated when the infant looked away for 1s or after a maximum looking time of 20s. The habituation criterion was satisfied when looking time across 3 consecutive trials decreased to 50% of the average looking time across the first 3 trials.

Test trials began immediately after the infant habituated or reached the maximum cutoff of 25 trials. During Same test trials, the infant viewed the label-object combinations that were presented during the habituation phase. During Switch trials, the labels for the two objects were switched; for example, Object B occurred with Label A. There were four Same and four Switch trials, in two counterbalanced testing orders. The dependent variable was the looking time difference for Same versus Switch test trials (Switch looking time – Same looking time). If infants learned the original label-object pairings, they should look longer during the Switch test trials in which the pairings are violated (e.g., [33]).

2.3. Results and Discussion

We first compared the looking-time differences for Same and Switch test trials for infants who heard the two counterbalanced familiarization languages (Language A and Language B). There were no significant looking-time differences between languages (all t-tests 2-tailed; all effect sizes reported are Cohen's d^1 [35]) in either the HTP condition [$t(18) = 1.18$, $p > .05$, $d = .37$] or LTP condition [$t(18) < 1.0$, $p > .05$, $d = .05$]. Therefore, performance was collapsed across the two languages in the subsequent analyses.

¹Cohen's d [35] was calculated using a pooled standard deviation for independent samples tests and an average standard deviation that corrects for the dependence between the means for paired-samples tests [36].

In both the HTP and LTP conditions, infants exhibited significantly longer looking times on Switch trials than on Same trials: HTP [$t(19) = 3.2, p < .05, d = .71$]; LTP [$t(19) = 4.3, p < .001, d = .98$] (see Figure 1). Sixteen of 20 infants in the HTP condition and 17 of 20 infants in the LTP condition looked longer during Switch trials. The mean number of trials to reach the habituation criterion did not differ across conditions [HTP: $M = 12.3, SD = 5.2$; LTP: $M = 12.4, SD = 5.1$; $t(38) < 1.0, p > .05$], nor did the total time to habituation [HTP: $M = 173.6, SD = 78.6$; LTP: $M = 141.9, SD = 91.0$; $t(38) < 1.0, p > .05$].

These results suggest that after exposure to fluent Italian sentences, infants can successfully learn to treat words from those sentences as labels for novel objects. The transitional probabilities internal to the labels did not appear to affect the learning outcomes; infants successfully mapped both HTP (TP = 1.0) and LTP (TP = 0.33) words as labels for objects. These results are surprising in light of the findings of Graf Estes et al. [20], who demonstrated that infants failed to learn label-object pairings when the labels comprised low probability sound sequences (Exp. 1 TP = 0.0; Exp. 2 TP = 0.5).

One potential explanation for the current findings is that the LTP words derived a boost from syllable-level familiarity; the first syllable of the LTP labels occurred more often throughout the corpus than did any of the other target words' syllables. We are not able to explicitly rule out this possibility given the constraints of manipulating transitional probability in natural language materials, while controlling for the frequency of the target words. However, Pelucchi and colleagues [26] provide a convincing demonstration that syllable-level familiarity does not drive 8-month-old infants' familiarity preferences for HTP-words over novel words in a very similar natural Italian corpus.

A second potential explanation for the current findings is that performance on the label-object association task was unaffected by prior familiarization with the Italian fluent speech. By 17 months of age, infants are fairly skilled at learning new label-object associations, even in the absence of previous experience with the labels in fluent speech (e.g., [37]). It is thus possible that infants in Experiment 1 did not need an extra boost from the word segmentation phase to successfully learn the labels. To rule out this possibility, we ran a control experiment in which infants completed the label-object association task from Experiment 1 without any initial familiarization with the Italian corpus. If infants can learn the label-object associations based solely on their experience during the habituation phase of the label-object association task, then infants should continue to succeed at learning the label-object pairings in Experiment 2.

3. Experiment 2

3.1. Introduction

In Experiment 2, we assessed infants' ability to map referents to words from an unfamiliar language without any previous experience with the language. Instead of being familiarized with the Italian corpus used in Experiment 1, infants listened to music before completing the label-object association task.

3.2. Method

3.2.1. Participants—Twenty 17-month-old monolingual English-learning infants (mean age 17.5 months, range: 17.1–17.9) participated in Experiment 2. The infants had the same characteristics as the sample from Experiment 1. Five additional infants were excluded due to fussiness (2), parental interference (1), or not paying attention (2).

3.2.2. Stimuli & Procedure—Infants first listened to a musical selection (*Eine kleine nacht musik*, W.A. Mozart) of the same duration to the corpus used in Experiment 1 (2 min

13 sec). Infants then completed the label-object association task used in Experiment 1. Half of the infants were trained with *casa* and *bici* as labels, and half were trained with *fuga* and *melo* as labels. All other stimuli and procedures were identical to Experiment 1.

3.3. Results and Discussion

We first compared the looking-time differences for Same and Switch test trials for infants trained on the two different pairs of words, *casa/bici* versus *fuga/melo*. There were no significant differences between test items: [$t(18) < 1.0, p > .05, d = .31$]. Therefore, performance was collapsed across the two groups in the subsequent analysis. A paired t-test revealed no significant difference between the looking times on Same vs. Switch trials: [$t(19) < 1.0, p > .05, d = .09$]. Ten of 20 infants looked longer during Switch trials (see Figure 2). Infants reached the habituation criterion after an average of 12.9 trials ($SD = 5.73$) and 158.2 s ($SD = 91.3$) of accumulated looking, which is comparable to Experiment 1 (HTP: 12.3 trials and 173.6 s; LTP: 12.4 trials and 144.0 s).

The results of Experiment 2 suggest that in the absence of prior experience, infants do not successfully map these labels to objects in the Switch task. Infants of this age (and younger) have been shown to readily associate novel labels with novel objects, even in the absence of prior experience with the labels (e.g., [37, 38, 39, 33]). In many of these studies, labels were produced in characteristic infant directed prosody. Although the speech streams used in the present experiments were spoken in a lively prosody, the labels used in the label-object association task were spoken as trochees with relatively monotone prosody. Graf Estes et al. [20] also found that infants failed to learn two-syllable sequences as object labels when the labels were presented in a monotone voice, in the absence of prior segmentation experience. However, Graf Estes [40] found that infants successfully learned the same labels when they were presented in infant-directed speech, even without prior segmentation experience. These considerations suggest that the Italian labels used in the current study, which were presented in citation form without infant-directed prosody, may be difficult to map to objects. Prior experience with these words in fluent speech may be particularly useful in such cases.

Taken together, Experiments 1 and 2 suggest that infants need the opportunity to segment the target words from fluent speech prior to the label-object association task in order to engage in successful word learning given these stimuli. An interesting puzzle emerges when comparing the label-learning patterns observed in the current study with the findings of Graf Estes et al. [20]. In both studies, infants learned labels following segmentation experience. In Experiment 1 of the current study, infants successfully mapped HTP and LTP Italian words as labels for objects. In contrast, Graf Estes et al. found that infants successfully learned high transitional probability sequences as object labels, but not low probability sequences that had occurred as part-words in an artificial language. One possible explanation for these diverging results is that the target words in Experiment 1 were trochees, whereas there were no prosodic cues present in the fluent speech used by Graf Estes et al. [20]. Thus, it is possible that our 17-month-old English-learning infants ignored the internal transitional probabilities of the target words and instead used a trochaic-based parsing strategy to segment words from the fluent speech.

A second possibility is that infants' success in both the HTP and LTP conditions in Experiment 1 (and failure in Experiment 2) was due to brief exposure to the sounds of Italian during familiarization, rather than to the specific sequential statistics exemplified in the corpora. Although the four Italian words used as labels are phonetically legal in English, familiarization with fluent Italian speech in Experiment 1 may have facilitated label-object learning by providing infants with exposure to the phonetic structure of Italian.

Another possible explanation concerns the directionality of the computation of the sequential statistics. In principle, transitional probability can be defined in two different and symmetrical ways, depending on how the normalization step is performed. In addition to forward transitional probabilities described in Eq. (1), it is also possible to compute backward transitional probabilities, as shown in Eq. (2), by measuring the likelihood of *X* preceding *Y*:

$$X|Y = \frac{[XY]}{[Y]} \quad (2)$$

Indeed, recent studies suggest that both adults [41, 42] and infants [27] can track backward transitional probabilities in fluent speech.

In the Italian corpora used in Experiment 1, we lowered the transitional probabilities within the LTP words by including additional occurrences of the first syllable of each of those words. For example, the transitional probability of *casa* was lowered (to 0.33) by including other words that contained the syllable *ca*. However, the backward transitional probability remained high, at 1.0 (see Table 1) because when *sa* occurred it was always preceded by *ca*. This manipulation suggests a potentially important difference from the low transitional probability part-word labels used by Graf Estes et al. [20], which contained low transitional probabilities in both directions (both forward and backward TPs = 0.5). It is thus possible that the high backward transitional probability of the LTP words played a crucial role in segmentation and subsequent word learning in Experiment 1.

To explore these issues, we designed a third experiment in which both forward and backward transitional probabilities were manipulated in the LTP labels. In the new counterbalanced corpora, the HTP items continued to contain high transitional probabilities (TPs = 1.0) in both directions, as in Experiment 1. However, both the forward and backward transitional probabilities for the LTP words were lowered to 0.33 by adding further occurrences of the words' first and second syllables elsewhere in the corpus (see Table 1).

If infants' success at mapping LTP words to objects in Experiment 1 was due solely to their ability to segment and associate meanings with trochees presented frequently in fluent speech, we would expect them to continue to successfully associate both the HTP and LTP labels with objects in Experiment 3. Similarly, if mere exposure to the phonetic structure of Italian was sufficient to facilitate object label learning, then we would expect infants to continue to successfully associate both the HTP and LTP labels with objects in Experiment 3. If, instead, the LTP words from Experiment 1 emerged as potential words due to their high backward transitional probabilities, the decreased backward transitional probabilities present in the LTP words in Experiment 3 should prevent infants from readily learning the labels.

4. Experiment 3

4.1. Introduction

In Experiment 3 we investigated infants' ability to learn object labels after manipulating both the forward and the backward transitional probabilities internal to the target words in fluent speech. The LTP words in Experiment 3 contained low internal forward and backward transitional probabilities. These materials thus diverged from the LTP words used in Experiment 1, which contained low transitional probabilities in the forward direction and high transitional probabilities in the backward direction.

4.2. Method

4.2.1. Participants—Forty 17-month-old monolingual English-learning infants (mean age 17.6 months, range: 17.0–18.0) participated in Experiment 3. All infants were recruited from the same pool and had the same characteristics as described in Experiment 1. Fourteen additional infants were excluded from the analysis because of fussiness (11), parental interference (1), or not paying attention (1). One additional infant showed looking time preferences > 3 SD from the mean and was excluded from the analyses.

4.2.2. Stimuli and Procedure—Materials were similar to those used in Experiment 1. There were 6 repetitions of each HTP word and 6 repetitions of each LTP word in the corpus. As in Experiment 1, the component syllables of two of the target words (HTP words) appeared nowhere else in the corpus, so that their forward and backward internal transitional probabilities were both 1.0. Critically, the transitional probabilities of the other two words (LTP words) were lowered to 0.33 in both directions by adding 12 further occurrences of their first syllable and 12 further occurrences of their second syllable (see Appendix B for sentence lists). The forward and backward transitional probabilities of the LTP words were thus both 0.33 (see Table 1). To control for possible arbitrary listening preferences, we created two counterbalanced sets of sentences, Language A and Language B, in which the HTP words and LTP words were switched. All other stimuli and procedures were identical to those described in Experiment 1.

4.3. Results and Discussion

We first compared the looking-time differences for Same and Switch test trials for infants who heard the two counterbalanced familiarization languages (Language A and Language B). There were no significant looking-time differences between languages in either the HTP [$t(18) < 1.0, p > .05, d = .01$] or the LTP condition [$t(18) = 1.34, p > .05, d = .42$]. Therefore, performance was collapsed across the two languages in the subsequent analyses.

Infants presented with HTP object labels exhibited significantly longer looking times on Switch than Same trials [$t(19) = 2.87, p < 0.05, d = .64$]. In contrast, infants presented with LTP labels showed no difference in looking time [$t(19) < 1.0, p > .05, d < .10$]. Fifteen of 20 infants in the HTP condition and 10 of 20 infants in the LTP condition looked longer on Switch trials (see Figure 3). The mean number of trials to reach the habituation criterion did not differ across conditions [HTP: $M = 11.9, SD = 5.0$; LTP: $M = 10.7, SD = 3.4$; $t(38) < 1.0, p > 0.05$]. The total time to habituation also did not differ across conditions [HTP: $M = 158.8, SD = 81.8$; LTP: $M = 127.5, SD = 45.3$; $t(38) = 1.50, p > 0.05$].

In our final analysis, we combined the results of the three experiments to ask whether words containing strong cues to word boundaries (TP of 1.0 in at least one direction: HTP words from Experiments 1 and 3; LTP words from Experiment 1) functioned as better labels for novel objects than words containing weak cues (TPs of 0.33 or below in both directions: words in Experiment 2; LTP words from Experiment 3). To do so, we performed a 2×2 ANOVA across all three experiments, examining the within-subjects effect of Trial Type (same vs. switch) and the between-subjects effect of Cue Strength (strong TP cues vs. weak TP cues) on looking time. There was a main effect of Trial Type [$F(1,98) = 20.76, p < .001$] revealing that when the data are collapsed across Cue Strength, infants looked longer to switch than to same trials. There was no main effect of Cue Strength [$F < 1, p > .35$]. Importantly, there was a significant Trial Type by Cue Strength interaction [$F(1,98) = 5.37, p < .05$]; Infants had more difficulty learning novel object labels when the labels had weak transitional probability cues [$t < 1, p > .9, d < .01$] than when the transitional probabilities between syllables were high in at least one direction [$t(59) = 5.71, p < .001, d = .74$].

Taken together, these results suggest that prior exposure to Italian fluent speech facilitated infants' ability to learn novel object labels, but only when the transitional probability between syllable sequences was high in at least one direction. When the forward and backward transitional probabilities were both lowered to 0.33, as in the LTP words used as labels in Experiment 3, infants failed to learn the label-object associations. This is despite the use of the favored trochaic stress pattern for all of the labels used in this experiment, suggesting that trochaic stress cues alone were insufficient to support segmentation and subsequent mapping of the LTP words as labels for novel objects, at least in this task.

5. General Discussion

In the current experiments, we manipulated the statistical information present in natural speech, and assessed the influence of these sequential statistics on subsequent word learning. The results suggest that the internal cohesiveness of novel words, as measured by the strength of their internal transitional probabilities, influence how readily infants map these words to novel objects. In Experiment 1, 17-months-old learned newly segmented words as labels for novel objects when their transitional probabilities were high in both directions (HTP words: forward and backward TPs = 1.0) and when their transitional probabilities were low in the forward direction, but high in the backward direction (LTP words: forward TP = 0.33, backward TP = 1.0). Experiment 2 ruled out the alternative hypothesis that the sequential statistics in the fluent speech stream were irrelevant to the word-learning task. Finally, in Experiment 3, infants again learned HTP words (forward and backward TPs both = 1.0), but failed to learn LTP words in which the internal transitional probabilities were low in both directions (forward and backward TPs both = 0.33). Across experiments, infants successfully learned new object labels after having the opportunity to segment the sound sequences from fluent speech. Importantly, the sequences that formed good labels had high internal transitional probabilities in at least one direction.

Our experiments offer further evidence that statistically-based word segmentation provides infants with adequate candidate labels for word learning, assuming that the statistics are sufficiently strong. These findings are consistent with Graf Estes et al.'s [20] findings, extending this line of research to include natural language materials, which are far closer to the kind of input that infants typically receive. Demonstrating that infants can track and use statistical regularities when faced with the complexity of natural language represents an important step in advancing our understanding of the role of statistics in natural language acquisition. Importantly, not only are infants sensitive to the internal transitional probabilities of words in natural speech, but they are also sensitive to both forward and backward transitional probability (see also [27]). Within words, low transitional probabilities in one direction appear not to hinder the process of segmenting words and associating them with objects, as long as high transitional probability information is available in the other direction. These results suggest that infants appear to possess a keen sensitivity to multiple statistical cues, including forward and backwards transitional probabilities. This flexibility may be useful for higher-level language acquisition, including learning grammatical gender, suffix, and determiner agreement, which rely on a subtle interplay between forward and backward transitional probability.

The present set of experiments also illuminates the importance of expanding beyond listening-time-based discrimination measures of learning. Most studies of word segmentation use preferential looking tasks in which infants must discriminate between sequences that differ in their relationship to an exposure corpus. For example, consider the study by Pelucchi et al. [26], Experiment 3, who used an exposure corpus identical to the current Experiment 1. Pelucchi et al. [26] found that 8-month-old infants can discriminate HTP words (forward and backward TPs both = 1.0) from LTP words (forward TP = 0.33,

backward transitional probability = 1.0) when they are pitted against each other using a typical preferential listening design. While these results demonstrated that infants discriminate between items with different internal transitional probability values, they did not give us any indication of the linguistic status of the LTP words. From the Pelucchi et al. [26] study, it might be reasonable to conclude that lowering the forward transitional probability of a word renders that word unsegmentable. However, our current results clearly demonstrate that these very same LTP words function as good candidate labels for objects, at least by the time that infants reach 17 months of age. Combining a segmentation task and a label-object association task permits a more nuanced look at the output of statistical learning.

It is important to note that although the internal transitional probability of target words provided infants with the only cue that they could have used to differentiate between HTP and LTP words, it was not the only cue to word boundaries in Experiments 1 and 3. All of the target words were trochees. It remains unclear whether infants would have been successful at this task had the stress pattern of the target words not provided a converging cue to word boundaries. Thus, we are not claiming that sequential statistics alone provide sufficient information for segmentation of natural speech, but rather that when multiple cues to word boundaries converge, sequential statistics are relevant (see also [21]). In natural languages, stress cues, sequential statistics, and other cues to word boundaries are not typically pitted against each other, but instead should work in concert to help infants segment words from fluent speech. It will be of great interest to replicate these experiments with a design in which stress and statistical cues are placed in conflict, for example by substituting trochees with iambs in the corpus.

Outside of the laboratory, infants are faced with highly complex input when learning their native language. Infants must not only detect words in noisy input, but they must also link those words to objects and events in their environments. By using a natural language to study the role of sequential statistics in the mapping of newly segmented words to meaning, we have more closely mimicked both the stimuli and the tasks that infants are faced with when learning their native languages. Future research will continue to explore other dimensions of complexity in natural speech. For example, in the current study, a transitional probability of 0.33 was considered to be “low”, particularly in comparison to other words in the corpus with perfect 1.0 transitional probabilities. But in natural languages, 0.33 is likely to be quite high compared to the surrounding statistical landscape. By investigating how infants’ sensitivity to sequential statistics influences segmentation, and how the output of segmentation influences downstream learning processes, research in this area has the potential to more clearly reveal the links between hypothesized learning mechanisms and the role they actually play in language acquisition.

Acknowledgments

This research was funded by grants from NICHD to JFH (F32-HD557032) and to JRS (R01HD37466), a grant from the James F. McDonnell Foundation to JRS, and by a NICHD core grant to the Waisman Center (P30HD03352). We would like to thank the members of the Infant Learning Lab and especially Rose Maier and Jessica Rich for their assistance in the conduct of this research, the families who generously contributed their time, and Alexa Romberg, Erik Thiessen, and three anonymous reviewers for helpful comments on a previous version of this manuscript.

Appendix A

Language 1A

Torno a casa con le bici cariche di frutta in bilico sulla sella.

La zia Carola si e` esibita in una fuga colla bici verde.
 Se porti il melo sulla bici forse cali un po' di chili.
 La bici ha subito un danno dentro la casa del capo di Lara.
 La cavia Bida è in fuga da casa per aver giocato con le bilie blu.
 La biscia in lenta fuga dal giardino capita in casa mia.
 Il tuo melo arcano fuga l'afa che debilita la folla.
 Arriviamo in bici fino al bivio del grande melo con un caro amico.
 Il picchio si abitua a fare la sua casa in ogni melo cavo e alto.
 Gusto i bigoli dentro casa o coricata all'ombra del melo verde.
 Di rado una bici in rapida fuga rincorre la moto bigia e rossa.
 Per ascoltare la fuga quasi cadi sul melo e inciampi sulla biro sull'erba.

Language 1B

Non è da me scendere dal melo in una futile fuga dalle api.
 Torno a casa dalla futa con la bici piena di mele mature.
 Il melo e diverse bici furono portate presso la mescita di vino.
 Zio Luigi Medo è in fuga colla bici verde.
 Vi fu l'età dei tentativi di fuga in bici verso il rifugio del melo antico.
 Il fu Romero Rossi temeva di andare in gita colla bici nuova.
 Dario fu l'ingenuo che portò una bici a casa il mese scorso.
 Una fuga da casa è il sogno della topina Mela verso la libertà.
 Il ratto Meco tentò la fuga da casa quando vi fu la tempesta.
 Il micio Refuso medita in casa o dimena la coda sotto al melo ombroso.
 Sui rami del melo che sembrano fusi c'è la casa del fuco solitario.
 La fuga della stella cometa si è fermata sul melo che fu della zia.

Appendix B

Language 2A

Spesso Lisa capita in fuga nella casa dove giaci gracile e tesa.
 Se cadi con la bici prima del bivio del melo cavo ti do dieci bigoli e una biro.
 Gli amici della cavia Bida poggiano le bici in bilico presso il melo per difesa dalla biscia.

Sovente carico la spesa nel vicinato dopo una fuga con la bici nuova.
 Carola si è esibita in una fuga verso il melo perché offesa dagli amici scortesì.
 Se vai a casa in bici ti debiliti ma cali e non sei più obesa.
 Dietro la casa del capo ho sprecato i ceci sotto al melo ombroso.
 Se cucì subito sulla divisa bigia il distintivo col melo vado in casa a dormire.
 Teresa si abitua alla fuga da casa con la vecchia bici senza luci posteriori.
 Taci sulla fuga di Marisa con il caro lattaio.
 Il bel melo sta tra la casa dei Greci e la chiesa arcana dove hai giocato con le bilie.
 I soci della ditta Musa si danno alla fuga con la bici della maglia rosa.

Language 2B

Roméro fu coinvolto in una futile fuga in bici verso il profumo del mélo ombroso.
 Il collega di Paolo Fusi trovò la bici per la fuga presso la casa del molo.
 La maga tiene in casa almeno un fuco, uno squalo e una tartaruga del Nilo.
 Il fuco procede parallelo alla casa sulla riga tracciata dalla cometa.
 Il gattone Refuso medita sul mélo presso casa ascoltando una fuga di Verdi.
 Il fu Medo Rossi rompe la braga nella bici il mese scorso durante la gara.
 Giga ogni mese paga con zelo l'affitto per la casa con il melo in fiore.
 meco prega il cielo che ogni fuga da casa termini sotto melo ombroso.
 Il delfino beluga si dimena tutto solo nella fuga verso il Nilo azzurro.
 Un pezzo di filo si è infilato nella bici appoggiata al melo dietro la méscita.
 Vi fu un tempo in cui la bici in lega non temeva il gelo del rifugio della Futa.
 La strega del melo fu vista in fuga sulla bici con un chilo di rametti.

References

1. van de Weijer, J. MPI Series in Psycholinguistics. Nijmegen; The Netherlands: 1998. Language input for word discovery..
2. Aslin, RN.; Woodward, JZ.; LaMendola, NP.; Bever, TG. Models of word segmentation in fluent maternal speech to infants.. In: Morgan, JL.; Demuth, K., editors. Signal to Syntax. LEA; Mahwah, NJ: 1996. p. 117-134.
3. Jones, D. An Outline of English Phonetics. 8th edition. Heffer & Sons; Cambridge: 1957.
4. Liberman AM, Cooper FS, Shankweiler DP, Studdert-Kennedy M. Perception of speech code. Psychological Review. 1967; 74:431–461. [PubMed: 4170865]

5. Saffran, JR.; Werker, JF.; Werner, LA. The infant's auditory world: Hearing, speech, and the beginnings of language.. In: Siegler, R.; Kuhn, D., editors. Handbook of child development. Vol. 6. Wiley; New York: 2006. p. 58-108.
6. Friederici AD, Wessels JM. Phonotactic knowledge and its use in infant speech perception. *Perception & Psychophysics*. 1993; 54:287–295. [PubMed: 8414887]
7. Mattys SL, Jusczyk PW. Phonotactic cues for segmentation of fluent speech by infants. *Cognition*. 2001; 78:91–121. [PubMed: 11074247]
8. Jusczyk PW, Hohne EA, Bauman A. Infants' sensitivity to allophonic cues for word segmentation. *Perception & Psychophysics*. 1999; 61:1465–1476. [PubMed: 10598463]
9. Christophe A, Dupoux E, Bertoncini J, Mehler J. Do infants perceive word boundaries? An empirical study of the bootstrapping of lexical acquisition. *Journal of the Acoustical Society of America*. 1994; 95:1570–1580. [PubMed: 8176060]
10. Jusczyk PW, Cutler A, Redanz NJ. Infants' preference for the predominant stress patterns of English words. *Child Development*. 1993; 64:675–687. [PubMed: 8339688]
11. Jusczyk PW, Houston DM, Newsome M. The beginnings of word segmentation in English-learning infants. *Cognitive Psychology*. 1999; 39:159–207. [PubMed: 10631011]
12. Morgan JL. A rhythmic bias in preverbal speech segmentation. *Journal of Memory and Language*. 1996; 35:666–688.
13. Polka, L.; Sundara, M.; Blue, S. The role of language experience in word segmentation: A comparison of English, French, and bilingual infants.. Paper presented at the 143rd Meeting of the Acoustical Society of America: Special Session in Memory of Peter Jusczyk; Pittsburgh, PA. 2002, June 3-7;
14. Johnson EK, Jusczyk PW. Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language*. 2001; 44:548–567.
15. Mattys SL, Jusczyk PW, Luce PA, Morgan JL. Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*. 1999; 38(4):465–494. [PubMed: 10334878]
16. Thiessen ED, Saffran JR. When cues collide: Use of stress and statistical cues to word boundaries by 7- to 9-month-old infants. *Developmental Psychology*. 2003; 39:706–716. [PubMed: 12859124]
17. Aslin RN, Saffran JR, Newport EL. Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*. 1998; 9:321–324.
18. Saffran JR, Aslin RN, Newport EL. Statistical learning by 8-month-old infants. *Science*. 1996; 274:1926–1928. [PubMed: 8943209]
19. Swingley D. Statistical clustering and the contents of the infant vocabulary. *Cognitive Psychology*. 2005; 50:86–132. [PubMed: 15556130]
20. Graf Estes K, Evans JL, Alibali MW, Saffran JR. Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological Science*. 2007; 18:254–260. [PubMed: 17444923]
21. Sahni SD, Seidenberg M, Saffran JR. Connecting cues: Overlapping regularities support cue discovery in infancy. *Child Development*. 2010; 81:727–736. [PubMed: 20573101]
22. Johnson EK, Tyler MD. Testing the limits of statistical learning for word segmentation. *Developmental Science*. 2010; 13(2):339–345. [PubMed: 20136930]
23. Mersad, K.; Nazzi, T. When mom helps statistics.. Poster presented at the ISIS Conference; Baltimore, MD. April, 2010; 2010.
24. Thiessen ED, Hill EA, Saffran JR. Infant-directed speech facilitates word segmentation. *Infancy*. 2005; 7(1):53–71.
25. Thiessen ED, Saffran JR. Learning to learn: Infants acquisition of stress-based strategies for word segmentation. *Language and Development*. 2007; 3:73–100.
26. Pelucchi B, Hay JF, Saffran JR. Statistical learning in a natural language by 8-month-old infants. *Child Development*. 2009; 80(3):674–685. [PubMed: 19489896]
27. Pelucchi B, Hay JF, Saffran JR. Learning in reverse: Eight-month old infants track backward transitional probabilities. *Cognition*. 2009; 113:244–247. [PubMed: 19717144]

28. Bortfeld H, Morgan JL, Golinkoff RM, Rathbun K. Mommy and me: Familiar names help launch babies into speech-stream segmentation. *Psychological Science*. 2005; 16:298–304. [PubMed: 15828977]
29. Mancini F, Voghera M. Lunghezza, tipi di sillabe e accento in italiano. *Archivio Glottologico Italiano*. 1994; 79(1):51–77.
30. De Mauro, T.; Mancini, F.; Vedovelli, M.; Voghera, M. *Lessico di frequenza dell'italiano parlato*. Etaslibri; Milano: 1993.
31. Cutler A, Carter DM. The predominance of strong initial syllables in the English vocabulary. *Computer Speech and Language*. 1987; 2:133–142.
32. Ramus R, Nespor M, Mehler J. Correlates of linguistic rhythm in the speech signal. *Cognition*. 1999; 73:265–292. [PubMed: 10585517]
33. Werker JF, Cohen LB, Lloyd VL, Casasola M, Stager CL. Acquisition of word-object associations by 14-month-old infants. *Developmental Psychology*. 1998; 34:1289–1309. [PubMed: 9823513]
34. Cohen, LB.; Atkinson, DJ.; Chaput, HH. *Habit X: A new program for obtaining and organizing data in infant perception and cognition studies (Version 1.0)* [Computer software]. University of Texas; Austin: 2004.
35. Cohen, J. *Statistical power analysis for the behavioral sciences*. 2nd ed.. Lawrence Earlbaum Associates; Hillsdale, NJ: 1988.
36. Morris SB, DeShon RP. Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*. 2002; 7:105–125. [PubMed: 11928886]
37. Werker JF, Fennell CT, Corcoran KM, Stager CL. Infants' ability to learn phonetically similar words: Effects of age and vocabulary size. *Infancy*. 2002; 3(1):1–30.
38. Curtin S. Twelve-month-olds learn word-object associations differing only in stress patterns. *Journal of Child Language*. 2009; 36:1157–1165. [PubMed: 19281635]
39. Mackenzie H, Graham SA, Curtin S. 12-month-olds privilege words over other linguistic sounds in an associative learning task. *Developmental Science*. 2011; 14(2):399–410.
40. Graf Estes, K. Learning words from infant-directed speech.. Paper presented at the International Conference on Infant Studies; Vancouver, British Columbia Canada. 2008, March;
41. Perruchet P, Desaulty S. A role for backward transitional probabilities in word segmentation? *Memory and Cognition*. 2008; 36:1299–1305.
42. Onnis, L. Statistical second language learning.. Paper presented at Georgetown University Round Table; Washington, DC. 2009, March 13–15;

Highlights 1

- We test the link between speech segmentation and subsequent label-object association.
- 17-month-olds were familiarized with an unfamiliar natural language.
- Infants mapped labels that contained strong internal statistics to novel objects.
- No mapping was observed when the internal statistics of the labels were weak.
- Speech stream statistics predict subsequent mapping of labels to referents.

Fig. 1. Experiment 1: mean looking times (± 1 SE) on Same and Switch test trials for infants in the HTP word and LTP word conditions.

Fig. 2. Experiment 2: mean looking times (± 1 SE) on Same and Switch test trials for infants in the control condition.

Fig. 3. Experiment 3: mean looking times (± 1 SE) on Same and Switch test trials for infants in the HTP word and LTP word conditions.

TABLE 1

Transitional probability values for the HTP and LTP words used in Experiment 1 and 3, in comparison with the “words” and “part-words” used in Graf Estes et al. (2007).

	High TP labels		Low TP labels	
	Forward TP	Backward TP	Forward TP	Backward TP
Experiment 1	1.0	1.0	0.33	1.0
Experiment 3	1.0	1.0	0.33	0.33
Graf Estes et al. (2007, Exp. 2)	1.0	1.0	0.5	0.5