

Lexical viability constraints on speech segmentation by infants

Elizabeth K. Johnson,^a Peter W. Jusczyk,^a Anne Cutler,^{b,*}
and Dennis Norris^c

^a *Departments of Psychology and Cognitive Science, Johns Hopkins University,
Baltimore, MD, USA*

^b *Max Planck Institute for Psycholinguistics, Nijmegen, Netherlands*

^c *Medical Research Council Cognition and Brain Sciences Unit, Cambridge, UK*

Accepted 4 April 2002

Abstract

The Possible Word Constraint limits the number of lexical candidates considered in speech recognition by stipulating that input should be parsed into a string of lexically viable chunks. For instance, an isolated single consonant is not a feasible word candidate. Any segmentation containing such a chunk is disfavored. Five experiments using the head-turn preference procedure investigated whether, like adults, 12-month-olds observe this constraint in word recognition. In Experiments 1 and 2, infants were familiarized with target words (e.g., *rush*), then tested on lists of nonsense items containing these words in “possible” (e.g., “nirush” [nip + rush]) or “impossible” positions (e.g., “prush” [p + rush]). The infants listened significantly longer to targets in “possible” versus “impossible” contexts when targets occurred at the end of nonsense items (*rush* in “prush”), but not when they occurred at the beginning (*tan* in “tance”). In Experiments 3 and 4, 12-month-olds were similarly familiarized with target words, but test items were real words in sentential contexts (*win* in “wind” versus “window”). The infants listened significantly longer to words in the “possible” condition regardless of target location. Experiment 5 with targets at the beginning of isolated real words (e.g., *win* in “wind”) replicated Experiment 2 in showing no evidence of viability effects in beginning position. Taken together, the findings suggest that, in situations in which 12-month-olds are required to rely on

* Corresponding author.

E-mail address: anne.cutler@mpi.nl (A. Cutler).

their word segmentation abilities, they give evidence of observing lexical viability constraints in the way that they parse fluent speech.

© 2003 Elsevier Science (USA). All rights reserved.

1. Introduction

The ease with which speech is understood tempts us to view word recognition as a relatively simple task: compare sound patterns of individual words against the set of form representations stored in the lexicon until an exact match is obtained, whereupon access to meaning results. However, the continuous nature of fluent speech poses potentially serious problems for identifying words in the acoustic signal. Fluent speech contains no analogue of the little white spaces between written words in text (Cole & Jakimik, 1980; Lehiste, 1972). Instead, the production of each spoken word blends into the next, obscuring the boundary where one word ends and another begins. The lack of clear pauses between words is most obvious when we hear speech in an unfamiliar language. Speakers sound as if they are talking very rapidly, running one word into the next. Yet speakers of other languages typically speak no faster than we do. Rather, the difficulty in parsing words from the unfamiliar language stems from the fact that the cues that we use to locate word boundaries in our native language are not necessarily the same as those which are the most reliable for segmenting words in other languages.

The fact that word boundary cues often differ from language to language raises some interesting issues with respect to how word segmentation abilities develop. Infants face essentially the same problems with utterances in their native language as do adult listeners with foreign language utterances—that is, they must discover the types of cues that mark word boundaries. Although it has sometimes been suggested that infants might circumvent the word segmentation problem by learning words in isolation, and then using their knowledge of particular words to recognize them and extract new words from longer utterances containing these words (e.g., Bloomfield, 1933; Brent & Cartwright, 1996; Pinker, 1984; Suomi, 1993), current evidence argues against this type of account. First, studies have indicated that input addressed to young language learners predominantly consists of multiword utterances (van de Weijer, 1998; Woodward & Aslin, 1990). Second, there is considerable evidence that infants begin to segment words from fluent speech as early as 7.5 months of age. For instance, Jusczyk and Aslin (1995) found that infants at this age who were familiarized with a pair of passages containing the target words *cup* and *dog* subsequently listened significantly longer to isolated productions of these words than to another pair, suggesting that they had segmented the target words

from the passages. Moreover, infants' ability to segment target words did not depend upon their prior knowledge of those words. In addition, infants' word segmentation abilities at this age allow them to extract not only monosyllabic words (Jusczyk & Aslin, 1995), but also bisyllabic (Jusczyk, Houston, & Newsome, 1999b), and even trisyllabic words (Santelman, Houston, & Jusczyk, 1997). Finally, 7.5-month-olds even display some ability to segment words under noisy conditions, such as when interference from a competing voice is present (Newman & Jusczyk, 1996).

Where then do these abilities come from? How do young infants *begin* to segment words from fluent speech? Certainly it appears that, during the second half of their first year, infants amass relevant knowledge about the sound organization of their native language. For example, between 6 and 9 months, infants develop sensitivity to prosodic and phonotactic characteristics of words in their native language. Thus although English-learning 6-month-olds do not display clear listening preferences for words with the predominant stress pattern of their native language (i.e., an initial strong syllable followed by one or more weak syllables), 9-month-olds favor words with the predominant stress pattern (Jusczyk, Cutler, & Redanz, 1993a). Similarly, at 6 months, English-learners are as content to listen to words consisting of sequences of phonetic segments which are impermissible within words of the native language as to words containing permissible sequences. However, by 9 months, not only do English-learners favor the words with the permissible sequences (Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk, 1993b), but they also favor those with frequently occurring phonotactic patterns over those with less frequent (but still permissible) patterns (Jusczyk, Luce, & Charles-Luce, 1994).

This knowledge is relevant for speech segmentation, because both sorts of information, prosody and phonotactics, have been proposed as potential word segmentation cues. Cutler and Carter (1987) noted the preponderance in English conversational speech of content words beginning with strong syllables. Subsequently, Cutler and Norris (1988) proposed the Metrical Segmentation Strategy (MSS), the notion that English listeners might use their knowledge of the predominant stress pattern of words as a first pass at locating the potential onsets of words in fluent speech. Considerable empirical evidence supports the view that, in languages such as English and Dutch (which has the same predominant word-stress pattern), listeners indeed follow the MSS in word segmentation (Cutler & Butterfield, 1992; Cutler & Norris, 1988; McQueen, Norris, & Cutler, 1994; Vroomen & de Gelder, 1995). Similarly, knowledge of native language phonotactic patterns has been suggested as another source of information about potential word boundaries in fluent speech (Brent & Cartwright, 1996; Cairns, Shillcock, Chater, & Levy, 1997; van der Lugt, 2001; McQueen, 1998; Vitevitch & Luce, 1998, 1999). For example, certain phonotactic patterns occur much more frequently between words than within the words of a language. Thus,

the appearance of such patterns in utterances signals the likely location of word boundaries. Once again, there is empirical evidence that adult listeners take advantage of such phonotactic patterns in segmenting and recognizing words (McQueen, 1998; Vitevitch & Luce, 1999). Finally, another source of information in the acoustic signal that potentially cues word boundaries resides in the positioning of certain allophones (Bolinger & Gerstman, 1957; Church, 1987; Hockett, 1955, Lehiste, 1960). Allophones of a given phoneme are often restricted with respect to their position within words. Thus, the aspirated allophone [t^h] of the English phoneme /t/ occurs at the beginning of stressed syllables, whereas the unaspirated allophone [t] is found in word-final position. Church (1987) has suggested that listeners might use information about the presence of such context-specific allophones in segmenting words from fluent speech.

Infants, too, appear to draw on all these potential sources of information when they begin segmenting words. Jusczyk et al. (1999b) investigated whether the preference that English-learners display for words with the predominant stress pattern might play a role in their segmentation of words in fluent speech. They familiarized 7.5-month-olds with words with strong/weak stress such as *kingdom* and *doctor*, or words with weak/strong stress such as *beret* and *guitar*. The infants were able to segment the strong/weak words correctly, but not the weak/strong words. In fact, consistent with what would be predicted if they were observing the MSS, the infants tended to mis-segment the weak/strong words at the onset of the strong syllable. Hence, they tended to segment *tar* from a fluent speech context that contained *guitar*. The tendency for English-learners to rely on the locus of strong syllables in word segmentation remains strong at 9 months (Echols, Crowhurst, & Childers, 1997; Morgan, 1996), so much so that they apply this strategy to segmenting words in an unfamiliar foreign language, Dutch (Houston, Jusczyk, Kuijpers, Coolen, & Cutler, 2000).

However, at this age they also appear to use phonotactic cues in segmenting words (Mattys & Jusczyk, 2001; Mattys, Jusczyk, Luce, & Morgan, 1999). In particular, 9-month-olds are better able to segment words from contexts that provide good phonotactic cues (Mattys & Jusczyk, 2001). Moreover, although English-learning 9-month-olds do not use allophonic cues to segment items such as *nitrates* and *night rates* from fluent speech contexts, English-learning 10.5-month-olds do (Jusczyk, Hohne, & Bauman, 1999a). Interestingly, at the latter age, when infants display sensitivity to both allophonic and phonotactic cues to word boundaries, they also show some ability to segment weak/strong words from fluent speech (Jusczyk et al., 1999b; Myers et al., 1996).

Therefore, the evidence suggests that infants' abilities to use different types of word segmentation cues that are available in the acoustic signal develop considerably between 7.5 and 10.5 months of age. Moreover, it is worth noting that these word segmentation abilities are developing well

before the point when infants have begun to link many sound patterns and their meanings in a lexicon. In fact, by 9 months, infants retain information about the sound patterns of frequently presented words in long-term memory for at least two weeks, even when no effort is made to teach them the meanings of these words (Jusczyk & Hohne, 1997). Thus, infants may not have to link the sound patterns of words with any particular meanings to begin storing the sound patterns in memory on their way to forming a lexicon.

Understanding how infants' early abilities for segmenting words develop into those of fluent speakers of a language requires the resolution of many issues. One of these has to do with how infants achieve a successful integration of the available sources of information that they have about potential word boundaries in fluent speech. Deciding on a likely location of a word boundary might not pose any special difficulty, if all available cues point to the same locus. However, the types of acoustic cues considered here—prosodic stress, phonotactics, and context-sensitive allophones—are probabilistic cues. None is completely sufficient for accurately segmenting the input. Hence, there will be many occasions in which the potential cues may conflict. For example, imagine infants hear the phrase *half tortoise*. Phonotactically, [ft] is common within words in English. However, [t] in *tortoise* will most likely be aspirated, and *tortoise* begins with a strong syllable. How do infants decide which cue is correct in such a situation? Some evidence suggests that, soon after English-learners begin to segment words, when prosodic stress cues conflict with other types of word boundary cues, infants tend to rely more heavily on the prosodic stress cues (Johnson & Jusczyk, 2001; Mattys et al., 1999). However, as noted above, by 10.5 months, English-learners do segment weak/strong words, suggesting that other types of cues can outweigh prosodic stress cues in some circumstances.

In models of adult word recognition, the issue of how listeners choose the appropriate segmentation of the input is resolved by assuming that two fundamental processes are involved: activation and competition (e.g., Luce & Pisoni, 1998; Marslen-Wilson, 1989; McClelland & Elman, 1986; Norris, 1994). The acoustic input activates some set of lexical candidates in memory.¹ Some models suggest that lexical candidates are activated only by certain portions of the signal, such as word beginnings (Marslen-Wilson, 1987, 1989). Yet, in English, words are often buried within other words, such as *grass* in *grasp* or *each* in *speech* (McQueen, Cutler, Briscoe, & Norris, 1995). In addition, spurious word forms that span existing word boundaries (such as *speech* in *grasp each*) may correspond to lexical items and lead to their activation in the candidate set (Gow & Gordon, 1995; Shillcock,

¹ Many of the details of the activation and competition process are still much debated, especially whether activation proceeds entirely in a bottom-up manner, based on information in the acoustic input or whether top-down (syntactic and semantic) information also is involved in activation (e.g., see Norris, McQueen, & Cutler, 2000, and in accompanying responses).

1990; Tabossi, Burani, & Scott, 1995; Vroomen & de Gelder, 1997). Consequently, other models have assumed that lexical candidates can be activated at any point in the speech stream (Luce, Goldinger, Auer, & Vitevitch, 2000; McClelland & Elman, 1986; Norris, 1994). Indeed, many words are apparently not properly identified until after their acoustic offset (Allopena, Magnuson, & Tanenhaus, 1998; Goldinger, Luce, & Pisoni, 1989; Luce et al., 2000). In any case, the competition among the activated candidates winnows the set to the word that is ultimately extracted and recognized.

Competition alone might often be sufficient to parse speech input successfully, assuming that all the words in the input were indeed in the listener's lexicon. But this assumption cannot be made for infant listeners; infants' lexicons consist at best of only a few words. Moreover, even adult listeners may be faced with input which includes items for which they have no lexical representation. Further research on adult word recognition has attempted to account for how everyday listeners successfully deal with unknown lexical items and rule out implausible segmentations of utterances containing such items. Norris, McQueen, Cutler, and Butterfield (1997) proposed that online segmentation and recognition of words is guided by a bias against considering any segmentation that leaves a single consonant as a residue.

Norris et al. initially proposed the constraint to account for the observation that listeners find it hard to detect words embedded in certain contexts. In a series of speeded word spotting experiments, Norris et al. found that adults more easily detected words such as *apple* in “vuffapple” than in “fapple” and *sea* in “seashub” than in “seash.” That is, the residual context determined whether the embedded word could be detected; if the residual context was a syllable, detection was easier than if the residual context was a single consonant. Norris et al. based a far-reaching claim on this finding. Syllables, they argued, have a special status in speech segmentation—not because they constitute an obligatory level of prelexical representation (indeed, there is evidence from a number of languages that this is not the case), but because they can be potentially words. Single consonants, on the other hand, cannot be words, in any language. They proposed, therefore, that their finding was evidence of a basic constraint in speech segmentation, which operates to rule out spurious word candidates activated by the input. Specifically, any word which would leave a residue of the input consisting only of consonants should not be accepted, because there would be no point looking such a residue up in the lexicon to see if it were a word—it never could be. Thus although the input *speech* might activate *each*, the latter would strand the residue *sp*, consisting only of consonants and hence not a possible word. Therefore, the candidacy of *each* should be discouraged. Norris et al. termed this constraint the Possible Word Constraint (PWC).

Norris et al. pointed out that the constraint clearly operates without reference to what actually *is* in the lexicon—*vuff* and *shub*, the residual contexts in “vuffapple” and “seashub” are no more English words than *f* and *sh* are.

Their only advantage is to be hypothetically possible words. Norris et al. also demonstrated that their results were not due to acoustic syllable boundary cues, because cross-spliced materials produced the same pattern of responding. Subsequent experiments established that it is not even important whether the residue be a possible stand-alone word in the vocabulary in question. Thus Cutler, Demuth, and McQueen (2002) showed that speakers of Sesotho spot words which leave a single syllable residue (which can *not* be a possible stand-alone word in Sesotho) just as easily as they spot words which leave a bisyllabic residue (a possible Sesotho word). And Norris, McQueen, Cutler, Butterfield, and Kearns (2001) showed that English listeners detect words which leave residual open syllables with short lax vowels (e.g., *væ* with the vowel of *vex*, not a possible English word) just as easily as words leaving a syllabic residue which could indeed be an English word (e.g., *voʊ*). These experiments together suggest that the constraint is universal in form and hence requires no acquaintance with a language-specific vocabulary to enable its operation.

Any constraint which does not depend on availability of a vocabulary could be, clearly, of great use to infants struggling to build an initial vocabulary from continuous-speech input. Cutler et al. (2002) indeed speculated that a universal PWC might be motivated by just such an initial use. If the form of the PWC is simply: reject parses containing any chunk without a vowel (which in turn is effectively a portion of speech with steady-state periodicity), then the constraint may be implemented in a relatively primitive manner. Interestingly, the PWC is similar to a proposal made by Brent and Cartwright (1996) to improve performance of their computational model of vocabulary acquisition with a corpus of child-directed speech. By rejecting all parses that posited words without vowels, vocabulary acquisition by the model was greatly improved.

In the present study, therefore, we explore whether speech segmentation by 12-month-old English-learners indeed gives evidence of being constrained by the PWC.

2. Experiment 1

If infants observe the PWC, they should have more difficulty spotting words in impossible contexts (contexts that leave a single consonant unaccounted for) than possible contexts (contexts that leave a syllable). To test this, we devised a procedure that was intended to parallel, as much as possible, the one used by Norris et al. with adults. Of course, we could not simply instruct infants to indicate whenever they detected a word inside a longer item. However, we could test the infants' abilities to detect familiar targets located inside longer items, using a version of the two-stage familiarization and test procedure used by Jusczyk and Aslin (1995) and in the subsequent

studies described above. We familiarized infants with two monosyllabic words such as *lop* and *rush*, and then tested them on four lists of words. Two word lists contained one of the familiarized words inside some longer items, plus a number of filler items. Similarly, the other two lists each had a novel word inside longer items and a number of filler items. Half of the infants were tested on lists in which segmenting the target from its context left a residue that could not stand as a possible word (*lop* from “plop”). The remaining infants were tested on lists in which segmenting the target from its context left residues that could stand as possible words (*lop* from “meplop”). If 12-month-olds’ speech segmentation is constrained by PWC, then they should listen significantly longer to test lists in which segmenting the familiarized target from its container left a possible, as opposed to impossible, word as a residue.

2.1. Method

2.1.1. Participants

Thirty-two infants (14 boys, 18 girls) were recruited from monolingual, English-speaking homes in the Baltimore, MD region. The infants were approximately 12 months old, with a mean age of 364 days (range = 350–387 days). The data from five additional infants were not included for the following reasons: crying (2), parental interference (2), or average looking times less than 3 s during the test phase (1). The parents of all participants gave informed consent.

2.1.2. Stimuli

The familiarization items consisted of four lists (mean length = 16.62 s; range = 16.03–16.94 s). Each list contained 15 repetitions of a single monosyllabic word: *lop*, *rush*, *rack*, or *win*. Each of the 8 test lists contained 15 items: 6 containers and 9 fillers (see Appendix A). Containers were items with target words within them (i.e., *rush* in “prush”). Containers never occurred in immediate succession in the lists. Half of the test lists included containers with familiarization words in “impossible” positions (i.e., *rush* in “prush”); the other half included containers with familiarization items in “possible” positions (i.e., *rush* in “nprush”). The latter were pronounced with trochaic stress. The mean length of the test lists containing items in “impossible” positions was 16.05 s (range = 16.0–16.11). The mean length of test lists containing items in “possible” positions was 21.35 s (range = 21.16–21.53). All stimuli were produced in an infant-directed manner by a female talker from New York City.

2.1.3. Design

Infants were randomly assigned to one of two familiarization conditions. Half of the infants were familiarized with *lop* and *rush* while the other half

were familiarized with *rack* and *win*. During the test phase, for each pair of familiarization words, half of the infants heard four test lists containing “impossible” targets and half heard four test lists containing “possible” targets. Each test list included containers with one of the target words—*lop*, *rush*, *rack*, and *win*—embedded within longer words.

2.1.4. Apparatus

The experiment was conducted in a three-sided test booth constructed out of white pegboard panels. The back of the center panel was lined with white cardboard to prevent infants from noticing the experimenter’s movements behind the booth. A small section of the center panel was left unlined to allow for the monitoring of infants’ headturns. The test booth had a red light and a loudspeaker mounted at eye level on each of the side panels, and a green light mounted on the center panel in front of the infant. Directly below the center light, a 5-cm hole accommodated the lens of a video camera used to record each test session. A computer terminal and a response box were located behind the center panel, out of view of the infant. The response box, connected to a Macintosh computer, was equipped with a series of buttons that started and stopped the flashing center and side lights, recorded the direction and duration of head turns, and terminated a trial when the infant looked away for more than 2 s. Data regarding the direction and duration of head turns for each trial were stored in a computer data file. Computer software was responsible for the selection and randomization of the stimuli and for the termination of the test trials. The computer calculated the average listening times for the test items at the completion of each testing session.

All stimuli were recorded in a sound attenuated room using Kay Elemetric’s Computerized Speech Lab (CSL). The recordings were made with a Shure microphone and sampled at a rate 20 kHz via a 16-bit analog-to-digital converter. Digitized versions of the readings were transferred to a Macintosh Quadra 650 for playback during the experiment. The computer controlled the presentation of the soundfiles during the experiment. The audio output was generated from the digitized waveforms of the samples. A 16-bit D/A converter was used to recreate the audio signal. The output was fed through anti-aliasing filters and a Kenwood audio amplifier (KA 57) to the two 7-in. Cambridge Soundworks loudspeakers mounted on the side walls of the test booth.

2.1.5. Procedure

Each infant was tested individually while seated in a caregiver’s lap. Both the caregiver and experimenter listened to masking music over Peltor Aviation headphones (model 7050) so that they could not hear the stimuli presented to the infant. The masker consisted of loud instrumental music, which had been recorded with few silent periods. At the beginning of the

familiarization phase, the green light on the front panel of the testing booth began to flash. Once the infant oriented toward the center light, the center light was extinguished and the red light above the loudspeaker on one of the side panels began to flash. When the infant turned at least 30° in the direction of the light, one of the two familiarization lists began to play and continued until either the entire 15-item list had been played or the infant failed to maintain the 30° headturn for 2 consecutive seconds. If the infant turned away from the target by 30° in any direction for less than 2 s and turned back again, the time spent looking away was not included in the orientation time. Thus, the maximum orientation time for a given trial was the duration of the entire speech sample. The red light flashed for the duration of the trial.

In the familiarization phase, two lists (*lop* and *rush* or *rack* and *win*) were presented in alternating order until the infants had accrued at least 30 s looking time towards each list. Immediately following familiarization, 12 test trials were presented (3 trials for each of the four test lists, blocked and presented in random order). Each test trial started with the green center light flashing, and the coordination between the lights and audio files was the same as during the familiarization period. Lists with containers including the familiar targets were presented during half of the trials, lists with containers including unfamiliar words were presented during the other half. The side from which the stimuli were played varied randomly from trial to trial, with the stipulation that no more than three trials in a row could be on the same side. The observer was not informed as to the familiarization condition to which the infant was assigned. Because both the observer and the infant's caregiver wore earplugs and listened to masking music over tight-fitting closed headphones, they were unaware of the ordering of the test samples.

2.2. Results and discussion

Mean orientation times to the two types of test lists, Familiar (with containers that included the targets from familiarization) and Unfamiliar (with containers that included novel words) were calculated for infants in both the Possible and Impossible conditions and are displayed in Fig. 1. In the Impossible condition, 8 of 16 participants had longer orientations towards the targets containing familiar items. In the Possible condition, 13 of 16 participants had longer orientations toward the targets containing familiar items. A mixed-design 2 (context: Possible versus Impossible) \times 2 (Test items: Familiar versus Unfamiliar) ANOVA was conducted, which revealed no significant main effect of Context Possibility, $F < 1$. However, there was a significant main effect of Familiarity, $F(1, 30) = 4.18, p < .05$. In addition, there was a significant interaction between Context Possibility and Familiarity, $F(1, 30) = 8.24, p < .01$. This interaction was further explored in t tests,




Fig. 1. Mean listening times of infants in Experiment 1 to the test lists with familiar and unfamiliar target words embedded at the end of longer nonsense words. The nonsense words were bisyllables for the Possible context and monosyllables for the Impossible context.

which revealed a significant Familiarity effect in the Possible condition ($t[15] = 3.11, p < .01$), but no significant difference in the Impossible condition ($t < 1$).

These findings suggest that 12-month-olds, like adults, find it easier to spot words within longer words if doing so does not strand a single consonant as residue. In short, 12-month-olds appear to exploit the PWC in speech segmentation.

It is interesting to compare infants' behavior here and in previous studies in which infants also had to match a familiarized target to a part of a larger word. Thus Jusczyk et al. (1999b) found that English-learning 7.5-month-olds who were familiarized with a strong syllable, such as *tar*, subsequently listened significantly longer to a passage with a weak/strong word containing the same syllable (i.e., *guitar*). However, infants at the same age did not show the same pattern of results when familiarized with a strong syllable, such as *king*, and tested on passages containing a strong/weak word containing the same strong syllable (i.e., *kingdom*). Jusczyk et al. (1999b) attributed this difference to the fact that in the strong/weak word, but not in the weak/strong word, the strong syllable always co-occurs with the same following syllable.² In fact, they demonstrated that when a weak/strong word was always followed by the same weak syllable (e.g., *guitar is*), infants no longer matched the familiarized strong syllable to the word in the passage. Jusczyk et al. (1999a) likewise found that 9-month-olds familiarized with *night* did

² The fact that in the weak/strong word, the strong syllable always co-occurs with the same preceding weak syllable is countered by infants' tendency to segment at the onsets of strong syllables at this age.

not match it to a passage in which *night* was always followed by *rates*, but did match it to a passage in which *night* was followed by six different words (e.g., *game*, *school*, *gown*, etc.). The fact that infants in Experiment 1 did find a match to the targets in the longer container words is similar, because the containers differed from each other in every instance. The only portion that remained the same was the section that matched the target.

However, an asymmetry between target words embedded finally versus initially in their containers cannot be ruled out in the present study. Although Norris et al. (1997) found that adult segmentation was constrained by the PWC regardless of whether target words occurred at the end of the longer nonsense item (*apple* in “fapple” or “vuffapple”) or at the beginning (*sea* in “seash” or “seashub”), the response time advantage of possible over impossible contexts which they observed was largest for words in final position, with the context preceding the word. Studies of adult listening with other techniques have also shown greater inhibition of activation for final-embedded words preceded by impossible contexts than for initial-embedded words followed by impossible contexts (Vroomen & de Gelder, 1997). The computational simulations with the Shortlist model (Norris, 1994) presented by Norris et al. supported the empirical results, in that they too showed a stronger effect for words in final position. This is because initially embedded words have an opportunity to reach high levels of activation before the PWC is triggered by the following context. In the case of finally embedded words, in contrast, the PWC comes into play at the very beginning of the word, thus preventing the embedded word from ever becoming highly activated.

Activation of finally embedded words is also subject to competition from rival words which begin with the same onset sequence as the container (e.g., the initial portion of “fapple” would activate words beginning *fa-* - *fat*, *fashion*, etc., which would compete with the embedded word *apple*). Competition will, therefore, also act to prevent finally embedded words from reaching a high level of activation. Words at the beginning of the container, in contrast, have time to reach high levels of activation before any inhibition takes effect as a result of competition from candidates supported by the later-occurring context. Thus both the timing of the PWC’s operation, and effects of competition, conspire to make initially embedded words easier to spot than finally embedded words.

Inter-word competition effects in infant listening will of course be trivial in comparison to those occurring in listening by adults with a fully stocked vocabulary. However, word form activation itself is undeniably affected by the simple fact that an initially embedded word is presented to the listener in its entirety before potentially inhibiting context arrives. Thus it is not necessarily the case that the pattern of results for initially embedded words can be predicted from that for finally embedded words. In our next experiment, we assess whether the PWC effect observed in Experiment 1 can also be ob-

served with words occurring at the beginning of their nonsense item containers, i.e., where the word itself precedes rather than follows the residual context.

3. Experiment 2

This experiment primarily differs from Experiment 1 in the location of the targets relative to the container items. The familiarized targets were here located at the beginning of the container (*dull* in “dulp” or “dullpick”) rather than at the end (*lop* in “plop” or “meplop” in Experiment 1). If 12-month-olds observe the PWC regardless of target location, then as in Experiment 1 they should listen significantly longer to test lists containing the familiarized targets in possible as opposed to impossible contexts.

3.1. Method

3.1.1. Participants

Thirty-two infants (12 boys, 20 girls) were recruited from monolingual, English-speaking homes in the Baltimore, MD region. The infants were approximately 12 months old, with a mean age of 365 days (range = 351–380 days). The data from nine additional infants were not included for the following reasons: crying (5), parental interference (2), or average looking times less than 3 s during the test phase (2). The parents of all participants gave informed consent.

3.1.2. Stimuli

The familiarization items consisted of four lists (mean length = 16.46 s; range = 16.03–16.92 s). Each list contained 15 repetitions of a single monosyllabic word: *dull*, *fill*, *moss*, or *tan*. It was necessary to change the target words from those in the previous experiment to create appropriate containers with phonetic sequences that are phonotactically permissible in English. As in Experiment 1, each of the 8 test lists contained 15 items: 6 containers and 9 fillers (see Appendix A). Containers were items with familiar target words within them (i.e., *dull* in “dulp”). Containers never occurred in immediate succession on the test lists. Once again, half of the test lists included containers with targets in “impossible” positions (i.e., *dull* in “dulp”), the other half had containers with targets in “possible” positions (i.e., *dull* in “dullpick”). Again, the bisyllabic items were pronounced with trochaic stress. The mean length of the test lists with targets in “impossible” positions was 15.92 s (range = 15.68–16.26). The mean length of test lists with targets in “possible” positions was 21.33 s (range = 21.17–21.44). All stimuli were produced in an infant-directed manner by the same female from New York City who had recorded the stimuli for Experiment 1.

3.1.3. Design

Infants were randomly assigned to one of two familiarization conditions. Half of the infants were familiarized with *dull* and *fill*, whereas the remaining infants were familiarized with *tan* and *moss*. During the test phase, for each pair of familiarization words, half of the infants heard four test lists with “impossible” targets, whereas the other half heard four test lists with “possible” targets. Each test list included one of the target words—*dull*, *fill*, *moss*, and *tan*—embedded within longer words.

3.1.4. Procedure, apparatus

The procedure and apparatus were identical to those in Experiment 1.

3.2. Results and discussion

Mean orientation times to the two types of test lists (familiar and unfamiliar) in the Possible and Impossible conditions are shown in Fig. 2. In the Impossible condition, 12 out of 16 participants had longer orientations towards the targets containing familiar items. In the Possible condition, 8 out of 16 participants had longer orientations towards the targets containing familiar items. An ANOVA was conducted as in Experiment 1; it revealed no significant main effect of Context Possibility, $F < 1$. As in Experiment 1, there was a significant main effect of Familiarity, $F(1, 30) = 4.3$, $p < .05$. However, there was this time no significant interaction between Context Possibility and Familiarity, $F(1, 30) = 3.0$, $p > .05$. Despite the insignificance of this interaction, we conducted t tests separately for the two context conditions. There was no effect of Familiarity for the Possible items

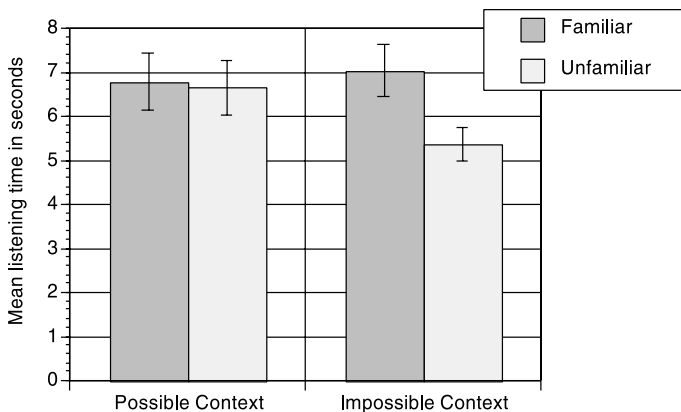


Fig. 2. Mean listening times of infants in Experiment 2 to the test lists with familiar and unfamiliar target words embedded at the beginning of longer nonsense words. The nonsense words were bisyllables for the Possible context and monosyllables for the Impossible context.

($t < 1$), but there was, directly counter to our predictions, a significant effect for the Impossible condition ($t[15] = 2.36, p < .05$).

These results provide no evidence that infants, like adults, observe the PWC regardless of where the impossible residue is located. In fact, infants here seemed to find the targets more easily when they were located within the impossible than the possible containers. This is exactly the reverse of what would be expected if infants were making use of the PWC.

As we pointed out above, words preceded by a silent pause may be highly activated before subsequent context arrives. The following context might then no longer exercise any effect on recognition. This does not itself explain why the predicted effect actually reversed in this experiment; but note that other research has suggested that infants are particularly sensitive to sound similarities at the beginnings of words (Goodman & Jusczyk, 2000; Jusczyk, Goodman, & Bauman, 1999). If infants' attention is focused on the beginnings of words, then they may simply ignore the final "p" in an item such as "dullp" (whereas a following syllable is harder to ignore, making the Possible context actually a more challenging container). We will return to this issue in Section 7. In Experiments 3 and 4, we explored the two types of embedding (final, initial) further, but rather than testing infants on word lists containing nonsense items, we used continuous passages containing real words. These experiments thus offer a closer approximation to a natural listening situation.

4. Experiment 3

This experiment is identical to Experiment 1 in all but two respects. First, the targets are real words buried in contexts that also form real words (*rest* in "crest" or "unrest"). Second, the test items are embedded in passages rather than word lists. If 12-month-olds can use the PWC for online segmentation of fluent speech, then as in Experiment 1, infants should listen significantly longer to test passages containing the targets when their segmentation from the containers leaves residues that are possible, as opposed to impossible, words.

4.1. Method

4.1.1. Participants

Forty infants (20 males, 20 females) were recruited from monolingual, English-speaking homes in the Baltimore, MD region. The infants were approximately 12 months old, with a mean age of 366 days (range = 353–383 days). The data from three additional infants were not included for the following reasons: crying (2) or average looking times less than 3 s during the test phase (1). The parents of all participants gave informed consent.

4.1.2. Stimuli

The familiarization items consisted of four lists (mean length = 16.22 s; range = 16.05–16.48 s). Each list contained 15 repetitions of a single monosyllabic word: *lay*, *rise*, *low*, or *rest*. In contrast to Experiments 1 and 2, the targets were embedded in passages during the test phase (see Appendix A). For four test passages, segmenting the target from the container would leave a possible word as a residue (i.e., *rest* in “unrest”), whereas for the other four test passages, the residue would constitute an impossible word (i.e., *rest* in “crest”). For each of the context possibility conditions, each of the four familiarization words occurred in one passage. Thus for any infant, the two target words heard during familiarization were each present inside longer words in one of the test passages, while the other two passages contained novel words embedded inside longer words. It is worth noting that the word containing the target was never repeated; rather, it varied from sentence to sentence in the passages. The mean length of the test passages that included targets and impossible words as residues was 22.45 s (range = 22.23–22.67). The mean length of test passages that included targets and possible words as residues was 22.33 s (range = 22.00–22.53). All stimuli were produced in an infant-directed manner by the same female from New York City who had recorded the stimuli in the previous experiments.

4.1.3. Design

Infants were randomly assigned to one of two familiarization conditions. Half of the infants were familiarized with *rise* and *lay* and half were familiarized with *rest* and *low*. Half of the infants were assigned to the “impossible” condition, where segmenting the target from the container left an impossible word as residue. The remaining infants were assigned to the “possible” condition where segmenting the target left a possible word as residue. In each condition, the four different passages included one of the four target words—*rise*, *lay*, *rest*, and *low*—embedded within longer words.

4.1.4. Procedure, apparatus

The procedure and apparatus were identical to those in Experiment 1. The only difference was the use of passages rather than word lists during the test phase.

4.2. Results and discussion

Mean orientation times to the two types of test passages (familiar and unfamiliar) in the Possible and Impossible conditions are shown in Fig. 3. In the Impossible condition, six of 20 participants had longer orientations towards the targets containing familiar items. In the Possible condition, 18 out of 20 participants had longer orientations towards the targets containing familiar items. The ANOVA again revealed no significant main effect of




Fig. 3. Mean listening times of infants in Experiment 3 to the passages with familiar and unfamiliar target words embedded at the end of longer words. The longer words were bisyllables for the Possible context and monosyllables for the Impossible context.

Context Possibility, $F < 1$. As in Experiment 1, there was a significant main effect of Familiarity, $F(1, 38) = 4.69$, $p < .05$ and a significant interaction between Context Possibility and Familiarity, $F(1, 38) = 6.97$, $p < .05$. Post hoc t tests revealed a significant effect of Familiarity for Possible contexts ($t[19] = 3.12$, $p < .01$), but not for Impossible contexts ($t < 1$).

These results suggest that even when infants are tested on fluent speech passages, they segment targets from their containers in accordance with what is predicted if they follow the PWC. Hence, like the findings of Experiment 1, the present findings lend support to the view that 12-month-olds observe the PWC in segmenting words. However, as with the comparison of Experiment 1 with Experiment 2, it is not necessarily the case that infants will show the same response pattern when targets are embedded at the beginning of longer words as when targets are embedded at the end of longer words. This issue is tested in the following experiment.

5. Experiment 4

The following experiment is identical to Experiment 3 in all but one respect, namely, the locus of the targets. As in Experiment 2, the familiarized targets were located at the beginning of the longer words (*win* in “wind” or “window”). If infants fail to observe the PWC, then it would appear that the results of Experiment 2 were not simply due to the format in which the stimuli were presented, i.e., as a list or in a passage. However, if 12-month-olds truly observe the PWC regardless of target location, then they should listen significantly longer to test passages for which the segmentation of targets

from their containers leaves residues that are possible, as opposed to impossible, words.

5.1. Method

5.1.1. Participants

Forty infants (18 males, 22 females) were recruited from monolingual, English-speaking homes in the Baltimore, MD region. The infants were approximately 12 months old, with a mean age of 369 days (range = 355–381 days). The data from five additional infants were not included for the following reasons: crying (2), parental interference (2), or average looking times less than 3 s during the test phase (1). The parents of all participants gave informed consent.

5.1.2. Stimuli

The familiarization items consisted of four lists (mean length = 15.94 s; range = 15.78–16.07 s). Each list contained 15 repetitions of a single monosyllabic word: *well*, *may*, *win*, or *rye*. As in Experiment 3, the targets were embedded in passages rather than lists (see Appendix A). The containers were words with the familiarized targets embedded within them (i.e., *win* in “wind”). For four of the test passages, segmenting the target from the containers would leave a possible word as a residue (i.e., *win* in “window”). For the other four passages, segmenting the target from the container left an impossible word as a residue (i.e., *win* in “wind”). Once again, the containers for the target varied from sentence to sentence in the test passages. The mean length of the test passages with targets and impossible words as residues was 20.03 s (range = 19.56–20.47). The mean length of passages with targets and possible words as residues was 19.8 s (range = 19.41–20.03). All stimuli were produced in an infant-directed manner by the same female from New York who had recorded the stimuli from the previous experiments.

5.1.3. Design

Infants were randomly assigned to one of two familiarization conditions. Half of the infants were familiarized with *well* and *may* and the other half with *win* and *rye*. Half of the infants were assigned to the “impossible” condition, where segmenting the target from the container left an impossible word as a residue. The other infants were assigned to the “possible” condition, where segmenting the target left a possible word as a residue. In each condition, the four different passages included one of the four target words—*well*, *may*, *win*, and *rye*—embedded within longer words.

5.1.4. Procedure, apparatus

The procedure and apparatus were identical to those in Experiment 3.

5.2. Results and discussion

Mean orientation times to the two types of test lists (familiar and unfamiliar) in the Possible and Impossible conditions are shown in Fig. 4. In the Impossible condition, nine of 20 participants had longer orientations towards the targets containing familiar items. In the Possible condition, 15 of 20 participants had longer orientations towards the targets containing familiar items. The ANOVA revealed once again no significant main effect of Context Possibility, $F < 1$. As in Experiments 1 and 3, there was a significant main effect of Familiarity, $F(1, 38) = 4.06, p < .05$ and a significant interaction between Context Possibility and Familiarity, $F(1, 38) = 7.39, p < .01$. Again the post hoc t tests revealed a significant effect of Familiarity for Possible contexts ($t[19] = 4.48, p < .001$), but no such effect for Impossible contexts ($t < 1$).

In contrast to the results of Experiment 2, these results indicate that 12-month-olds do seem to observe the PWC when targets are embedded at the beginning of a longer word. The major difference between these two experiments concerned the format in which the information was presented (i.e., lists in Experiment 2 and fluent speech passages in Experiment 4). One might, therefore, argue that infants engage the PWC when detecting targets in fluent speech but do not engage segmentation processes to the same extent in lists. However, this does not explain why infants observed the PWC in Experiment 1 but not in Experiment 2 given that the test stimuli were presented in lists in both experiments. The answer must thus lie also in the position of the target relative to its container. Only when test items are presented in lists, *and* with initially embedded targets, do we fail to observe a significant effect of the PWC. However, we decided to further test the robustness of the Experiment 2 result, in case it should prove to be due to

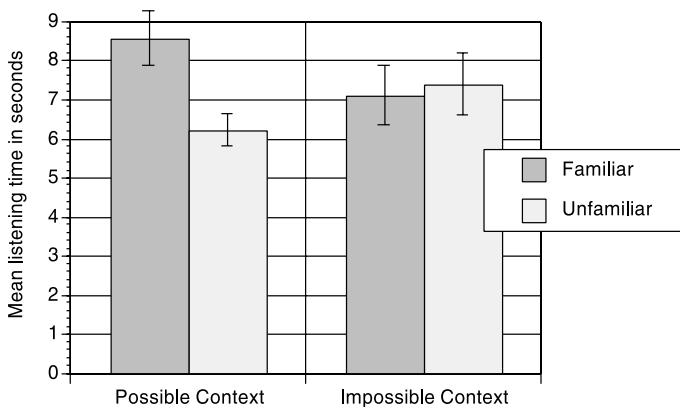


Fig. 4. Mean listening times of infants in Experiment 4 to the passages with familiar and unfamiliar target words embedded at the beginning of longer words. The longer words were bisyllables for the Possible context and monosyllables for the Impossible context.

some experiment-specific artifact. For instance, the talker might accidentally have produced the nonsense containers of Experiment 2 less fluently than the real word containers of Experiment 4, thus inadvertently providing juncture information at the offset of the target which would obviate the need for segmentation. In Experiment 5, then, we replicated Experiment 2 with the real-word materials of Experiment 4.

6. Experiment 5

The present experiment used the same real-word containers as in Experiment 4; but instead of appearing in fluent speech passages, the containers were presented in test lists, as in Experiment 2.

6.1. Method

6.1.1. Participants

Thirty-two infants (19 males, 13 females) were recruited from monolingual, English-speaking homes in the Baltimore, MD region. The infants were approximately 12 months old, with a mean age of 367 days (range = 352–389 days). The data from six additional infants were not included for the following reasons: crying (3), parental interference (1), or average looking times less than 3 s during the test phase (2). The parents of all participants gave informed consent.

6.1.2. Stimuli

The familiarization items consisted of four lists (mean length = 14.73 s; range = 14.58–14.84 s). Each list contained 15 repetitions of a single monosyllabic word: *well*, *may*, *win*, or *rye*. As in Experiments 1 and 2, the targets were embedded in test lists (see Appendix A). Containers were items with targets embedded within them (i.e., *well* in “welt”). Containers never occurred in immediate succession on the test lists. Once again, half of the test lists included containers with targets in “impossible” positions (i.e., *may* in “mail”), the other half had containers with targets in “possible” positions (i.e., *well* in “welter”). The mean length of the test lists with the targets in “impossible” positions was 14.0 s (range = 13.94–14.03). The mean length of test lists with the targets in “possible” positions was 15.09 s (range = 15.03–15.15). Because the talker who had recorded the stimuli for the previous experiments was no longer available, the new stimuli were produced in an infant-directed manner by a female from Buffalo, New York.

6.1.3. Design

Infants were randomly assigned to one of two familiarization conditions. Half of the infants were familiarized with *well* and *may* and the other half

were familiarized with *win* and *rye*. During the test phase, for each pair of familiarization words, half of the infants heard four test lists with “impossible” targets, whereas the other half heard four test lists with “possible” targets. Each test list included one of the target words—*well*, *may*, *win*, and *rye*—embedded within longer words.

6.1.4. Procedure, apparatus

The procedure and apparatus were identical to those in Experiment 1.

6.2. Results and discussion

Mean orientation times to the two types of test lists (familiar and unfamiliar) in the Possible and Impossible conditions are shown in Fig. 5. In the Impossible condition, seven of 16 participants had longer orientations towards the targets containing familiar items. In the Possible condition, six of 16 participants had longer orientations towards the targets containing familiar items. The ANOVA revealed no significant main effect of Context Possibility, $F < 1$. In contrast to Experiment 2, there was also no significant main effect of Familiarity, $F < 1$. Most importantly, the interaction between Context Possibility and Familiarity did not approach statistical significance, $F < 1$, and separate t tests for each context condition were likewise insignificant (both $t < 1$).

The lack of a significant Familiarity effect in the present experiment was surprising, since this effect had appeared in all previous experiments, including Experiment 2. To check for possible differences in how infants had responded to the targets in that experiment and the present one, we




Fig. 5. Mean listening times of infants in Experiment 5 to the test lists with familiar and unfamiliar target words embedded at the beginning of longer words. The longer words were bisyllables for the Possible context and monosyllables for the Impossible context.

conducted an omnibus ANOVA combining the data from both experiments. None of the main effects of Experiment ($F < 1$), Context Possibility ($F < 1$), or Familiarity ($F[1, 61] = 2.17, p > .1$) was significant, and nor were the interactions of Familiarity with Experiment ($F[1, 61] = 3.32, p > .05$) and with Context Possibility ($F[1, 61] = 2.98, p > .05$). Hence, there is no evidence that infants in Experiment 5 behaved differently than those in Experiment 2.

In two experiments, then, infants failed to observe the PWC when familiarized items were located at the beginning of longer words presented in lists. Of course, the experiments differ in that the statistical reversal of the PWC in Experiment 2 was not replicated in Experiment 5. We have no definitive explanation for why the effect reversed in Experiment 2, although we cannot rule out some as yet unknown aspect of our materials selection. However, we doubt that the impossible contexts in Experiment 2 were simply less perceptible than those in Experiment 5, given that the list duration was if anything rather longer in Experiment 2. We noted that only Experiment 5 contained target words with CV structure, which might in principle be harder to extract from a following context than CVC words (which was the structure of all four words in Experiment 2). In this case Experiment 5 might show different results for the target words *rye* and *may* (with CV structure) versus for *well* and *win* (with CVC structure). Therefore, we conducted a post hoc analysis comparing the results for the CV versus the CVC subsets; there was no significant difference in looking time between the two subsets and no interaction of target word structure with Familiarity.

However, the result of Experiment 5 definitely suggests that whether or not infants observe the PWC with initially embedded targets is a function of the kinds of materials used (i.e., lists versus passages). Experiment 2 was not anomalous and its results were not purely due to some artifact of the materials. The failure of PWC effects to show up in experiments using list presentation with initially embedded targets is replicable; infants are more likely to show effects of the PWC with such items when they are presented in fluent speech passages.

7. General discussion

Lexical viability constrains the segmentation of speech input even for listeners who are not yet in a position to make full use of a lexicon. Our study has clearly shown that the Possible Word Constraint plays a role in segmentation of speech by infant listeners, just as it constrains speech segmentation by adult listeners. At issue in the present five experiments were the abilities of 12-month-olds to detect the occurrence of familiarized target words embedded within longer words. The critical variables that were manipulated throughout the experiments concerned the types of contexts in which targets were embedded and the positioning of the targets within these contexts. Spe-

cifically, we examined whether targets would be easier to find when the residual context constituted lexically viable chunks (i.e., possible words). Our results confirmed that this was the case, at least in those contexts in which segmentation abilities are most called upon.

The test materials for Experiments 1 and 2 consisted of lists of either monosyllabic or bisyllabic words. In Experiment 1, the target words occurred at the end of the longer words (e.g., *rush* in “prush” or “niprush”). In line with the PWC, infants were better able to segment the targets from the latter context than the former. However, in Experiment 2, when the targets occurred at the beginning of the longer words, segmentation was no better for the targets inside bisyllables than for targets inside monosyllables. We argue that the format of the materials, as lists of words, affected the extent to which the PWC had scope for application in this experiment. In Experiments 3 and 4, we explored infants’ ability to segment the targets from monosyllabic and bisyllabic longer words in fluent speech passages. The results of these experiments indicated that, regardless of whether the targets were embedded at the ends (Experiment 3) or beginnings of words (Experiment 4), infants were better able to detect the targets in the bisyllabic words. Finally, to determine the robustness of the different result in Experiment 2, Experiment 5 replicated that design, but with the targets inside real words. We obtained the same pattern of results as in Experiment 2, suggesting that the PWC is indeed less helpful with following contexts in word lists. Instead, it is most useful in fluent speech contexts and with preceding contexts—situations which may be said to make heavier demands on segmentation abilities.

Word recognition and word segmentation are intimately related. Fluent speech is continuous (i.e., without clear pauses between words). Yet understanding requires that speech be mapped to lexical entries. Adult listeners can rely to a considerable extent on processes of automatic lexical activation and inter-word competition. But there are also other sources of information which point to locations in the signal that could correspond to the onsets or offsets of distinct words and thereby contribute to the activation of lexical candidates. These types of information (prosodic stress, phonotactics, context-sensitive allophones, and so on) are all exploited efficiently by adult listeners. And infants within the first year of life likewise exploit prosodic information (Jusczyk et al., 1999b), phonotactic sequence constraints (Mattys et al., 1999; Mattys & Jusczyk, 2001), context-sensitive allophones (Jusczyk et al., 1999a), and even statistical cues to word boundaries (Saffran, Aslin, & Newport, 1996) and talker-specific idiosyncasies (Houston, 1999; Nazzi, Jusczyk, & Johnson, 2000). All of these types of information can be learned from experience with language input and all can help with the problem of segmenting words from context.

The PWC is a further constraint that can help with this problem. The PWC reduces the number of lexical candidates to be considered by

penalizing any parse (and the lexical candidates associated with it) which leaves a residue that cannot stand as a possible word. The results of Cutler et al. (2002) and Norris et al. (2001), described in Section 1, suggest that the PWC is independent of language-specific vocabularies, and hence putatively universal across languages. A procedure which requires no prior knowledge of a particular language is arguably even more useful for segmenting new words from fluent speech contexts than the vocabulary-dependent procedures (such as use of phonotactic sequence constraints, or of the prosodic characteristics of words) which we know can be acquired in infancy. A universal constraint can help acquire whatever language might be presented in the environment, from the earliest exposure onwards. It can, for instance, minimize unhelpful false-alarm candidacy of accidentally recurring strings in the input (such as *bl-* in “black and blue”).

The five experiments presented in this paper provide support for the proposal that English-learning 12-month-olds indeed rely on the PWC in parsing speech. Experiments 1 and 3 indicated that when target words are embedded at the end of longer words, infants were better able to recognize the targets when segmenting them left residues that are possible (e.g., *rest* in “unrest”) as opposed to impossible words (e.g., *rest* in “crest”). Similarly, in Experiment 4, which used fluent speech passages, infants were better able to detect target words embedded at the beginning of longer words when segmenting them left residues that are possible as opposed to impossible words. The cases for which there was no indication that infants observed the PWC (Experiments 2 and 5) both involved the segmentation of target words embedded in longer words, when the test items were presented in the form of lists. Contrary to our original predictions, infants did not use the PWC to recognize words embedded at the beginning of longer words (e.g., *win* in “wind” or “window”). However, the list materials of Experiments 2 and 5 may not have provided adequate scope for infants to exploit the operation of the PWC.

As we pointed out above, adult speech perception also shows weaker PWC effects with initially embedded than with finally embedded words when the containers are presented in isolation. In Norris et al.’s (1997) original study, responses to bisyllabic words (e.g., *sugar*) were in fact in one of their experiments fractionally slower in possible following contexts (“sugarthig”) than in impossible following contexts (“sugarth”). Vroomen and de Gelder (1997) likewise reported cross-modal priming effects from embedded words with following impossible contexts (such as *bell* in “belk”)—although note that their experiment was actually carried out in Dutch), but no such effects—i.e., no effective activation—of embedded words with preceding impossible contexts (such as *wine* in “dwine”). In Norris et al.’s computational simulations, in which activation of word candidates in the Shortlist model was halved if an impossible residue would be stranded, the level of activation of finally embedded words was lowered by

competition from words activated by the preceding context (*fat*, etc. in the case of “fapple”). Once halved, this activation thus became very low indeed. For initially embedded words, on the other hand, the activation reached a high level before arrival of the impossible context; halving still left it higher than activation of the finally embedded words, and thus Norris et al.’s simulations correctly captured the response pattern from the experiments.

Infant speech perception is less troubled by competition. But as our study has shown, infant speech perception exhibits the same asymmetry of initially versus finally embedded words when the containers are presented in isolation. We suggest that here too it is the degree of activation of the word form which is crucial—preceded by a pause, a word can be clearly perceived and its form activated in memory to a level which is reasonably resistant to subsequent inhibition by the PWC when an impossible context follows. When an impossible context precedes the word, however, the activation of its form can be immediately inhibited by the PWC because it will be obvious from the outset that the resulting residue would be unacceptable.

In continuous speech, our infant subjects found both initially and finally embedded words significantly harder to recognize in impossible contexts. This suggests that the processing of the context as a whole delayed activation of the initially embedded word compared with the case where the word was preceded by a pause; this delay allowed the PWC to come into operation as effectively as with finally embedded words. Note that no adult study has so far investigated the PWC’s operation in continuous speech; it would be interesting to observe whether in adults, too, the asymmetry between initially versus finally embedded words would disappear under such conditions.

The form of the PWC is simple: parse speech into chunks such that each chunk contains at least one vowel. Thus, any word candidate should be rejected if accepting it would leave a vowel-less residue. Since vowels can potentially be identified by a quite primitive mechanism, the PWC in this form could be acquired on the basis of the very earliest exposure to speech. The effect of a vowel-based constraint is to allow only chunks which are minimally a syllable in length, since the defining characteristic of a syllable is that it must contain a vowel (though it may optionally also contain consonants around this obligatory vowel nucleus). Across languages, the smallest acceptable stand-alone word is a syllable. As we pointed out above, this does not mean that any syllable can be a word in any language. Many languages have severe restrictions on what may be a stand-alone word—thus English accepts no monosyllables ending with short full vowels as stand-alone words, and Sesotho allows no stand-alone monosyllables at all. But every syllable can be a word in some language, somewhere. Thus the PWC allows only chunks which are, universally,

possible words. In this respect the PWC could be particularly useful in early vocabulary acquisition.

Note that there is other evidence that infants are sensitive to syllables but not to vowel-less strings, and even that such sensitivity is in place at birth. Thus neonates will suck more strongly given a signal which activates a recording of their mother's voice than given a signal which precedes silence, and the three-phoneme syllables *pat*, *tap* work as discriminably separate signals in this paradigm, whereas the three-phoneme nonsyllables *pst*, *tsp* do not (Moon, Bever, & Fifer, 1992). Mehler, Dupoux, and Segui (1990) indeed proposed a model of infant speech perception in which syllables are the basic building blocks from which a lexicon is constructed. However, it is hard to justify such a model given that it would presumably apply only to infant vocabulary acquisition. Because syllabic structure varies widely across languages and many languages do not have easily perceptible syllable boundaries, syllabic segmentation of speech would be impracticable universally; apart from this, there is now abundant evidence that speech processing in adulthood is truly continuous, without intermediate recoding into units such as syllables (Cutler & Otake, 2002; Marslen-Wilson & Warren, 1994; McQueen, Norris, & Cutler, 1999). Postulating a level of representation to obtain in infancy only is not parsimonious.

As Cutler, McQueen, Norris, and Somejuan (2001) observed, the PWC goes a step further than any syllabically based hypothesis in accounting for the special status of the syllable in perception. It offers a rationale for this status. Syllables are acceptable chunks into which speech can be parsed, precisely because they are lexically viable in a way that vowel-less chunks are not. In this way both infant and adult speech perception can be seen to be driven primarily by the goal of deriving lexical meaning from speech.

Like adults, 12-month-old infants disfavor parses that posit single consonants as candidate words. Thus, as our findings have indicated, they can more readily detect a monosyllabic target embedded within a bisyllabic word than the same target within a longer monosyllabic word. In other words, infants observe the PWC when parsing words from fluent speech, just as adult listeners, in experiments with many languages, have been shown to do. The fact that infants without a well-stocked lexicon can make use of the PWC suggests that this constraint, which appears to have a universally constant implementation, may have its roots in the initial process of acquiring a vocabulary.

Acknowledgments

This investigation was supported by a Research Grant from NICHD (#15795) and a Senior Scientist Award from NIMH (#01490) to PWJ. The authors wish to thank Ann Marie Jusczyk and two anonymous review-

ers for helpful comments on previous versions of the manuscript, and Mat-tijn Morren for help in preparing stimulus materials. Previous reports of part of this work were presented to the 22nd Annual Conference of the Cognitive Science Society, Philadelphia, in August 2000 and to the 140th Meeting of the Acoustical Society of America, Newport Beach, CA, in November 2000.

Appendix A. Stimulus materials

A.1. Experiment 1

rush in impossible context: θæp, θʌʃ, kʌæv, bɪf, pʌʃ, mæf, plɛv, ɡʌʃ, jɛθ, tʌʃ, matʃ, klɪv, dʌʃ, ɡlɛf, kʌʃ

win in impossible context: mʌθ, twɪn, tʃæɡ, θɪv, klʌd, dwɪn, fɪp, θwɪn, tɪv, kwɪn, mʌv, ɡwɪn, ʃɛk, ɡlʌm, swɪn

lop in impossible context: plʌp, θɪɡ, flʌp, ʃɪb, dʌæv, tʃʌn, ɡlʌp, wæf, dʒʌz, blʌp, ɡlɛv, kʌz, klʌp, lʌtʃ, slʌp

rack in impossible context: ʃjæk, stæv, nɪɡ, bjæk, plʌf, flʌtʃ, fjæk, fɛp, pɪɛm, tjæk, tʃʌv, ɡjæk, ʃɪɡ, tʃɪm, pjæk

rush in possible context: θæplʌt, nɪpʌʃ, kʌævbaɪp, bɪfmat, sæɡʌʃ, mæfɪt, plɛvntɛp, bɛtʌʃ, jɛθnʌʃ, hɛθʌʃ, matʃbas, klɪvpaɪt, ʃadʌʃ, ɡlɛfzɛm, pʌkʌʃ

win in possible context: mʌθɡʌt, mʌtwɪn, tʃæɡpʊt, θɪvtɪn, klʌdmɪk, hʊdwɪn, fɪpɪk, nɛθwɪn, tɪvnaɪt, sɛkwɪn, mʌvbaɪt, zʌɡwɪn, ʃɛksɪt, ɡlʌmtɪv, pæsɪvɪn

lop in possible context: mɛplʌt, θɪɡʌv, hæflʌp, ʃɪblʌt, dʌævpaɪn, tʃʌntɪk, dʒɪɡlʌp, wæfnʌɡ, dʒʌzɪt, næblʌp, ɡlɛvɪk, kʌzɪn, zʌklʌp, lʌtʃmʌs, flʌslʌp

rack in possible context: ɡapjæk, stævɪk, nɪɡɪt, mɪfjæk, plʌftɪp, flʌtʃkʌt, zɛɡjæk, fɛptɪs, pɪɛmtɪk, nʌtjæk, tʃʌvtɪn, tʃɪbjæk, ʃɪɡbaɪp, tʃɪmtɪl, sæfjæk

A.2. Experiment 2

dull in impossible context: θæp, dʊlf, kʌæv, bɪθ, dʊlp, mæf, plɛv, dʊls, jɛθ, dʊlt, matʃ, klɪv, dʊʃ, ɡlɛf, dʊlk

fill in impossible context: fɪlθ, θɪɡ, fɪlp, ʃɪb, dʌæv, tʃʊn, fɪlk, waf, dʒʌz, fɪʃ, ɡlɛv, kʌz, fɪlt, lʌtʃ, fɪls

moss in impossible context: mʌsk, stæv, nɪɡ, mʌsp, plʌf, vʌtʃ, mʌst, fɛp, prɛm, mʌsk, tʃʊv, mʌsp, ʃɪɡ, tʃɪm, mʌst

tan in impossible context: nʌθ, tæns, tʃæɡ, θɪv, klʌd, tænt, fɪp, tæntʃ, tɪv, tæns, mʌv, tænt, ʃɛk, ɡlʌm, tæntʃ

dull in possible context: dʊlpɪk, θɪɡjʌv, dʊlfɛn, fɪblat, dʒævpʌn, tʃʌntɪk, dʊltʌp, wæfnʌɡ, dʒʌzɪt, dʊlfæt, glɛvɪk, kʌztɪn, dʊlsɪm, lʌtʃmʌs, dʊlkɛf

fill in possible context: θæplʌt, fɪlθɛk, kʒævbʌp, bɪθmat, fɪlsʌp, mæfɪt, plɛvtep, fɪljæm, jɛθnʌʃ, fɪlɪt, matʃbas, klɪvpat, fɪlkɛs, glɛfzɛm, fɪltʌf

moss in possible context: mʌskɪv, stævɪk, nɪɡɪt, mʌstɛb, plʌftɪp, vʌtʃkʌt, mʌspæn, fɛptɪs, pɛmtɪk, mʌskɪv, tʃʊvɪn, mʌstɛb, fɪɡbʌp, tʃɪmtɪl, mʌspæn

tan in possible context: nʌθɡʌt, tænfɛl, tʃæɡpʊt, θɪvtɪn, klʌdmɪk, tæntʊp, fɪpɪk, tænsʌm, tɪvɪnʌt, tænfɛl, mʌvbʌt, tæntʊp, fɛksɪt, glʌmtɪv, tænsʌm

A.3. Experiment 3

low in impossible context: That slow turtle claims to have beaten the speedy hare in a race. Mr. Matt Plo did not take the news of his defeat very well. The unexpected blow to his ego was a real shock to him. I put the star shaped sticker on the ceiling and watched it glow. The shy winner quailed at the huge flow of racing enthusiasts.

rest in impossible context: The proud young robin puffed out its vivid scarlet breast and sang. Our uniforms were taken to the dry cleaner's to be cleaned and pressed. The little bird with the odd crest really drew our attention. We had dressed warmly for our weekend hiking trip in the mountains. Our guide had stressed the importance of bringing plenty of water.

lay in impossible context: At harvest time we worked alongside the farmers that flay the wheat. I think we all felt that the first play we saw was simply riveting. Afterwards we made all sorts of silly creatures with the bright clay. Her hair would splay across her shoulders whenever she brushed it. A white sleigh carried the exuberant actors off of the stage.

rise in impossible context: Some fresh hot fries would taste wonderful with a chocolate milkshake. I found a wonderful prize in my box of caramel popcorn. Sun dries our laundry as quickly as any machine can do it. Our most recent announcements were met with happy shouts and cries. Every parent tries to teach their children good eating habits.

low in possible context: My boss granted me the furlough I spent in the Great Lakes Region. In Buffalo there is a gigantic grocery store called Wegman's. I say hello to the friendly stock boy every time I go there. Its often below zero degrees up there near Canada. The tall willow I like parking under shelters my car from the snow.

rest in possible context: The warm breeze caressed us as we neared the summit that afternoon. We progressed on schedule until we saw the wild billy goats. Our presence caused some unrest when the herd leader spotted

us watching. The new bulb I put in our aquarium hood fluoresced brightly. Now I am saving my nickels and dimes to vacation in Bucharest.

lay in possible context: They should relay the message to our parents in the countryside. During the winter we go to our private chalet to relax. We were expecting the delay to put us a bit behind schedule. I will mislay the directions if I am not terribly careful. The skilled dentist carefully repaired the milkman's gold inlay.

rise in possible context: A surprise awaited us once we reached the top of the steep hill. We had to be up by sunrise in order to hear the birds sing. Problems arise when I try to read a compass before breakfast. I was overwhelmed by the detailed map I had to memorize. They will authorize another hiking adventure next summer.

A.4. Experiment 4

well in impossible context: We won't welt in the sun because they made a lemonade stand. The folks who donated money for the swings have great wealth. They always bring their cute Welsh Terriers when they visit us. The construction workers must weld together the metal frame. A boy playing in the sand found the remains of a white whelk.

win in impossible context: The cold wind caused icicles to form along the rooftop edge. Every time we reached the top of the big green hill I would wince. The first team to reach the summit and get the puzzle assembled wins a prize. We dragged a heavy toolbox and winch through the obstacle course. Whoever lost would not cry and whinge much because it was fun.

rye in impossible context: Our inherent right to travel in space is often denied. I envy the astronauts that ride to the big yellow moon. I have heard that it is made of cheese and somewhat rife with holes. If the entire place is made of food it must smell quite ripe. Their space shuttles will go over that rise in the distance.

may in impossible context: The old mate living on the ship never recycles his cans. The weekend maid told me that she was a bit surprised by this. It would make her wonder if the old man cared about the sea. Next time I get a chance I will have to send mail to her. She seems happiest when riding her horse with the tattered mane.

well in possible context: That welder is putting together our brand new monkey bars. A short Welsher fellow raised the money for the new playground. My teacher said he is quite wealthy but usually cheap. The mayor's attendance at the ground breaking will cause a great welter. My best friend will arrive early because she must welcome him.

win in possible context: The last time we visited Windsor it was rather chilly out. Even though the weather was nippy we all avoided whingeing. A bald winsome fellow was jogging in shorts and a t-shirt. Travelling in

the fall and winter is easiest for us. We watched the parade from that incredibly grand window.

rye in possible context: If you look towards the stage you can see the pilot rising now. He was such a great writer even before he learned to fly planes. I guess when a man practices so much his skills must ripen. In a small jet the unfortunate rider has little space. The bright rifle case he had was just a toy made out of plastic.

may in possible context: The card mailer wrote to let us know it would rain cats and dogs. So the highly anticipated maiden voyage was postponed. The party we have planned will surely lead to shipyard mania. The crew chief encouraged making this celebration a big deal. The bird mated and raised its chicks in a nest below the bridge.

A.5. Experiment 5

well in impossible context: can, welt, dab, wealth, sing, broom, weld, tote, whelk, soul, pop, welsh.

win in impossible context: tame, wind, belt, whinge, stall, lap, wins, ten, winch, monk, wince, land.

rye in impossible context: paint, right, sub, ride, zone, pot, rife, beg, ripe, board, gem, rise.

may in impossible context: bike, mate, leaves, maid, grand, bell, make, tart, mail, craft, mane, mutt.

well in possible context: candle, welter, dabble, welsher, icing, broomstick, wealthy, total, welder, solar, popcorn, welcome.

win in possible context: tamer, windsor, beltway, whingeing, stalwart, laptop, winsome, tension, winter, monkey, window, lantern.

rye in possible context: painter, writer, subject, rider, ozone, spotting, ripen, beggar, rifle, border, gemstone, rising.

may in possible context: bikerack, mated, leaves, maiden, grandeur, belfry, making, starting, mania, mutter, mailer, crafty.

References

- Allopena, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, *38*, 419–439.
- Bloomfield, L. (1933). *Language*. New York: Holt.
- Bolinger, D. L., & Gerstman, L. J. (1957). Disjuncture as a cue to constraints. *Word*, *13*, 246–255.
- Brent, M. R., & Cartwright, T. A. (1996). Distributional regularity and phonotactic constraints are useful for segmentation. *Cognition*, *61*, 93–125.
- Cairns, P., Shillcock, R., Chater, N., & Levy, J. (1997). Bootstrapping word boundaries: A bottom-up corpus-based approach to speech segmentation. *Cognitive Psychology*, *33*, 111–153.

- Church, K. W. (1987). *Phonological parsing in speech recognition*. Dordrecht: Kluwer Academic Publishers.
- Cole, R., & Jakimik, J. (1980). A model of speech perception. In R. Cole (Ed.), *Perception and production of fluent speech* (pp. 136–163). Hillsdale, NJ: Erlbaum.
- Cutler, A., & Butterfield, S. (1992). Rhythmic cues to speech segmentation: Evidence from juncture misperception. *Journal of Memory and Language*, *31*, 218–236.
- Cutler, A., & Carter, D. M. (1987). The predominance of strong initial syllables in the English vocabulary. *Computer Speech and Language*, *2*, 133–142.
- Cutler, A., Demuth, K., & McQueen, J. (2002). Universality and language-specificity in listening to running speech. *Psychological Science*, *13*, 258–262.
- Cutler, A., McQueen, J. M., Norris, D. G., & Somejuan, A. (2001). The roll of the silly ball. In E. Dupoux (Ed.), *Language, brain and cognitive development: Essays in honor of Jacques Mehler* (pp. 181–194). Cambridge, MA: MIT Press.
- Cutler, A., & Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception & Perception*, *14*, 113–121.
- Cutler, A., & Otake, T. (2002). Rhythmic categories in spoken-word recognition. *Journal of Memory and Language*, *46*, 296–322.
- Echols, C. H., Crowhurst, M. J., & Childers, J. B. (1997). Perception of rhythmic units in speech by infants and adults. *Journal of Memory and Language*, *36*, 202–225.
- Goldinger, S. D., Luce, P. A., & Pisoni, D. B. (1989). Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory and Language*, *28*, 501–528.
- Goodman, M. B., & Jusczyk, P. W. (2000). The onset of sensitivity to internal syllable structure. In *Papers in Laboratory Phonology V: Acquisition and the Lexicon* (pp. 228–239). Cambridge: Cambridge University Press.
- Gow, D. W., & Gordon, P. C. (1995). Lexical and prelexical influences on word segmentation: Evidence from priming. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 344–359.
- Hockett, C. A. (1955). *A manual of phonology*. Baltimore, MD: Waverly Press.
- Houston, D. M. (1999). *The role of talker variability in infant word representations*, Ph.D. Thesis in Psychology. Baltimore, MD: Johns Hopkins University.
- Houston, D. M., Jusczyk, P. W., Kuijpers, C., Coolen, R., & Cutler, A. (2000). Cross-language segmentation by 9-month-olds. *Psychonomic Bulletin & Review*, *7*, 504–509.
- Johnson, E. K., & Jusczyk, P. W. (2001). Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language*, *44*, 548–567.
- Jusczyk, P. W., & Aslin, R. N. (1995). Infant's detection of sound patterns of words in fluent speech. *Cognitive Psychology*, *29*, 1–23.
- Jusczyk, P. W., Cutler, A., & Redanz, N. (1993a). Preference for the predominant stress patterns of English words. *Child Development*, *64*, 675–687.
- Jusczyk, P. W., Friederici, A. D., Wessels, J., Svenkerud, V. Y., & Jusczyk, A. M. (1993b). Infants' sensitivity to the sound patterns of native language words. *Journal of Memory and Language*, *32*, 402–420.
- Jusczyk, P. W., Goodman, M., & Bauman, A. (1999). 9-month-olds' attention to sound similarities in syllables. *Journal of Memory and Language*, *40*, 62–82.
- Jusczyk, P. W., & Hohne, E. A. (1997). Infants' memory for spoken words. *Science*, *277*, 1984–1986.
- Jusczyk, P. W., Hohne, E. A., & Bauman, A. (1999a). Infants' sensitivity to allophonic cues for word segmentation. *Perception & Psychophysics*, *61*, 1465–1476.
- Jusczyk, P. W., Houston, D., & Newsome, M. (1999b). The beginnings of word segmentation in English-learning infants. *Cognitive Psychology*, *39*, 159–207.
- Jusczyk, P. W., Luce, P. A., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language*, *33*, 630–645.
- Lehiste, I. (1960). *An acoustic-phonetic study of internal open juncture*. New York: S. Karger.

- Lehiste, I. (1972). The timing of utterances and linguistic boundaries. *Journal of the Acoustical Society of America*, *51*, 2018–2024.
- Luce, P. A., Goldinger, S. D., Auer, E. T., & Vitevitch, M. (2000). Phonetic priming, neighborhood activation, and PARSYN. *Perception & Psychophysics*, *62*, 615–625.
- Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing*, *19*, 1–36.
- van der Lugt, A. (2001). The use of sequential probabilities in the segmentation of speech. *Perception & Psychophysics*, *63*, 811–823.
- Marslen-Wilson, W. D. (1989). Access and integration: Projecting sound onto meaning. In W. D. Marslen-Wilson (Ed.), *Lexical access and representation* (pp. 3–24). Cambridge, MA: Bradford.
- Marslen-Wilson, W. D., & Warren, P. (1994). Levels of perceptual representation and process in lexical access: Words, phonemes, and features. *Psychological Review*, *101*, 653–675.
- Mattys, S. L., & Jusczyk, P. W. (2001). Phonotactic cues for segmentation of fluent speech by infants. *Cognition*, *78*, 91–121.
- Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*, *38*, 465–494.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, *18*, 1–86.
- McQueen, J. M. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language*, *39*, 21–46.
- McQueen, J. M., Cutler, A., Briscoe, T., & Norris, D. (1995). Models of continuous speech recognition and the contents of the vocabulary. *Language and Cognitive Processes*, *10*, 309–331.
- McQueen, J. M., Norris, D. G., & Cutler, A. (1994). Competition in spoken word recognition: Spotting words in other words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 621–638.
- McQueen, J. M., Norris, D. G., & Cutler, A. (1999). Lexical influence in phonetic decision-making: Evidence from subcategorical mismatches. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1363–1389.
- Mehler, J., Dupoux, E., & Segui, J. (1990). Constraining models of lexical access: The onset of word recognition. In G. T. M. Altmann (Ed.), *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives* (pp. 236–262). Cambridge, MA: MIT Press.
- Moon, C., Bever, T. G., & Fifer, W. P. (1992). Canonical and non-canonical syllable discrimination by two-day-old infants. *Journal of Child Language*, *19*, 1–17.
- Morgan, J. L. (1996). A rhythmic bias in preverbal speech segmentation. *Journal of Memory and Language*, *35*, 666–688.
- Myers, J., Jusczyk, P. W., Kemler Nelson, D. G., Charles-Luce, J., Woodward, A., & Hirsh-Pasek, K. (1996). Infants' sensitivity to word boundaries in fluent speech. *Journal of Child Language*, *23*, 1–30.
- Nazzi, T., Jusczyk, P. W., & Johnson, E. K. (2000). Language discrimination by English-learning 5-month-olds: Effects of rhythm and familiarity. *Journal of Memory and Language*, *43*, 1–19.
- Newman, R. S., & Jusczyk, P. W. (1996). The cocktail party effect in infants. *Perception & Psychophysics*, *58*, 1145–1156.
- Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition*, *52*, 189–234.
- Norris, D., McQueen, J. M., & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary. *Brain and Behavioral Sciences*, *23*, 299–325.
- Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology*, *34*, 191–243.

- Norris, D., McQueen, J. M., Cutler, A., Butterfield, S., & Kearns, R. (2001). Language-universal constraints on speech segmentation. *Language and Cognitive Processes*, 16, 637–660.
- Pinker, S. (1984). *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Santelman, L., Houston, D., & Jusczyk, P. W. (1997). 7.5-month old's segmentation of multisyllabic words in fluent speech. In *Proceedings of the 21st annual BU conference on language development, Boston, MA*.
- Shillcock, R. C. (1990). Lexical hypotheses in continuous speech. In G. T. M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computation perspectives* (pp. 24–49). Cambridge, MA: MIT Press.
- Suomi, K. (1993). An outline of a developmental model of adult phonological organization and behavior. *Journal of Phonetics*, 21, 29–60.
- Tabossi, P., Burani, C., & Scott, D. (1995). Word identification in fluent speech. *Journal of Memory and Language*, 34, 440–467.
- van de Weijer, J. (1998). *Language input for word discovery*, Unpublished Ph.D. University of Nijmegen, Nijmegen.
- Vitevitch, M. S., & Luce, P. A. (1998). When words compete: Levels of processing in the perception of spoken words. *Psychological Science*, 9, 325–329.
- Vitevitch, M. S., & Luce, P. A. (1999). Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory and Language*, 40, 374–408.
- Vrooomen, J., & de Gelder, B. (1995). Metrical segmentation and lexical inhibition in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 98–108.
- Vrooomen, J., & de Gelder, B. (1997). Activation of embedded words in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 710–720.
- Woodward, Z. N., & Aslin, R. N. (1990). Segmentation cues in maternal speech to infants. *Poster presented at the International Conference on Infancy Studies*, Montreal, Canada.