Lexical support for phonetic perception during nonnative spoken word recognition

Arthur G. Samuel^{1,2,3} · Ram Frost^{3,4}

Published online: 23 October 2015
© Psychonomic Society, Inc. 2015

Abstract Second language comprehension is generally not as efficient and effective as native language comprehension. In the present study, we tested the hypothesis that lower-level processes such as lexical support for phonetic perception are a contributing factor to these differences. For native listeners, it has been shown that the perception of ambiguous acoustic–phonetic segments is driven by lexical factors (Samuel *Psychological Science*, 12, 348–351, 2001). Here, we tested whether nonnative listeners can use lexical context in the same way. Native Hebrew speakers living in Israel were tested with American English stimuli. When subtle acoustic cues in the stimuli worked against the lexical context, these nonnative speakers showed no evidence of lexical guidance of phonetic perception. This result conflicts with the performance of native speakers, who demonstrate lexical effects on phonetic perception even with conflicting acoustic cues. When stimuli without any conflicting cues were used, the native Hebrew subjects produced results similar to those of native English speakers, showing lexical support for phonetic perception in their second language. In contrast, native Arabic speakers, who were less proficient in English than the native Hebrew speakers, showed no ability to use lexical activation to support phonetic perception, even without any conflicting cues. These

results reinforce previous demonstrations of lexical support of phonetic perception and demonstrate how proficiency modulates the use of lexical information in driving phonetic perception.

Keywords Top-down lexical support · Spoken word recognition · Nonnative listening

Experience using a nonnative language (L2) leads to greater proficiency, but certain aspects seem to be difficult to master. For example, achieving unaccented speech is notoriously difficult for anyone who does not begin learning an L2 at a young age. Similarly, certain syntactic forms, such as the use of articles and prepositions or gender marking, are also often imperfect, even in relatively proficient bilinguals. Here, we investigated whether these limitations extend down to the fundamental, low-level processes that underlie the perception of speech sounds. Specifically, we tested the extent to which nonnative listeners can use lexical information to support phonetic perception. There is substantial evidence that in native language listening, such support plays an important role.

Two classic effects illustrate lexical support for phonetic encoding: the Ganong effect and phonemic restoration. Ganong (1980) created stimuli with an ambiguous initial phoneme (e.g., a sound midway between /d/ and /t/), followed by one of two syllable endings (e.g., “_ask” and “_ash”). When an ambiguous /d/–/t/ sound preceded “ask,” people tended to report hearing “task,” whereas when it preceded “ash,” they more often reported “dash.” Thus, how the ambiguous sound is heard depends on the lexical context. Warren (1970) showed that when part of a word is replaced by noise, listeners usually hear the word as intact (“phonemic restoration”). Restoration is stronger in real words than in matched

✉ Arthur G. Samuel
a.samuel@bcbl.eu

¹ Department of Psychology, Stony Brook University, Stony Brook, NY, USA

² IKERBASQUE, Basque Foundation for Science, Bilbao, Spain

³ Basque Center on Cognition Brain and Language, Donostia-San Sebastian, Spain

⁴ Department of Psychology, Hebrew University, Jerusalem, Israel

pseudowords, supporting a lexical influence on the perceived phonemes (Samuel, 1981).

A longstanding controversy concerns whether these effects provide unambiguous proof that lexical representations affect the activation of phonemic representations. Interactive models (e.g., McClelland & Elman, 1986) posit such lexical activation of phonemic representations, but autonomous models (e.g., Norris, McQueen, & Cutler, 2000) only allow for information flow from lower-level to higher-level representations. In autonomous models, the Ganong effect and phonemic restoration are attributed to a postperceptual decision stage, rather than to feedback from the lexical down to the phonemic level. It has proven very difficult to empirically distinguish between these views, but one type of test provides a clear separation: If a lexical influence on phonemic perception is found using a measure that does not require overt identification, interactive models are supported over autonomous ones, because the latter focus on a decision-stage effect that is precluded when such decisions are not made.

Samuel (1997, 2001) provided two such tests. In one, the lexical items supported phonetic perception through phonemic restoration; in the second, the Ganong effect provided the lexical support. Both studies showed that in native language listening, lexical context can determine phonetic perception without involving overt identification. In the present study, we tested whether nonnative listening also has this property, or whether instead such processes are among those that are difficult to acquire after the native language is established. Presumably, aspects that are difficult to implement in the L2 are those that need to be fine-tuned and secured during initial acquisition, to meet challenging processing requirements.

As in Samuel (2001), our approach involved the combination of two speech phenomena: the Ganong effect and selective adaptation (Eimas & Corbit, 1973). To ensure that responses were not based on postperceptual decision processes, our methodology focused on lexical influences on phonetic perception that are revealed by a *consequence* of perception, rather than by a decision regarding phonemic identity. This was done using the selective-adaptation paradigm. In adaptation studies, listeners identify items from a continuum of speech sounds. For example, different mixtures of “iss” and “ish” can be constructed, with the most extreme “iss” being heavily weighted toward “iss,” and the most extreme “ish” being heavily weighted toward “ish.” These mixtures are played in a random order, many times, and listeners identify each token as “iss” or as “ish.” Eimas and Corbit (1973) showed that if a listener then hears a sound repeatedly, the person’s phoneme boundary shifts. For example, if “iss” is heard repeatedly, there will be fewer “iss” reports than on an initial baseline.

Samuel (2001) tested whether lexically driven (Ganong) phonetic percepts can produce adaptation. Listeners identified

“iss”–“ish” syllables in two adaptation conditions. In one, the repeated items were words like “diminish”—words that ended in “sh.” In the other, the repeated words were ones that ended in “s” (e.g., “bronchitis”). However, the final segment in each word was actually an ambiguous mixture of “s” and “sh” and was identical across the two conditions. Because of the Ganong effect, listeners generally heard this sound as “s” in one condition, and as “sh” in the other. Critically, listeners never made any judgments about the repeating words—they just listened. The measurement was on the “iss”–“ish” test syllables presented *afterward*, a consequential effect. The ambiguous mixtures should reduce later “s” reports if that portion of the waveform is *perceived* as “s,” and should reduce “sh” reports if the lexicon pushes perception toward “sh.” In multiple experiments, the predicted adaptation shifts were found. Comparable findings were obtained using lexically driven phonemic restoration (Samuel, 1997).

Here, we tested whether nonnative listeners rely on lexical support for phonetic perception, as native listeners do. Using the English stimuli that had successfully been used before, we tested Hebrew speakers in Israel. Hebrew includes consonants that are quite similar to English “s” and “sh,” so that Hebrew speakers are familiar with the “s”–“sh” contrast. Children in Israel learn English in school, usually starting around age 10.

Our stimuli were designed with a fundamental property of spoken language in mind: Each vowel or consonant blended into the ones before and after it, a phenomenon known as *coarticulation*. Because of coarticulation, when an experimenter replaces a speech segment with another segment, decisions must be made about where to cut into any adjacent segments. For example, if the “sh” in “diminish” is to be replaced by an ambiguous “s”–“sh” mixture, some of the preceding vowel must also be replaced. Stimulus construction thus involves trade-offs between eliminating residual coarticulatory cues in the vowel and stimulus naturalness.

In Experiment 1, we controlled for coarticulation by using adaptors made from words that were intentionally mispronounced. For example, “diminiss” and “bronchitish” were recorded, rather than “diminish” or “bronchitis.” When the final fricative was replaced by an ambiguous mixture of “s” and “sh,” any residual coarticulatory cues would produce shifts in the direction opposite those that lexical support would produce. For example, any residual cues from “diminiss” would cause adaptation consistent with the “s” that was originally recorded, but the lexical influence would support the “sh” in “diminish.” Samuel (2001) found that these adaptors produced smaller shifts than the original stimuli, but that the shifts were still significant and in the lexically based direction; there were some small residual cues, but the effect was primarily driven by lexical activation.

Experiment 1

In Experiment 1, we tested whether nonnative listeners demonstrate lexical support of phonetic perception. If they do, we can conclude that knowledge of a language comes with lexical support for the perceptual encoding of sublexical information. If nonnative listeners do not produce native-like adaptation shifts, then lexical support would best be viewed as a property (perhaps like the ability to produce nonaccented speech) that comes with the initial formation of the language-processing system.

Method

Subjects Thirty native Hebrew speakers were recruited. We obtained both an objective measurement of English vocabulary and a set of self-report measures of English use. Gollan et al.'s (2012) Multilingual Naming Test (MINT) was used to assess English vocabulary knowledge. The MINT consists of a set of 68 black-and-white line drawings that subjects are asked to name, starting with very familiar objects and progressively moving to more obscure ones. As the context for assessing the English level of our bilinguals, Gollan et al. reported correct report of about 55 items by Spanish-dominant bilinguals living in California (and about 65 items by monolingual English speakers). The subjects also completed the Language Experience and Proficiency Questionnaire (LEAP-Q; Marian, Blumenfeld, & Kaushanskaya, 2007), providing information about what languages a person knows, when they were learned, frequency of exposure to each, and self-assessments of reading, listening, and speaking proficiency.

Table 1 provides the average numbers of words successfully named in English on the MINT, and the most relevant LEAP-Q information. The English exposure measure provides a subject's percentage of daily English language use. The average is rather low because our bilinguals live in a country in which English is not the dominant language. Age of acquisition (AoA) is the subject's report of the age at which they started to learn English, and the three remaining scores are self-assessments, on a 10-point scale, of speaking, understanding, and reading ability. MINT scores were positively correlated with self-reported English-speaking ability ($r = 0.48$, $p = .02$).

Stimuli The stimuli were digital copies of those used by Samuel (2001, Exp.3). Two types of stimuli were used—test items on a continuum from “iss” to “ish,” and words that served as the adaptors. All stimuli were made by digitally editing recordings (16-kHz sample rate) made by a monolingual male speaker of American English.

An eight-step test continuum was made by digitally mixing varying weightings of the frication from “s” with the frication

from “sh,” in 5% increments (e.g., 35% “s” with 65% “sh,” 40% “s” with 60% “sh,” etc.). The adaptors were based on eight words, four ending in “s” (“bronchitis,” “embarrass,” “malpractice,” and “tremendous”) and four ending in “sh” (“abolish,” “demolish,” “diminish,” and “replenish”). As we noted above, the original recordings were mispronounced versions of each of these words, with “s” instead of “sh” and vice versa. Eight versions of each adaptor were constructed, replacing the final fricative with “s”–“sh” mixtures in 5% increments. A portion of the vowel preceding the final fricative was replaced by white noise in order to reduce the conflicting coarticulatory information, but the preceding vowel (/I/) was easily perceived from its initial pitch periods. Each subject only heard a single mixture (e.g., the one based on 45% “s” and 55% “sh”), with that same mixture being used for all eight adaptors (see below).

Apparatus and procedure Subjects came to the laboratory twice, each time doing an identification task followed by an adaptation task. The identification task was used to familiarize the subject with the stimuli and to estimate each subject's perceptual boundary between “s” and “sh.” The eight members of the “iss”–“ish” continuum were randomized 16 times and presented to listeners over high-quality headphones. Subjects identified each token as ending in either “s” or “sh” by pushing response buttons. One second after a response, the next syllable was presented, with a 2,500-ms timeout.

After the identification task during a subject's first session, the experimenter examined the percentage of “sh” responses to each of the eight members of the continuum. The experimenter identified that subject's boundary between “s” and “sh”—the token that received a response rate closest to 50% “sh” (and thus, 50% “s”)—in order to individually select the tokens for the adaptation task. The idea was to pick the mixture of “s” and “sh,” for each subject, that was most ambiguous, and thus most subject to being influenced by the lexical context. Critically, whatever mixture was selected for a given subject was the mixture used for *all* adapting words, in both sessions—the boundary was established once and was not changed.

During the second session, the identification task was only used to give the subject a warm-up period. In both sessions, the identification task was followed by an adaptation task that included 14 passes, with each pass consisting of an adaptation phase followed by a response phase. During one session, the adaptation task was based on the four words that ended in “s,” and during the other, it was based on the “sh” words. Order was counterbalanced across subjects. An adaptation phase consisted of ten randomizations of the four words, with 300 ms of silence between successive adaptor words. Thus, subjects heard 40 tokens during each adaptation phase over the course of approximately 40 s. Subjects simply listened during adaptation—no responses were made. After a 1-s

Table 1 Multilingual Naming Test (MINT) scores and self-report values from the Language Experience and Proficiency Questionnaire, for the native Hebrew-speaking subjects

L1	MINT	English exposure (%)	AoA	English speaking	English understanding	English reading
Hebrew	49.4	18.3	6.3	7.8	8.7	7.7

pause at the end of an adaptation phase, a response phase began. This phase included one randomization of the eight “iss”–“ish” stimuli. Subjects identified each token as ending with “s” or “sh.” One second after the eighth response, the next adaptation phase began.

Results and discussion

On the initial “iss”–“ish” labeling task, seven subjects failed to identify members of the test series consistently; they were not included in the data analyses. For the remaining subjects, two response functions were computed, one based on responses during the “s” adaptation session, and one from the “sh” adaptation session. Each function traced the increase in reports of “sh” as the tokens in the test continuum changed from “s” to “sh.” If listeners used the lexical context to encode the ambiguous phonetic information (despite potentially misleading coarticulatory cues from the mispronounced words), then the “sh” adaptation function should lie below the one from the “s” condition. Figure 1 (left panel) presents these functions for the native Hebrew speakers; for comparison, the results for native English speakers (from Samuel, 2001) are shown (middle panel).

There is a weak trend in the lexically driven direction for the Hebrew speakers, but the difference is quite small. For each subject, the average report of “sh” across the middle four members of the test continuum was computed, for both the “sh” adaptation condition and the “s” adaptation condition. The average difference in “sh” reports as a function of adaptor condition was

then tested with within-subjects paired *t* tests. Adaptation effects are typically largest near the category boundary, making this metric a sensitive one (Samuel, 1997, 2001). As Fig. 1 shows, there was no significant shift for the native Hebrew speakers [$t(22) = 1.165$, n.s.].

This null effect suggests that native language proficiency is needed for lexical activation to intervene in the encoding of phonetic information. However, before accepting this conclusion, we should be certain that native speakers can in fact use lexical information to support phonetic encoding when faced with input containing potentially misleading coarticulatory cues. Samuel (2001) reported evidence consistent with this conclusion, but there was only a single test under these conditions, and the effect was significant but small. Therefore, we ran a replication of that experiment with American listeners. The stimuli were the same ones used in the present experiment. A group of 35 native speakers of American English were tested at Stony Brook University. Seven failed to label the “iss”–“ish” test items consistently and were not included in the analyses. As Fig. 1 (right panel) shows, the new results are extremely similar to the original ones (middle panel)—a small but significant shift [$t(27) = 1.703$, $p = .05$], despite the potential influence of residual cues that would push identification in the opposite direction. Thus, our findings demonstrate a difference in phonetic perception between native and nonnative speakers: Whereas native listeners use lexical context to aid phonetic encoding, nonnative listeners do not have the capacity to do so, at least not with residual acoustic–phonetic information that opposes the lexical influence.

Fig. 1 Identification functions after adaptation with words that end in either “s” (open triangles) or “sh” (filled triangles), for adaptors that were originally intentionally mispronounced in order to eliminate any potentially helpful coarticulatory cues. *Left panel* Results for native

Hebrew speakers. *Middle panel* Results from the native English-speaking subjects in Samuel (2001, Exp.3). *Right panel* Results from the native English-speaking subjects in the replication experiment of the present study

Experiment 2

The stimuli in Experiment 1 did not merely deprive listeners of potentially helpful coarticulatory information. Rather, they included potentially *conflicting* coarticulatory cues: Words that should end in “s” were recorded with “sh,” and vice versa. This method provides the most rigorous test of lexical influences, but it might underestimate listeners’ use of lexical support. In Experiment 2, we tested whether nonnative listeners can use lexical support to drive phonetic encoding in the absence of any conflicting coarticulatory information.

Method

Subjects Twenty new subjects were recruited from the same population tested in Experiment 1.

Stimuli The stimuli were digital copies of those used in Experiment 1 of Samuel (2001). The same eight-step “iss”–“ish” continuum was used, and the same set of eight words served as adaptors. However, normal versions of the words were used to construct the adaptors, rather than versions with mispronunciations. See Samuel (2001) for details.

Apparatus and procedure The same apparatus and procedures were used as in Experiment 1.

Results and discussion

On the “iss”–“ish” identification task, two speakers failed to identify members of the test series consistently; they were not included in the analyses. Figure 2 (left panel) presents the labeling functions for the native Hebrew speakers; for comparison, the results for native English speakers (Samuel, 2001) are shown (right panel).

Fig. 2 Identification functions after adaptation with words that end in either “s” (open triangles) or “sh” (filled triangles), for adaptors that were originally pronounced normally. *Left panel* Results from native Hebrew-

The findings of Experiment 2 are straightforward: Adaptation from words without conflicting residual cues produced a robust effect for the native Hebrew speakers [$t(17) = 2.86, p < .01$], quite similar to the pattern for the native English speakers. The strong effect for the native Hebrew speakers demonstrates that even nonnative listeners can use lexical context to drive phonetic perception.

General discussion

The present study provides new evidence that lexical context affects the *perception* of phonetic input, and adds two critical new findings. First, even nonnative speakers are capable of using lexical context to support phonetic encoding. Second, although present, this ability appears to be weaker in nonnative language processing, since only the native speakers employed it in the face of conflicting acoustic information.

Taken together, the data of Experiments 1 and 2 suggest that the ability to use lexical activation to support phonetic encoding is correlated with proficiency—more proficient (native) subjects are able to do so under a wider set of conditions. To explore this notion, we selected a new population of nonnative English speakers—L1 Arabic speakers. These were individuals of the same age, living in the same location, but with a different language history. In particular, L1 Arabic speakers who know English are generally significantly less proficient than comparable L1 Hebrew speakers; the former typically must learn Modern Standard Arabic (which differs from spoken Arabic), Hebrew, and English, providing less experience with English. We recruited 25 L1 Arabic speakers and tested them with the stimuli and procedures of Experiment 1; Table 2 presents their scores on the MINT and the LEAP-Q. Importantly, as desired, the MINT scores for these L1 Arabic subjects were significantly lower than the scores for the L1 Hebrew group, $F(1, 38) = 10.45, p < .005$. To provide a more complete comparison to the more proficient

speaking subjects. *Right panel* Results from the native English-speaking subjects in Samuel (2001, Exp. 1)

Table 2 Multilingual Naming Test (MINT) scores and self-report values from the Language Experience and Proficiency Questionnaire, for the native Arabic-speaking subjects

L1	MINT	English exposure (%)	AoA	English speaking	English understanding	English reading
Arabic	39.5	13.8	7.2	7.5	8.5	8.8

L1 Hebrew listeners, we recruited a second group of L1 Arabic speakers ($N = 23$) from the same population to test with the stimuli and procedures of Experiment 2. If a high level of nonnative proficiency is needed to capitalize on lexical support for phonetic perception, then not only should these listeners fail to do so when there are conflicting acoustic cues, they should also be less able to do so with normal coarticulatory cues.

Eight listeners in the experiment with conflicting cues, and four listeners in the experiment with normal cues, failed to identify the test items consistently and were not included in the analyses. The results for both groups, shown in Fig. 3, are very clear: With both conflicting [$t(16) = 0.387$, n.s.] and normal [$t(18) = 0.15$, n.s.] coarticulation, these less-proficient nonnative listeners showed no evidence of being able to use lexical support in phonetic encoding. Thus, we see a very systematic pattern: Native speakers enjoy lexical support of phonetic encoding even with problematic input. Relatively proficient nonnative speakers receive such support only when the acoustics can contribute (or at least not conflict) with coarticulatory cues. Finally, less-proficient nonnative speakers cannot rely on the lexicon for such perceptual assistance. The results of the present study do not allow us to say whether the native Hebrew speakers needed the subtle coarticulatory cues that were available in Experiment 2 or whether instead, being nonnative English speakers, their ability to use lexical context was blocked in Experiment 1 by the presence of conflicting coarticulatory cues. This requires further investigation.

Another unknown is what aspect of proficiency is critical in the ability of lexical representations to support phonetic encoding. One aspect of proficiency is the size of the lexicon, and it is possible that more of the L1 Hebrew population than the L1 Arabic population knows the relatively low-frequency adaptor words. It is plausible that a few of them (perhaps “diminish,” “replenish,” and “malpractice”) might have been unfamiliar to some of the less-proficient listeners. Although this is possible, it seems unlikely to account for the results. Most of the words are ones that this college-educated population would be likely to have come across, and even if a few were unfamiliar, the others should have been effective (any unfamiliar ones should have been neutral in the adaptation procedure). In addition, the L1 Hebrew subjects across the two experiments produced different adaptation patterns as a function of the coarticulatory details, with the same adapting words.

Our results suggest that increasing L2 proficiency leads to the development of fully functional lexical representations, which in turn lead to changes in low-level phonetic processing. In Kroll’s revised hierarchical model (Kroll, Van Hell, Tokowicz, & Green, 2010), less-proficient bilinguals rely on L1 mediation for mapping a word onto its meaning; more-proficient operation of L2 allows for a direct mapping from the input word to its conceptual representation. Our results can be viewed in a similar way, but with lexical development having “downward” effects on the ability to engage phonetic encoding, rather than “upward” effects to conceptual representations. In both cases, the notion is that a less-proficient L2

Fig. 3 Identification functions after adaptation with words that end in either “s” (open triangles) or “sh” (filled triangles), for native Arabic speakers. *Left panel* Results for adaptors that were originally

intentionally mispronounced in order to eliminate any potentially helpful coarticulatory cues. *Right panel* Results for adaptors that were originally pronounced normally

speaker gets less support in word recognition from lexical representations than does a more-proficient speaker. From this perspective, a critical aspect of increased proficiency in a non-native language is the development of lexical representations that are fully functional—that is, able to engage with representations at lower (e.g., phonetic) levels, at the same (lexical) level, and at higher (e.g., conceptual) levels. This suggestion resonates with parallel claims regarding lexical quality in the domain of orthographic representations (e.g., Perfetti, 2007). The growing literature on the development of lexical representations (e.g., Dumay & Gaskell, 2007; Gaskell & Dumay, 2003; Leach & Samuel, 2007) has shown that this kind of lexical engagement is separable from simply accumulating knowledge about a word. Our results suggest that a critical aspect of language proficiency is the development of such fully functional lexical representations.

References

- Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science*, 18, 35–39. doi:10.1111/j.1467-9280.2007.01845.x
- Eimas, P. D., & Corbit, J. D. (1973). Selective adaptation of linguistic feature detectors. *Cognitive Psychology*, 4, 99–109.
- Ganong, W. F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 110–125. doi:10.1037/0096-1523.6.1.110
- Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, 89, 105–132. doi:10.1016/S0010-0277(03)00070-2
- Gollan, T., Weissberger, G. H., Runnqvist, E., Montoya, R. I., & Cera, C. M. (2012). Self-ratings of spoken language dominance: A Multilingual Naming Test (MINT) and preliminary norms for young and aging Spanish-English bilinguals. *Bilingualism: Language and Cognition*, 15, 594–615.
- Kroll, J. F., Van Hell, J. G., Tokowicz, N., & Green, D. W. (2010). The revised hierarchical model: A critical review and assessment. *Bilingualism: Language and Cognition*, 13, 373–381.
- Leach, L., & Samuel, A. G. (2007). Lexical configuration and lexical engagement: when adults learn new words. *Cognitive Psychology*, 55, 306–353. doi:10.1016/j.cogpsych.2007.01.001
- Marian, V., Blumenfeld, H. K., & Kaushanskaya, M. (2007). The Language Experience and Proficiency Questionnaire (LEAP-Q): Assessing language profiles in bilinguals and multilinguals. *Journal of Speech, Language, and Hearing Research*, 50, 940–967.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1–86. doi:10.1016/0010-0285(86)90015-0
- Norris, D., McQueen, J. M., & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary. *Behavioral and Brain Sciences*, 23, 299–325. doi:10.1017/S0140525X00003241
- Perfetti, C. (2007). Reading ability: Lexical quality to comprehension. *Scientific Studies of Reading*, 11, 357–383.
- Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, 110, 474–494. doi:10.1037/0096-3445.110.4.474
- Samuel, A. G. (1997). Lexical activation produces potent phonemic percepts. *Cognitive Psychology*, 32, 97–127.
- Samuel, A. G. (2001). Knowing a word affects the fundamental perception of the sounds within it. *Psychological Science*, 12, 348–351. doi:10.1111/1467-9280.00364
- Warren, R. M. (1970). Perceptual restoration of missing speech sounds. *Science*, 167, 392–393.