Pattern Induction by Infant Language Learners

Jenny R. Saffran and Erik D. Thiessen
University of Wisconsin—Madison

How do infants learn the sound patterns of their native language? By the end of the 1st year, infants have acquired detailed aspects of the phonology and phonotactics of their input language. However, the structure of the learning mechanisms underlying this process is largely unknown. In this study, 9-month-old infants were given the opportunity to induce specific phonological patterns in 3 experiments in which syllable structure, consonant voicing position, and segmental position were manipulated. Infants were then familiarized with fluent speech containing words that either fit or violated these patterns. Subsequent testing revealed that infants rapidly extracted new phonological regularities and that this process was constrained such that some regularities were easier to acquire than others.

Months before infants speak their first words, they have acquired extensive and detailed knowledge about the sound patterns of their native language. Indeed, sounds are the infant's entrance point into spoken language acquisition, beginning with the rhythmic patterns of the infant's language—knowledge acquired before birth (e.g., Mehler et al., 1988)—and extending during the 1st year to include language-typical phonological patterns such as lexical stress cues (Jusczyk, Cutler, & Redanz, 1993; Jusczyk, Houston, & Newsome, 1999), the phonemic inventory of the native language (e.g., Kuhl, Williams, Lacerda, Stevens, & Lindblom, 1992; Werker & Tees, 1984), language-specific phonotactic patterns governing the placement and combinations of phonemes (Friederici & Wessels, 1993; Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk, 1993; Jusczyk, Luce, & Charles-Luce, 1994), and the relationship between all of these types of information and the newly emerging lexicon (Mattys & Jusczyk, 2001; Mattys, Jusczyk, Luce, & Morgan, 1999).

Although numerous studies document infants' attainments in native-language acquisition, relatively little research has addressed the question of how this learning occurs. One promising type of methodology for investigating learning mechanisms in infancy is the use of artificial languages—miniature systems mirroring some aspect of natural languages—in which specific cues of interest are isolated (for a review, see Gomez & Gerken, 2000). After a brief exposure to an artificial language, learners are tested (a) to determine whether the cues in the input were sufficient for acquisition or (b) to assess the relative strengths of different cues. Adapting this method from the adult learning literature, researchers have used artificial languages to investigate the learning capacities underlying the discovery of words in fluent speech, including the use of statistical and phonological cues to word boundaries (e.g., Goodsitt, Morgan, & Kuhl, 1993; Johnson & Jusczyk, 2001; Morgan & Saffran, 1995; Saffran, Aslin, & Newport, 1996; Thiessen & Saffran, 2001). Similarly, artificial-language methods have been profitably applied to the beginnings of syntax acquisition by infants (e.g., Gomez, 2002; Gomez & Gerken, 1999; Marcus, Vijayan, Bandi Rao, & Vishton, 1999) and are beginning to be used to investigate the acquisition of phonemic contrasts (e.g., Maye, Werker, & Gerken, 2002).

In the current research, we applied the logic of the artificial-language methodology to the question of infant phonological and phonotactic learning. Phonological regularities govern many aspects of sound structure, including constraints on syllable structures, constraints on sound patternings across syllable boundaries, and stress rules involving multiple syllables. Phonotactic regularities constrain the positions of particular phonemes relative both to syllable boundaries and to one another. For example, /fs/ can occur at the end, but not the beginning, of syllables in English. Phonotactics are not a simple function of pronounceability; sequences that are legal in some languages are illegal in others. Despite the complexity of phonotactic structure, it is acquired during the 1st year of life; 9-month-old infants prefer to listen to phonotactically legal sequences (Friederici & Wessels, 1993; Jusczyk, Friederici, et al., 1993) and even prefer to listen to frequently exemplified phonotactic structures over equally legal but rarer structures (Jusczyk et al., 1994). By the end of the 1st year, infants are able to use their developing knowledge about phonotactic structure as a cue to word boundaries, segmenting novel words out of continuous speech more readily at points of low phonotactic probability (Mattys & Jusczyk, 2001; Mattys et al., 1999). These generalizations require abstraction across a corpus of linguistic input and differ substantially from language to language, requiring a powerful learning process. At the same time, there are many types of sound patterns that never occur cross-linguistically, suggesting a role for constraints on the learning process.

To address the acquisition of phonotactic knowledge in infancy, we combined the artificial-language-learning methodology with a rather different technique used to study the acquisition of phonotactic patterns by adults. Dell and colleagues, using a word production paradigm, embedded experiment-wide phonotactic constraints within a reading task (Dell, Reed, Adams, & Meyer, 2000; Gupta & Dell, 1999). Adult participants spoke lists of syllables that conformed to both language-wide phonotactics and experiment-wide phonotactics; for example, half the participants experienced /f/ as an onset and /s/ as a coda, and the other half experienced the reverse. The speech errors of participants reflected both the phonotactics of their native language and the experiment-induced phonotactics, which suggests that participants implicitly learned the phonotactic regularities that characterized the syllables they repeated (for related recent research, see Onishi, Chambers, & Fisher, 2002).

In the current experiments, we asked whether 9-month-old infants could learn experiment-wide phonological and phonotactic structures. By manipulating the experiment-wide patterns, we hoped to begin to uncover the types of cues detected and used by infant learning mechanisms. The structures exemplified in these experiments represent a subset of acceptable English patterns. Infants were first exposed to a list of nonsense words conforming to the phonological or phonotactic structure of interest. This portion of the experimental session, referred to as the pattern induction phase, served as a brief learning experience during which infants had the opportunity to induce phonological or phonotactic patterns characteristic of possible words. Infants then received a word segmentation task designed to ask whether they had learned the patterns present during the pattern induction phase. These materials consisted of four new nonsense words (not heard during the pattern induction phase), two that conformed to the template from the pattern induction phase and two that did not. These words were presented continuously as fluent speech, with no acoustic markers at word boundaries.

If infants are unable to acquire the patterns during the pattern induction phase of the experiment, then the four new words should be equally difficult to segment from fluent speech during the segmentation phase of the experiment. However, if infants successfully induce these patterns, then the two novel words consistent with the newly acquired patterns should be more readily detected and segmented from fluent speech than should the other two words. To test this hypothesis, when infants entered the third phase of the experiment, we tested them with repetitions of each of the four words heard during the segmentation phase. If infants more rapidly segmented the words that fit the newly induced pattern than the words that did not, we expected them to discriminate between the two types of words during testing, as measured by different listening times.

Because prior studies have not included multiple learning phases (first pattern induction, then segmentation, followed by testing) within a single session, Experiment 1 was designed to investigate the induction of a phonological pattern that should be salient to infant learners: syllable structure. In Experiment 2, we asked whether infants could acquire a phonotactic generalization concerning restrictions on voiced and voiceless stops as a function of position within the syllable. Finally, in Experiment 3 we sought to uncover the nature of the learning mechanism by assessing the ease with which infants acquired different types of sound patterns.

This last question is of considerable theoretical interest because constraints on infants' language-learning capacities may serve as an important force in determining the course of language development. In addition, constraints on infant learners may play a role in shaping natural languages: Generalizations that are easier to acquire may be more likely to emerge and persist in the inventory of natural languages (e.g., Bever, 1970; Christiansen & Devlin, 1997; Ellefson & Christiansen, 2000; Newport, 1982, 1990; Saffran, 2002).

Experiment 1

As a first step in the investigation of the learning mechanisms underlying linguistic sound patterns, we asked whether 9-month-old infants could acquire restrictions on syllable structure. Languages vary greatly in the inventory of syllable structures permitted by their phonology. English is a permissive language, allowing syllable structures (V = vowel; C = consonant) ranging from open syllables with zero consonants (V, as in the word *I*) or one initial consonant (CV, as in the word *she*) to closed syllables containing multiple consonant clusters (CCCVCCC, as in the word *splints*). Other languages are far more restrictive; for example, the syllable inventory of Somali is limited to {V, CV, VC, CVC} (Kenstowicz, 1994). These differences are likely to be salient to infant learners given their well-documented interest in the rhythmic properties of speech. We designed Experiment 1 to ask whether infants exposed to a particular syllable structure could rapidly induce that structure and apply it to novel words in a word segmentation task. During the pattern induction phase, infants listened to a list of nonsense words that conformed to a specific syllable template. Infants were subsequently familiarized with continuous speech containing novel words that either did or did not conform to the syllable template heard during pattern induction. We then used the head-turn-preference procedure to determine whether infants more readily segmented words from continuous speech when they conformed to the syllable structures heard earlier in the experiment. If infants are able to use their knowledge of recently heard syllable structures in processing novel words, then we would expect listening times during testing to differ depending on whether test items conformed to the syllable structures heard during pattern induction.[1]

Method

Subjects. Thirty full-term 9-month-old monolingual infants with no history of recurrent ear infections were tested (mean age = 9 months 0 weeks; range = 8 months 3 weeks to 9 months 2 weeks). Half of the infants were randomly assigned to the CVCV condition, and half were assigned to the CVCCVC condition. Twenty-two additional infants were tested but not included in the analysis for the following reasons: fussiness (7), looking times averaging less than 3 s to one or both sides (6), parental interference

---

[1] We used the preferential listening methodology as an index of what infants could learn, as assessed by discrimination between test items. The *direction* of preference reflects other factors not relevant to the hypotheses tested here, such as the speed of the infant's learning and the degree of habituation (e.g., Hunter & Ames, 1988). It is the *consistency* of the direction of listening preferences across infants in an experiment, rather than the direction itself (familiarity vs. novelty), that is relevant for data analysis and interpretation in these experiments.

(4), experimenter error (3), falling asleep (1), and not looking at the side lights (1). All infants in this and the subsequent experiments were solicited from local birth announcements and hospital records, and parental consent was obtained prior to testing in accordance with the guidelines of the local human subjects review committee and the principles of ethical treatment established by the American Psychological Association.

Stimuli. For the pattern induction phase, we generated two lists of bisyllabic nonsense words. One of the lists consisted of 30 CVCV words (e.g., *boga, diku*), and the other list consisted of 30 CVCCVC words (e.g., *bikrub, gadkug*). Both lists contained the same set of phonemes (see Appendix A for word lists). A synthesized monotone female voice read each list twice, with a 0.5-s silence between words (MacinTalk running on an Apple Quadra 650 computer, with the Victoria voice). The output of the synthesizer was digitized using SoundEdit and stored on disk at a 22-kHz sampling rate, with the CVCV and CVCCVC lists in separate files. The pattern induction stimuli thus consisted of a digitized 60-word list conforming to either the CVCV template or the CVCCVC template and lasting approximately 2 min (CVCV, 2 min 2 s; CVCCVC, 2 min 11 s).

For the segmentation phase, we generated a continuous speech stream consisting of four new bisyllabic nonsense words not heard during the pattern induction phase. The same synthesis procedures were used as in the pattern induction phase; no editing was necessary to create the continuous speech. Two of the words fit the CVCV template (*baku, dola*), and two of the words fit the CVCCVC template (*tupgod, girbup*). The four words were repeated in random order, with the constraint that the same word not occur twice in a row. Each word occurred 26 times, for a total of 104 words. The synthesis procedures used to create the pattern induction materials were used to generate a continuous speech stream with no acoustic or prosodic cues to word boundaries (e.g., *tupgodbakugirbupdolabaku . . .*). The speech stream lasted approximately 60 s and was digitized for playback during the experimental session.

The test items included the four words used during the segmentation phase: *baku, dola, tupgod,* and *girbup.* Each test trial consisted of repetitions of a single word, with a 0.5-s silence between repetitions. For infants in the CVCV condition, *baku* and *dola* conformed to the familiar syllabic template heard during the pattern induction phase, and *tupgod* and *girbup* did not; infants in the CVCCVC condition experienced the opposite pattern of familiarity and novelty. This between-subjects counterbalanced design ensured that any observed preferences were due to learning. Each test item was synthesized and digitized for playback during testing.

Procedure. Infants were tested individually while seated in a parent's lap in a sound-attenuated booth. An observer outside the booth monitored the infant's looking behavior on a closed-circuit TV system and coded the infant's behavior using a button-box connected to a PC. This button-box was used to initiate trials and to enter the direction of the infant's head turns, which controlled the duration of each test trial. Both the parent and the observer listened to masking music over headphones to eliminate bias.

During the 2-min pattern induction phase, either the CVCV or the CVCCVC word list played continuously from two loudspeakers (one located on each of the two side walls in the booth). No lights were used during this portion of the procedure so that infants would not become disinterested in the lights before reaching the test phase.

At the beginning of the 1-min segmentation phase, the infant's gaze was first directed to a blinking light on the front wall in the testing booth. Then the continuous speech stream was presented without interruption from the two loudspeakers. To acclimate the infants to the lights used during testing, we lit and extinguished a blinking light above one of the two loudspeakers (randomly selected) depending on the infant's looking behavior. When this blinking sidelight was extinguished, the central blinking light was illuminated until the infant's gaze returned to center, and another blinking sidelight was then presented to elicit the infant's gaze. During the entire segmentation phase, there was no contingency between lights and sound, which played continuously.

Immediately after the segmentation phase, 12 test trials were presented. The trials were broken into three blocks, each consisting of one trial testing each of the four test items; trials were randomized within each block. Each test trial consisted of a single test word repeated with a 0.5-s silence between repetitions. All infants heard the same 12 trials: six test items conforming to the phonological template heard during pattern induction (familiar words) and six test items not conforming to the phonological template heard during pattern induction (novel words). Each test trial began with the blinking light on the front wall. When the observer signaled the computer that the infant was fixating this central light, one of the lights on the two side walls began to blink and the central light was extinguished. When the observer judged that the infant had made a head turn of at least 30° in the direction of the blinking sidelight, a button press signaled to the computer that one of the test items should be presented from the loudspeaker adjacent to the blinking light. This test item was repeated until the observer coded the infant's head turn as deviating away from the blinking light for 2 consecutive seconds. When this look-away criterion was met, the computer extinguished the blinking sidelight, turned off the test stimulus, and turned on the central blinking light to begin another test trial. The computer randomized the order of test trials (three for each of the four test items) and cumulated the total looking time to each test item. Test trials automatically ended after a maximum of 12 word repetitions.

Results and Discussion

The first analysis contrasted infants assigned to the two pattern induction conditions, the CVCV condition and the CVCCVC condition. Because listening-time differences between familiar and novel words did not differ as a function of pattern induction condition, $t(28) = 1.2$, $p = .21$ (all tests are two-tailed), the data from the CVCV and CVCCVC conditions were pooled in the subsequent analyses.

The principal hypothesis concerned the difference in listening times between familiar and novel words—those items from the segmentation phase that conformed to the phonological template heard during the pattern induction phase and those that did not.[2] There was a significant difference in listening times for familiar versus novel words: $t(28) = 2.2$, $p < .05$. Infants listened longer to familiar words ($M = 7.52$ s, $SE = 0.40$ s) than to novel words ($M = 6.86$ s, $SE = 0.41$ s). This difference suggests that infants did induce the syllable structures present during pattern induction: Infants exposed to CVCV words during pattern induction showed a pattern of test performance (longer times to CVCV items than to CVCCVC test items) different from that of infants exposed to CVCCVC words during pattern induction (longer times to CVCCVC test items than to CVCV test items).

The results of Experiment 1 suggest that infants can acquire knowledge about phonological structure following a very brief exposure period. Infants showed a different pattern of test performance depending on the syllable structures presented in the word list played during pattern induction. We interpret these results as indicating that infants brought knowledge acquired during pattern induction to bear on the segmentation task: The words in fluent speech that best fit infants' expectations about word forms (based

---

[2] Both types of test items were equally familiar with respect to the segmentation phase: Infants heard all four words equally often in the speech stream. Instead, "familiarity versus novelty" refers to the syllable structure of the test items, which either conformed (familiar) or did not conform (novelty) to the syllable structure heard during the pattern induction phase.

on learning during the pattern induction phase) were located and segmented most rapidly, with later segmentation of the words containing relatively novel syllable structures. According to this hypothesis, listening times for familiar and novel test words differed because infants segmented the former more rapidly than the latter and consequently received different amounts of familiarization for the two types of words during the segmentation phase. Another possibility is that the pattern induction learning experience may have served to prime responses during testing: Exposure to one type of word form may have influenced test preferences for that form over other forms regardless of the presence of the segmentation task (see also Aslin, 2000).[3] Both types of explanations support our principal hypothesis: Infants can extract phonological regularities from a list of novel words presented during pattern induction. Note that these results reflect the infants' in-laboratory experiences; infants in the CVCCVC condition have certainly heard many English words with CVCV syllable structure, and vice versa. The pattern induction experience apparently serves to influence infants' subsequent processing, at least in the very short term, despite months of prior exposure to many different types of patterns.

Although these results suggest that the pattern induction phase affected subsequent performance, it remains unclear exactly which aspects of the patterns were acquired by the infants. As suggested above, the infants may have attended to information about syllable structure. Alternatively, infants may have acquired knowledge of expected lexical duration. CVCCVC items are inherently longer than CVCV items; in the case of the present materials, the average duration of CVCCVC items was 759 ms, whereas the average duration of CVCV items was 625 ms. Thus, it is unclear whether infants' differential test performance was due to the induction of syllable structure or to the relative familiarity or novelty of test item duration. Nevertheless, infants must have acquired phonological and/or acoustic patterns during the induction phase, and this knowledge influenced subsequent performance.

Syllable structure and timing differences are a fairly coarse-grained aspect of linguistic sound structure. The literature suggests that infants are able to represent other more subtle aspects of the sound structure of languages that enter into the phonotactic constraints governing sound patternings in the native language. For example, 6-month-old infants recognize the acoustic categories correlated with vowel identities despite the dissimilarity of vowels produced by different speakers and in different contexts (e.g., Kuhl, 1979, 1983). Similarly, infants can organize syllables according to the first consonant's place (e.g., Hillenbrand, 1984) and manner (e.g., Jusczyk, Goodman, & Baumann, 1999) of articulation despite acoustic differences between consonant tokens.

One important aspect of sound patterning that infants must learn concerns consonant voicing (e.g., the difference between the voiceless /p/ and the voiced /b/).[4] Infants are attuned to voicing differences between minimal pairs, such as /pa/ versus /ba/, at 1 month of age for stops (e.g., Eimas, Siqueland, Jusczyk, & Vigorito, 1971) and at 4 months of age for fricatives such as /s/ versus /z/ (e.g., Eilers, 1977; Eilers, Wilson, & Moore, 1977). Infants must also learn that the voicing status of a consonant affects its use. Languages contain phonological generalizations involving voicing, such as the English plural, in which the voiceless /s/ becomes the voiced /z/ when following a voiced consonant, such as /dogz/. Such patterns can be quite language-specific; for example, only voiceless fricatives such as /s/ and /z/ can precede liquids like /l/ in English, whereas other languages permit voiced fricatives in the preliquid position. In order to learn such patterns, learners must attend to consonants' voicing status.

In Experiment 2, we used the pattern induction methodology to assess the acquisition of voicing restrictions on segments in specific syllable positions. Infants received a pattern induction word list in which phoneme voicing was manipulated, with restrictions on voicing in particular syllable positions. Infants next received a segmentation task containing four novel words, two of which conformed to the voicing patterns heard during pattern induction. Infants were subsequently tested to determine whether the voicing patterns present during pattern induction affected their listening preferences. Experiment 3 was designed to uncover the specific aspects of the phonological structure acquired by infants in Experiment 2: Did infants learn restrictions on consonant placement as a function of phoneme voicing—restrictions involving categories of phonemes—or did infants learn positional restrictions for specific phonemes? Experiments 2 and 3 were designed in tandem to ask not just what types of regularities could be detected and learned by infants, but how those regularities were detected. These experiments, taken together, allowed us to directly investigate the types of information used by infants in the phonotactic learning task, which in turn may suggest possible constraints on phonological learning by infants.

Experiment 2

Infants in Experiment 2 were confronted with a particular pattern of consonant voicing in bisyllabic words during pattern induction. Half of the infants heard words in which syllable-initial consonants were voiceless and syllable-final consonants were voiced (e.g., *todkad*, referred to below as the −V+V condition), and the other half of the infants heard words with the opposite pattern of voicing (e.g., *dakdot*, referred to below as the +V−V condition). After exposure to these pattern induction materials, infants next entered the segmentation phase, in which they were familiarized with continuous speech containing four new words, two consistent with each of the two voicing patterns. As in Experiment 1, infants were then tested on all four words from the segmentation phase. We hypothesized that if infants learned the voicing patterns present during pattern induction, they should show a difference in listening times for words fitting the familiar pattern and words fitting the novel pattern. Like the previous experiment, this study was designed to look for discrimination between the novel and familiar patterns as assessed by a consistent direction of preference across infants; the actual direction of the preference is affected by other task variables, such as the difficulty of the segmentation task and test discrimination, that are not relevant to our hypothesis.

Method

Subjects. Thirty full-term 9-month-old monolingual infants with no history of recurrent ear infections were tested (mean age = 9 months 0

---

[3] We thank LouAnn Gerken for this suggestion.

[4] Note that we are not claiming that the phonetic feature "± voicing" is psychologically real. Instead, we are using voicing as shorthand for an acoustic dimension of phonetic similarity.

weeks; range = 8 months 3 weeks to 9 months 2 weeks). Half of the infants were randomly assigned to the –V+V condition, and half were assigned to the +V−V condition. Twenty-nine additional infants were tested but not included in the analysis for the following reasons: fussiness (8), looking times averaging less than 3 s to one or both sides (7), parental interference (5), not looking at the sidelights (3), insufficient data for one trial type (2), and participation in two or more prior studies in the laboratory (4).[5]

Stimuli. For the pattern induction phase, we generated two lists of bisyllabic nonsense words. One of the lists consisted of 30 words in which syllables began with voiceless stops and ended with voiced stops (e.g., *todkad, pibtug*). These materials are referred to as the –V+V condition. The second list consisted of the same 30 words with phonemes in reverse order such that the positions of the syllable-initial and syllable-final phonemes were reversed (e.g., *dakdot, gutbip*). These materials are referred to as the +V−V condition. Both lists contained the same set of phonemes (see Appendix B for word lists). The lists were synthesized and digitized as in Experiment 2, with each audiofile consisting of the 30-word list repeated twice. The pattern induction stimuli thus consisted of a digitized 60-word list conforming to either the −V+V template or the +V−V template and lasting approximately 2 min 12 s.

In natural speech, vowels before voiced consonants are longer than vowels preceding voiceless consonants (e.g., Crystal & House, 1988; DiSimoni, 1974). This was also the case for the synthetic stimuli used here: Vowels preceding voiced consonants (–V+V condition) were 14 ms longer, on average, than vowels preceding voiceless consonants (+V−V condition). Because the average vowel duration in these stimuli was 183 ms, the change in vowel duration that was due to the voicing of the following consonant was small (13%), which is consistent with the results of Crystal and House (1988). However, other researchers have found that vowels preceding voiced consonants in natural speech can be up to twice as long as vowels preceding voiceless consonants, and reports of a 3:2 ratio are common (e.g., DiSimoni, 1974; House, 1961; Peterson & Lehiste, 1960). Although differences of the magnitude reported by Crystal and House (1988) have not been found to play an important role in the perception of the voiced/voiceless contrast, differences on the order of 3:2 and greater can have a significant effect on the perception of the following consonant (Krause, 1982); for example, adults and children can use vowel length alone to distinguish between voiced and voiceless consonants (e.g., Hogan & Rozyspal, 1980; Raphael, 1972).

Although vowel lengthening is a natural side effect of voicing, it presents a possible confound in the pattern induction materials used in this experiment. Rather than learning consonant voicing patterns, infants might have considered the test items to be familiar or novel on the basis of their relative vowel durations: Items corresponding to the –V+V template contained slightly longer vowels than items corresponding to the +V−V template. However, the brevity and variability of the vowel-length differences make this possibility unlikely. Eilers (1977) found that 3–4-month-old infants provided no evidence of discrimination of a 100-ms vowel duration difference between /at/ and /a:t/. Even by 12 months of age, infants discriminated this kind of contrast very poorly (Eilers et al., 1977). Similarly, although 5–11-month-old infants were able to discriminate between vowel pairs when one was 33% longer than the other (300 ms vs. 400 ms), their performance was poor (Eilers, Bull, Oller, & Lewis, 1984). These differences are markedly larger than the duration differences between vowels preceding voiced and voiceless consonants in the current stimuli. In addition, overall word length was equivalent in the two conditions because of the variability of the intrinsic vowel lengths and the phonetic contexts of the nonsense words (–V+V condition = 729.3 ms; +V−V condition = 729.1 ms). We thus consider it unlikely that infants ignored the consistent consonant voicing cues and instead acquired patterns and performed test discriminations based on vowel length, which would have required them to ignore word length and focus on vowel length with greater perceptual acuity than has been shown in any previous study.

As in Experiment 1, materials for the segmentation phase consisted of a continuous speech stream consisting of four new bisyllabic nonsense words not heard during pattern induction. Two of the words fit the –V+V template (*kibpug, pagkob*), and two of the words fit the +V−V template (*bupgok, gikbap*). Each of the four words was repeated 20 times, in random order, with the constraint that the same word not occur twice in a row, for a total of 80 words. The same synthesis procedures used for the pattern induction stimuli were used to generate a continuous speech stream with no acoustic or prosodic cues to word boundaries. The speech stream lasted approximately 60 s and was digitized for playback during the experimental session.

The test items included the four words used during the segmentation phase: *kibpug, pagkob, bupgok,* and *gikbap*. Each test trial consisted of a single word repeated with a 0.5-s silence between repetitions. For infants in the –V+V condition, *kibpug* and *pagkob* conformed to the familiar syllabic template heard during pattern induction, and *bupgok* and *gikbap* did not; infants in the +V−V condition experienced the opposite pattern of familiarity and novelty. Each test item was synthesized and digitized for playback during testing.

Procedure. The procedure was identical to that in Experiment 1.

Results and Discussion

The first analysis contrasted infants assigned to the two pattern induction conditions, the –V+V condition and the +V−V condition. Because listening-time differences between familiar and novel words did not differ as a function of pattern induction condition, $t(28) = 1.07$, $p = .29$ (all tests are two-tailed), the data from the –V+V and +V−V conditions were pooled in the subsequent analyses.

The principal hypothesis concerned the difference in listening times between familiar and novel words—those that conformed to the phonological template heard during pattern induction and those that did not. There was a significant difference in listening times for familiar versus novel words: $t(29) = 2.68$, $p = .01$. Infants listened longer to novel words ($M = 7.06$ s, $SE = 0.36$ s) than to familiar words ($M = 6.05$ s, $SE = 0.38$ s). This difference suggests that infants' pattern induction experiences affected their subsequent listening preferences: Infants exposed to –V+V words during pattern induction showed a pattern of segmentation different from that of infants exposed to +V−V words.[6]

Although the direction of preference is not crucial to interpreting these results, it is interesting that infants in Experiment 2 showed a novelty preference whereas infants in Experiment 1 showed a familiarity preference. A number of factors influence infants' direction of preference, including the complexity of the learning problem (e.g., Aslin, 2000; Hunter & Ames, 1988), the match between the familiarization and test stimuli (e.g., Thiessen

---

[5] Pilot data for Experiment 2 revealed differences between infants for whom this was their first or second experimental participation in our laboratory and infants who had additional experience in the laboratory. The analyses for Experiments 2 and 3 thus excluded infants with extensive prior experience with our testing procedure. No infants in Experiment 1 had participated in more than one prior experiment.

[6] The novelty preference is unlikely to be due to a priming effect from the pattern induction materials, because priming should result in familiarity preferences. This suggests that the results of Experiment 1 are more parsimoniously interpreted as the function of differential learning during the segmentation phase, as per our original hypothesis, than as the result of priming between the pattern induction and test phases.

& Saffran, in press), and the amount of processing allowed by the task (e.g., Roder, Bushnell, & Sasseville, 2000). In general, novelty preferences are associated with easy learning tasks, and familiarity preferences are associated with more difficult tasks. Here, infants showed a familiarity preference when asked to learn about syllable structure (Experiment 1) and a novelty preference when asked to learn about the relation between consonant voicing and syllable position (Experiment 2). This suggests that learning about voicing was the easier of the two tasks.

There are (at least) two possible explanations for this pattern of results. One is that the pattern induction materials used in Experiment 1 offered infants less information than those used in Experiment 2. There are four redundant sources of information available regarding the relationship between voicing and syllable position in a word such as *dutgap* (from Experiment 2): the voiced consonants that begin the two syllables and the voiceless consonants that end the two syllables. In Experiment 1, only two sources of information distinguished the CVCCVC and the CVCV conditions: the consonant (or lack thereof) at the end of the first syllable and the consonant (or lack thereof) at the end of the second syllable. In addition, the words in Experiment 2 contained relevant information in the word-initial position, a position that is likely to be more perceptually salient than word-medial or word-final positions (Cole, 1981; Cutting, 1974). In Experiment 1, no word-onset information was present to distinguish the CVCCVC and the CVCV conditions. The presence of less information overall, and less information in the salient word-initial position, may have conspired to make the patterns available in Experiment 1 more difficult to learn than the patterns available in Experiment 2.

Another possible explanation for the observed differences in direction of preference concerns the variability of the stimuli. In Experiment 2, both syllable structure (all of the words were CVC-CVC) and consonant voicing positions were consistent. However, in Experiment 1, only syllable structure was consistent: The words in the pattern induction period of Experiment 1 were more variable in other phonological dimensions. It is thus possible that the patterns in Experiment 1 were more difficult to learn because of the increased variability of the stimuli; for example, there were no consistent voicing patterns. Regardless of the reason for the shift in direction of preference, it is evident that infants in both experiments were able to learn the patterns heard during the pattern induction period, as evidenced by the significant discrimination performance observed in both tasks.

The results of Experiment 2 suggest that infants can detect and learn relatively fine-grained phonotactic generalizations. To do so, infants must do more than notice which phonemes occur or do not occur in words; infants must track the correlations between phonemes and their positions within syllables and/or words. Infants in the current experiments could not have shown the pattern of performance they did just by remembering the phoneme /p/—they were also required to represent the possible positions of the phoneme /p/. How much detail the infants acquired in this experiment is unknown. For example, they may have relied on the phonemes' occurring in the initial position, which is consistent with research suggesting that infants are more likely to notice syllable-initial similarities than syllable-final similarities (e.g., Jusczyk, Goodman, & Baumann, 1999). Regardless, phonological learning must have occurred during the pattern induction phase of the experiment.

Experiment 3 was designed to probe more deeply into the learning process underlying the acquisition of new phonotactic patterns. The particular structures hypothesized to be acquired by infants in Experiment 2 concern voicing patterns—generalizations regarding which syllable positions are voiced and which are voiceless. However, infants may have acquired the materials in Experiment 2 in a fundamentally different way. Instead of learning restrictions on consonant positions as a function of voicing, infants might have learned about the privileges of occurrence of particular segments. That is, rather than learning that syllables begin with voiceless consonants and end with voiced consonants in the −V+V condition, infants might have learned something qualitatively different: Syllables can begin with /p/, /t/, and /k/ and end with /b/, /d/, and /g/. The distinction lies between acquiring a pattern that supports subsequent generalizations, such as the expectation that syllables can begin with voiceless consonants, and acquiring a pattern limited to specific exemplars—namely, the particular consonants heard during exposure.

Acquiring a voicing pattern and acquiring knowledge of specific segment placements are two very different ways of learning phonological structure. In the former case, the similarity between items with respect to voicing (or other phonetic or acoustic dimensions) serves as a cue for learning. A learner designed to detect such patterns of similarity would presumably learn more easily when such patterns are present, as they are in all phonological systems. In the latter case, similarity between exemplars is less important; the learner would be adept at acquiring local patterns such as the distribution of /k/ and the distribution of /p/ but would be less apt to notice and exploit the similarity between /k/ and /p/. Languages also contain a great deal of information of this type, most notably in word formation: One must learn that /tap/ is a word and that replacing the initial segment with another voiceless stop like /kap/ changes the meaning of the word. At the word level, acquiring knowledge about particular segment identities and positions is critical. In acquiring phonology, however, one must be able to generalize beyond specific phonemes.

To investigate what types of information infants attend to when confronted with a list of novel words, we made a minor alteration in the design of Experiment 2. Infants in Experiment 2 had two possible options for what they could learn. They may have acquired voicing patterns of the type –V+V, as hypothesized above. Alternatively, they may have been acquiring the particular phoneme tokens presented during pattern induction, learning that syllables can begin with /p/, /t/, and /k/ and end with /b/, /d/, and /g/. In order to distinguish between these two hypotheses, we altered the materials used in Experiment 2 by flipping the positions of /t/ and /d/ in the pattern induction word lists. This served to disrupt the voicing regularities: It was no longer possible to characterize these materials as –V+V or +V–V. The only remaining basis for learning consisted of the identities of the phonemes occurring in each position: In one condition, syllables could begin with /p/, /d/, and /k/ and end with /b/, /t/, and /g/, and in the second condition, the reverse was true. Infants then received the segmentation and test materials from Experiment 2. Because these materials contained neither /d/ nor /t/, they maintained in Experiment 3 the same familiarity versus novelty distinctions relative to the pattern induction materials as they did in Experiment 2.

The critical difference between Experiments 2 and 3 is that whereas infants in Experiment 2 could use voicing as a cue for

learning the phonotactic regularities during pattern induction, infants in Experiment 3 could not. If infants successfully discriminated between familiar and novel patterns during testing in Experiment 3, it would suggest that they are able to use segmental position in the absence of consistent voicing cues to acquire phonotactic patterns. However, if infants failed to discriminate in Experiment 3, it could not be attributed to the difficulty of the segmentation and test phases, because these were identical to those in Experiment 2, in which infants successfully discriminated during testing. Instead, failure in Experiment 3 would have to be due to the increased difficulty of the pattern induction learning phase. Because the only difference between the pattern induction learning phases in the two experiments concerned the presence of consistent voicing cues, failure in Experiment 3 would suggest that the voicing pattern assisted learners in Experiment 2.

Experiment 3

As in Experiment 2, infants were confronted with a list of bisyllabic words. Half of the infants heard words in which syllables began with /p/, /d/, and /k/ and ended with /b/, /t/, and /g/ (e.g., *dotkat*, referred to as Condition 1); the other half of the infants heard the opposite pattern (e.g., *taktod*, referred to as Condition 2). Unlike in Experiment 2, the only regularity in these pattern induction materials concerned the identity of segments; no generalizations about voicing could be drawn. After the pattern induction phase, infants were familiarized with the continuous speech stream from Experiment 2, which contained four new words, two consistent with each pattern induction condition. As in the previous experiments, infants were then tested on all four words heard during the segmentation phase. We hypothesized that if infants could learn the segment positions played during the pattern induction phase, they should show a difference in listening times for words fitting the familiar pattern and words fitting the novel pattern. However, if infants were assisted by the voicing regularities in Experiment 2, the lack of voicing regularities in these materials should make the pattern induction task more difficult, lowering the chances of successful test discrimination.

Method

Subjects. Thirty full-term 9-month-old monolingual infants with no history of recurrent ear infections were tested (mean age = 9 months 0 weeks; range = 8 months 3 weeks to 9 months 2 weeks). Half of the infants were randomly assigned to Condition 1, and half were assigned to Condition 2. Thirty-five additional infants were tested but not included in the analysis for the following reasons: fussiness (15), looking times averaging less than 3 s to one or both sides (6), parental interference (2), not looking at the sidelights (1), insufficient data for one trial type (1), and participation in two or more prior studies in the laboratory (10).

Stimuli. For the pattern induction phase, we used the two word lists from Experiment 2 (−V+V and +V−V) and altered them by changing /d/ to /t/ and /t/ to /d/. Condition 1 refers to the materials altered from the −V+V condition, and Condition 2 refers to the materials altered from the +V−V condition. Both lists contained the same set of phonemes (see Appendix C for word lists). The lists were synthesized and digitized as in Experiments 1 and 2, with each audiofile consisting of the 30-word list repeated twice. The pattern induction stimuli thus consisted of a digitized 60-word list conforming to either the Condition 1 template or the Condition 2 template and lasting approximately 2 min 11 s. The segmentation and test materials were those used in Experiment 2.

Procedure. The procedure was identical to those in Experiments 1 and 2.

Results and Discussion

The first analysis contrasted infants assigned to the two pattern induction conditions, Conditions 1 and 2. Because listening-time differences between familiar and novel words did not differ as a function of pattern induction condition, $t(28) = 1.29$, $p = .21$ (all tests are two-tailed), the data from the two conditions were pooled in the subsequent analyses.

The principal hypothesis concerned the difference in listening times between familiar and novel words—those that conformed to the phonological template heard during the pattern induction phase and those that did not. There was no significant difference in listening times for familiar versus novel words: $t(29) = 0.33$, $p = .74$. Unlike in Experiment 2, there was no consistent pattern of listening preferences; whereas 22 of the 30 infants in Experiment 2 listened longer to the novel words, only 13 of the 30 infants in Experiment 3 listened longer to the novel words. These results suggest that the pattern induction experiences of infants in Experiment 3 did not assist them in segmenting familiar patterns more readily than novel patterns.

To further compare the performance of infants in Experiments 2 and 3, we entered the data from the two experiments into an analysis of variance in which the factors were test item type (familiar vs. novel) and experiment (Experiment 2 vs. 3). Although the main effects of item type, $F(1, 58) = 2.77$, $p = .10$, and experiment, $F(1, 58) = 1.18$, $p = .28$, were not significant, the interaction between item type and experiment was significant, $F(1, 58) = 4.54$, $p < .05$. The difference in performance across experiments cannot be attributed to the difficulty of the segmentation and test phases themselves, because identical materials were differentially segmented and discriminated by the infants in Experiment 2. Instead, the results suggest that a minor change in the pattern induction materials—switching the positions of /d/ and /t/—had a major impact on infants' learning outcomes. Unlike the discriminations tested by Experiments 1 and 2, Experiment 3 did not provide evidence that infants can acquire lexical patterns based on the specific segments with which syllables can begin and end (see Figure 1 for the results from all three experiments).

General Discussion

Although a number of recent studies have demonstrated that infants learn a great deal about native-language phonological patterns during the 1st year, little research has addressed the nature of this learning process. The current study takes a first step in this direction, demonstrating rapid acquisition by infants of the types of phonological and phonotactic patterns that characterize human languages. In Experiment 1, a brief exposure to 60 new words containing a specific syllable structure affected infants' subsequent listening preferences, with different listening times found for novel words that conformed to the recently heard syllable structures and for novel words that did not. Infants may have been using their currently most highly activated phonological expectations to segment novel words from continuous speech despite the fact that all of the tested syllable structures were present and legal in their native language. Alternatively, the listening experience gained
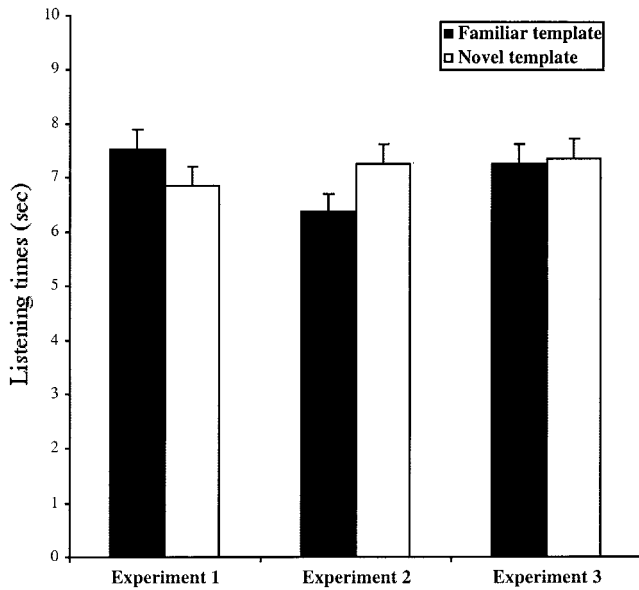
Figure 1. Listening times to test words consistent (familiar template) and inconsistent (novel template) with the phonological structures heard during the pattern induction phase of Experiments 1, 2, and 3.

during the pattern induction phase of the experiment may have served to prime infants' successful discrimination during the test, without any influence of the segmentation task. In either case, the successful test discrimination observed in this experiment required the infants to acquire the patterns presented during the pattern induction phase and to generalize them to include the new words presented during the segmentation and/or test phase.

In Experiments 2 and 3, we used this methodology to explore the nature of infants' phonological learning: Do infants attend primarily to generalizations in sound structure that obtain across the corpus, or do they attend to the privileges of occurrence of individual segments? Both types of information were available in Experiment 2: Infants could acquire patterns of voicing, or they could learn the positions of individual segments. Infants' pattern induction experiences affected their test performance, as measured by successful discrimination between words that fit the newly induced pattern and words that did not. However, no voicing patterns were available as cues for learners in Experiment 3; the positions of individual segments were the only useable patterns. In the absence of voicing regularities, infants failed to discriminate words consistent with the materials from the pattern induction phase of the experiment from other words despite the fact that the segmentation and test portions of the experiment were identical to those in Experiment 2.

The results of Experiments 2 and 3, taken together, suggest that inconsistent sound patterns are harder to learn than consistent sound patterns. Infants are able to capitalize on voicing regularities and, presumably, more generally on other types of phonological information available in native-language input. However, learning the regularities present in Experiment 3—that, for example, syllables could begin with /p/, /d/, or /k/—was more difficult for the infants. Although it is difficult to interpret null results, the fact that the segmentation and test stimuli were identical to the materials

used in Experiment 2 strongly suggests that something about the pattern induction materials in Experiment 3 impaired learning. Because the only difference between the two sets of materials concerned the presence of voicing generalizations, the difference between the experimental results is likely to be a function of the removal of consistent patterns useful to infant learners.

This hypothesis is particularly interesting when one considers the types of patterns found in natural languages. Regularities concerning voicing are prevalent cross-linguistically; for example, English and many other languages include voicing assimilation in their phonology and morphophonology (e.g., the plural morpheme /s/ undergoes voicing assimilation, as in /cats/ vs. /dogz/). However, the segmental regularities offered to infants in Experiment 3 (e.g., syllables can begin with /p/, /d/, and /k/ but not /b/, /t/, and /g/) are highly unlikely to occur in languages. To follow this type of regularity, a language would have to allow /pat/ and /dat/ as possible words but not /gat/. Phonological systems tend not to have restrictions on individual segments. Instead, phonological systems place restrictions on *types* of segments. Although it is certainly the case that infants must attend to the identities of individual segments in order to learn words—that is, to notice that /pat/ and /kat/ are different lexical items—the acquisition of phonological systems requires generalization from attested patterns to expected patterns. To the extent that patterns that do not occur in natural languages are more difficult to acquire, we may consider the possibility that constraints on how infants learn may have served to shape the phonology of natural languages. Patterns that are difficult to acquire are less likely to persist cross-linguistically than those that are easily learned. Thus, languages may exploit devices such as voicing regularities in part because they are readily acquired by young learners.

The current experiments represent an early stage in understanding phonological learning during the 1st year. Infants were exposed only to a single generalization at a time, whereas in natural languages, infants are exposed to numerous generalizations. Similarly, the experimental materials were deterministic rather than probabilistic: Regularities never occur 100% of the time in real language input as they did in these experiments. In addition, prior to participating in our experiments, these infants spent 9 months beginning to learn the sound patterns of English, which may have influenced the types of information they prioritized when inducing patterns in the laboratory. Nevertheless, the results suggest interesting avenues for future research. For example, if infants are truly generalizing the phonological patterns in the input, they should show the same pattern of performance when given totally novel materials; for example, after hearing voicing regularities in Experiment 2, infants should extend their newly acquired phonological knowledge to include phonemes not heard earlier in the experimental session.

The infants tested in this experiment were 9 months of age, a point during development at which infants are gathering detailed knowledge of the phonological systems that characterize their native language (for a review, see Jusczyk, 1997). We might therefore expect younger infants not to show the same level of specificity in the types of phonological patterns they are willing to construct—or, more likely, younger infants may not acquire phonological generalizations at all. Older infants, however, raise conflicting predictions. Increased native-language knowledge might predict increased specificity of the patterns readily acquired with

age. It is also possible that the increased information-processing capacity that emerges over development might lead older infants to more readily acquire arbitrary combinations of sounds with age, as suggested by results from Chambers, Onishi, and Fisher (2002). Future research is needed to disentangle these divergent developmental predictions and to provide a better understanding of the relationship between the acquisition of a more nuanced linguistic system and increased general aptitude for learning.

With the current results in hand, we can begin to explore other constraints on learning present during infancy. To the extent that infants are not open-minded in the types of generalizations they rapidly acquire, we may find explanations both for the course of language development and for cross-linguistic regularities. Languages may not contain patterns like those found in the materials used in Experiment 3 because infants find these regularities relatively difficult to detect. Similarly, certain patterns may not occur in word formation (e.g., Newport & Aslin, 2000) and syntax (e.g., Gomez, 2002; Saffran, 2002) because of the difficulties they pose for human learners. Exploring the types of patterns infants *cannot* learn may be just as informative as documenting their impressive feats of learning.

References

Aslin, R. N. (2000, July). *Interpretation of infant listening times using the headturn preference technique.* Talk presented at the International Conference on Infancy Studies, Brighton, England.

Bever, T. (1970). The cognitive basis for linguistic structures. In J. Hayes (Ed.), *Cognition and the development of language* (pp. 279–352). New York: Wiley.

Chambers, K. E., Onishi, K. H., & Fisher, C. (2002, April). *Phonotactic acquisition in infancy.* Poster presented at the International Conference on Infant Studies, Toronto, Ontario, Canada.

Christiansen, M. H., & Devlin, J. T. (1997). Recursive inconsistencies are hard to learn: A connectionist perspective on universal word order correlations. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 113–118). Hillsdale, NJ: Erlbaum.

Cole, R. A. (1981). Perception of fluent speech by children and adults. *Annals of the New York Academy of Sciences*, 379, 92–109.

Crystal, T. H., & House, A. S. (1988). Segmental durations in connected-speech signals: Syllabic stress. *Journal of the Acoustical Society of America*, 83, 1574–1585.

Cutting, J. E. (1974). Two left-hemisphere mechanisms in speech perception. *Perception & Psychophysics*, 16, 601–612.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1355–1367.

DiSimoni, F. G. (1974). Influence of consonant environment on duration of vowels in the speech of three-, six-, and nine-year-old children. *Journal of the Acoustical Society of America*, 55, 362–363.

Eilers, R. E. (1977). Context-sensitive perception of naturally produced stops and fricative consonants by infants. *Journal of the Acoustical Society of America*, 61, 1321–1336.

Eilers, R. E., Bull, D. H., Oller, K., & Lewis, D. L. (1984). The discrimination of vowel duration by infants. *Journal of the Acoustical Society of America*, 75, 1213–1218.

Eilers, R. E., Wilson, W. R., & Moore, J. M. (1977). Developmental changes in speech discrimination in infants. *Journal of Speech and Hearing Research, 18,* 158–167.

Eimas, P. D., Siqueland, E. R., Jusczyk, P. W., & Vigorito, J. (1971). Speech perception in infants. *Science, 171,* 303–306.

Ellefson, M. R., & Christiansen, M. H. (2000). Subjacency constraints without universal grammar: Evidence from artificial language learning and connectionist modeling. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 645–650). Hillsdale, NJ: Erlbaum.

Friederici, A. D., & Wessels, J. M. I. (1993). Phonotactic knowledge and its use in infant speech perception. *Perception & Psychophysics, 54,* 287–295.

Gomez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science, 13,* 431–436.

Gomez, R. L., & Gerken, L. A. (1999). Artificial grammar learning by one-year-olds leads to specific and abstract knowledge. *Cognition, 70,* 109–135.

Gomez , R. L., & Gerken, L. A. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences*, 4, 178–186.

Goodsitt, J. V., Morgan, J. L., & Kuhl, P. K. (1993). Perceptual strategies in prelingual speech segmentation. *Journal of Child Language, 20,* 229–252.

Gupta, P., & Dell, G. S. (1999). The emergence of language from serial order and procedural memory. In B. MacWhinney (Ed.), *Emergentist approaches to language, 28th Carnegie Mellon Symposium on Cognition* (pp. 447–481). Mahwah, NJ: Erlbaum.

Hillenbrand, J. (1984). Speech perception by infants: Categorization based on nasal consonant place of articulation. *Journal of the Acoustical Society of America*, 75, 1613–1622.

Hogan, J. T., & Rozyspal, A. J. (1980). Evaluation of vowel duration as a cue for the voicing distinction in the following word-final consonant. *Journal of the Acoustical Society of America*, 67, 1764–1771.

House, A. S. (1961). On vowel duration in English. *Journal of the Acoustical Society of America, 33,* 1174–1178.

Hunter, M. A., & Ames, E. W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. *Advances in Infancy Research, 5,* 69–95.

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. (1997). *The discovery of spoken language*. Cambridge, MA: MIT Press.

Jusczyk, P. W., Cutler, A., & Redanz, L. (1993). Infants' sensitivity to predominant stress patterns in English. *Child Development, 64,* 675–687.

Jusczyk, P. W., Friederici, A. D., Wessels, J. M., Svenkerud, V. Y., & Jusczyk, A. M. (1993). Infants' sensitivity to the sound patterns of native language words. *Journal of Memory and Language*, 32, 402–420.

Jusczyk, P. W., Goodman, M. B., & Baumann, A. (1999). 9-month-olds' attention to sound similarities in syllables. *Journal of Memory and Language*, 40, 62–82.

Jusczyk, P. W., Houston, D. M., & Newsome, M. (1999). 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.

Kenstowicz, M. (1994). *Phonology in generative grammar.* Cambridge, MA: Blackwell.

Krause, S. E. (1982). Vowel duration as a perceptual cue to postvocalic consonant voicing in young children and adults. *Journal of the Acoustical Society of America, 71,* 990–995.

Kuhl, P. K. (1979). Speech perception in early infancy: Perceptual constancy for spectrally dissimilar vowel categories. *Journal of the Acoustical Society of America*, 66, 1668–1679.

Kuhl, P. K. (1983). Perception of auditory equivalence classes for speech in early infancy. *Infant Behavior and Development, 6,* 263–285.

Kuhl, P. K., Williams, K. A., Lacerda, F., Stevens, K. N., & Lindblom, B.

(1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science, 255,* 606–608.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science, 283,* 77–80.

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 & prosodic effects on word segmentation in infants. *Cognitive Psychology, 38,* 465–494.

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition, 82,* B101-B111.

Mehler, J., Jusczyk, P. W., Lambertz, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. *Cognition, 29,* 144–178.

Morgan, J. L., & Saffran, J. R. (1995). Emerging integration of segmental and suprasegmental information in prelingual speech segmentation. *Child Development, 66,* 911–936.

Newport, E. L. (1982). Task specificity in language learning? Evidence from speech perception and American Sign Language. In E. Wanner & L. R. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 450–486). Cambridge, England: Cambridge University Press.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science, 14,* 11–28.

Newport, E. L., & Aslin, R. N. (2000). Innately constrained learning: Blending old and new approaches to language acquisition. In *Proceedings of the 24th Annual Boston University Conference on Language Development* (pp. 1–21). Somerville, MA: Cascadilla Press.

Onishi, K. H., Chambers, K. E., & Fisher, C. (2002). Learning phonotactic constraints from brief auditory experience. *Cognition, 83,* B13-B23.

Peterson, G. E., & Lehiste, I. (1960). Duration of syllable nuclei in English. *Journal of the Acoustical Society of America, 32,* 693–703.

Raphael, L. J. (1972). Preceding vowel duration as a cue to the perception of the voicing characteristic of word-final consonants in English. *Journal of the Acoustical Society of America, 51,* 1296–1303.

Roder, B. J., Bushnell, E. W., & Sasseville, A. M. (2000). Infants' preference for familiarity and novelty during the course of visual processing. *Infancy, 1,* 491–507.

Saffran, J. R. (2002). Constraints on statistical language learning. *Journal of Memory and Language, 47,* 172–196.

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

Thiessen, E. D., & Saffran, J. R. (2001, April). *Infant use of stress cues versus statistical cues to word segmentation at 7, 9, and 11 months.* Poster presented at the Biennial Meeting of the Society for Research in Child Development, Minneapolis, MN.

Thiessen, E. D., & Saffran, J. R. (in press). When cues collide: Use of stress and statistical cues to word boundaries by 7- to 9-month-old infants. *Developmental Psychology.*

Werker, J., & Tees, R. C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development, 7,* 49–63.

Appendix A

Experiment 1 Stimuli

CVCV condition			CVCCVC condition		
boga	rati	dibo	tilgud	pudrat	ludkab
giku	gado	tila	pikbak	lokdip	rotgul
koga	kuga	dari	lirkag	galtup	gikrub
pabu	bugo	bidu	dubtit	buldar	kobkig
tala	bagi	paba	rukpig	bikrub	takbid
piro	toda	tori	ditbul	togbut	tidgar
podu	pula	lapu	rabgag	kadlok	dokbab
kabi	rapi	balu	daklit	larbid	gadkug
lido	lita	roko	birkig	bagrit	tikdup
tuka	gapi	dibu	takbut	dagkot	bapgut

Note. Pronunciation: /i/ = "ee"; /u/ = "oo"; /a/ = "ah"; and /o/ = "oh".

(*Appendixes continue*)

Appendix B

Experiment 2 Stimuli

−V + V condition			+V − V condition		
todkad	tigpod	kigtod	dakdot	dopgit	dotgik
kigpid	kogpog	pobtag	dipgik	gopguk	gatbop
pibtog	tibpid	kogpub	gotbip	dipbit	bupgok
kobtig	pabkib	podtib	gitbok	bikbap	bitdop
pidkad	kagpud	pabtug	gakdip	dupgak	gutbap
tudpig	pagtud	tadkud	gipdut	dutgap	dukdat
pudkad	podkug	kubkag	dakdup	dukbop	gakbuk
tabtod	tidkob	tidpog	dotbat	bokdit	gopdit
tobkid	kabtag	tudkob	dikbot	gatbak	bokdut
kubtib	togpud	padpug	bitbuk	dupgot	gupdap

Note. Pronunciation: /i/ = "ee"; /u/ = "oo"; /a/ = "ah"; and /o/ = "oh".

Appendix C

Experiment 3 Stimuli

Condition 1			Condition 2		
dotkat	digpot	kigdot	taktod	topgid	todgik
kigpit	kogpog	pobdag	tipgik	gopguk	gadbop
pibdog	dibpit	kogpub	godbip	tipbid	bupgok
kobdig	pabkib	potdib	gidbok	bikbap	bidtop
pitkat	kagput	pabdug	gaktip	tupgak	gudbap
dutpig	pagdut	datkut	giptud	tudgap	tuktad
putkat	potkug	kubkag	taktup	tukbop	gakbuk
dabdot	ditkob	ditpog	todbad	boktid	goptid
dobkit	kabdag	dutkob	tikbod	gadbak	boktud
kubdib	dogput	patpug	bidbuk	tupgod	guptap

Note. Pronunciation: /i/ = "ee"; /u/ = "oo"; /a/ = "ah"; and /o/ = "oh".