J. Linguistics **41** (2005), 353–387. © 2005 Cambridge University Press
doi:10.1017/S0022226705003294 Printed in the United Kingdom

Use of complex phonological patterns in speech processing: evidence from Korean[1]

NATASHA WARNER

Department of Linguistics, University of Arizona

JEESUN KIM

Department of Psychology, University of Melbourne

CHRIS DAVIS

Department of Psychology, University of Melbourne

ANNE CUTLER

Max Planck Institute for Psycholinguistics, Nijmegen

(Received 26 June 2003; revised 1 November 2004)

Korean has a very complex phonology, with many interacting alternations. In a coronal-/i/ sequence, depending on the type of phonological boundary present, alternations such as palatalization, nasal insertion, nasal assimilation, coda neutralization, and intervocalic voicing can apply. This paper investigates how the phonological patterns of Korean affect processing of morphemes and words. Past research on languages such as English, German, Dutch, and Finnish has shown that listeners exploit syllable structure constraints in processing speech and segmenting it into words. The current study shows that in parsing speech, listeners also use much more complex patterns that relate the surface phonological string to various boundaries.

1. Introduction

How does the wide variety of phonological patterns found in the world's languages affect native listeners' processing of speech and recognition of words and morphemes? Research in formal phonology has elucidated many types of changes that can happen to sounds between the underlying representation and the surface form. Research in psycholinguistics has shown that listeners can use phonotactic constraints to help them locate syllable and word boundaries. But how do these two types of research come together?

[1] This work was supported by a SPINOZA award from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek to the fourth author. We would like to thank Taehong Cho, Adam Ussishkin, James McQueen, Mike Hammond, Rachel Hayes-Harb, and Gwanhi Yun for useful discussion of this paper, as well as two anonymous *JL* referees for their helpful feedback. We are also very grateful to Prof. Mahn Young Lee for enabling us to conduct Experiment 1 at Korea University, and to Maarten Jansonius for research assistance.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Most past studies of listeners' use of phonological patterns for speech segmentation have examined relatively simple phonological patterns, such as constraints on which consonants can form a syllable onset cluster in a particular language. In many languages, the distribution of segments at the surface level is determined not only by syllable structure constraints, but also by a wide variety of interacting phonological phenomena. In Korean, the subject of this paper, many phonological patterns play a role in the mapping between underlying and surface forms. This paper investigates whether, and if so how, such complex phonology is also used in speech segmentation.

1.1 *Previous work on phonotactics and segmentation*

A common means of investigating speech segmentation is the 'word-spotting' task, in which a real word is embedded within a larger nonsense word. Listeners are asked to press a button as quickly as possible whenever they detect a real word within a stimulus, and to then say the word they detected. For example, a listener might hear the stimulus /vʌfæpʌl/, press the response button, and say the word *apple*. The environment in which the real word is embedded is manipulated to determine what features of the environment make it more or less difficult to segment a word out of surrounding speech. Reaction times and error rates (failure to detect target words) are measured, and faster reaction times and lower error rates are taken to indicate easier recognition of words (McQueen 1996).

McQueen (1998) investigated the influence of phonotactic constraints on speech segmentation in Dutch. He found that listeners can recognize a word more easily if phonotactic constraints indicate that syllable boundaries coincide with the boundaries of the word. For example, Dutch listeners were faster and more accurate in recognizing the word *rok* 'skirt' when it was embedded in a nonsense sequence /fim.**rɔk**/ than when it was embedded in a nonsense sequence such as /fi.d**rɔk**/. In Dutch, /mr/ cannot form a syllable onset, so the syllable boundary in /fim.**rɔk**/ must be aligned with the beginning of the real word. However, because Dutch has syllable-final devoicing, a syllable cannot end in /d/, so the syllable boundary in /fi.d**rɔk**/ cannot be at the beginning of the embedded word. (/dr/ is a possible onset cluster in Dutch.) This effect also held for words embedded before, rather than after, nonsense material. This indicates that listeners are sensitive to phonotactic constraints: phonotactic constraints provide information about syllable boundaries, and listeners use syllable boundaries in the task of locating word boundaries.[2] Dumay, Frauenfelder & Content

[2] These results could potentially stem from either the phonotactic constraints themselves, or from allophonic variation in the segments conditioned by syllable position. McQueen (1998) finds that the largest part of the effect is not acoustic, and Weber (2001) similarly finds very little in the phonotactic segmentation results that can be traced to syllable-conditioned allophonic differences.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

(2002) investigated phonotactic cues to syllable boundaries in French. They found the same effect for real words embedded at the onset of items, and a similar but weaker effect for words embedded at the end of items.

Weber (2001) replicated the phonotactic effect for American English. She found that English listeners are faster and more accurate at spotting real words such as *length* in sequences where the word onset is aligned with the syllable boundary, such as /zarʃ.**lɛŋθ**/ and /fun.**lɛŋθ**/, than in sequences where the word onset may not be at a syllable boundary, such as /fuk**lɛŋθ**/ and /joʊs**lɛŋθ**/. Syllable boundary information comes from the fact that sequences such as /ʃl/ and /nl/ cannot be syllable onsets in English, whereas /kl/ and /sl/ can. (Unlike McQueen's (1998) Dutch experiment, these sequences are simply ambiguous as to syllable boundary, rather than overtly misaligned with the word onset.)

Weber (2001) also investigated second language learners' use of phonotactic constraints for segmentation by testing native German speakers' ability to spot the same English words in the same materials. (The German speakers were all proficient in English as a second language.) Weber used the fact that German and English phonotactic constraints differ to construct an interesting set of comparisons. In one condition, the real word was aligned with the syllable boundary in both languages (e.g. /fun.**lɛŋθ**/, since /nl/ cannot be an onset in either language). In another, the word is aligned with the syllable boundary in English but is ambiguous in German (e.g. /zarʃ.**lɛŋθ**/, since /ʃl/ can be an onset in German but not in English). In a third, the word onset is aligned in German but is ambiguous in English (e.g. /joʊs**lɛŋθ**/, since /sl/ can be an onset in English but not in German). In the final condition, the syllable boundary is ambiguous in both languages (e.g. /fuk**lɛŋθ**/, since /kl/ is a possible onset for both). Weber found that the native German (but proficient in English) listeners were influenced by phonotactic constraints of both languages. They spotted words most quickly and accurately when the phonotactic constraints of both languages indicated that the syllable boundary was at the onset of the word, least quickly and accurately when neither language provided phonotactic cues, and intermediately when just one language provided phonotactic cues to the syllable boundary.

A study by Suomi, McQueen & Cutler (1997) tested listeners' use of another type of phonological regularity in segmentation: vowel harmony in Finnish. They found that Finnish listeners could recognize a real word more easily if it followed a nonsense syllable with a vowel from a different harmony class than if it followed a nonsense syllable with a vowel from the same harmony class as the vowels of the word itself. For example, listeners recognized a word such as *hymy* 'smile' faster and more accurately in a disharmonious sequence *puhymy* (/u/ and /y/ cannot occur within the same word in native Finnish words) than in a harmonious sequence such as *pyhymy*. That is, violations of vowel harmony give the listener information about the presence

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

of a word boundary. In parsing speech, listeners make use of the patterns in the surface string created by the phonology.

However, listeners do not make use of all language-specific phonological patterns in speech segmentation. Norris, McQueen, Cutler, Butterfield & Kearns (2001) and Cutler, Demuth & McQueen (2002) have shown that the 'Possible Word Constraint' (PWC) depends on a universal, rather than a language-specific, definition of what constitutes a possible word. Listeners find it more difficult to recognize a real word if they are presented with an item that has just a consonant (not an entire syllable) next to the real word, and this effect has been termed the Possible Word Constraint (Norris, McQueen, Cutler & Butterfield 1997). For example, for English listeners it is harder to find the word *apple* in the string /f**æpʌl**/, where recognizing it would require leaving the lone consonant /f/ unparsed, than in a string /vʌf**æpʌl**/, where an entire syllable (which could possibly be a word) is left over.

The criterion that at least a syllable must be left over in order to avoid hindering recognition appears to be quite general, even though particular languages may have additional constraints on what can actually form a word. For example, even though English words cannot end in a lax vowel, English listeners recognize a word embedded after a single CV syllable with a lax vowel just as easily as a word embedded after a CV tense vowel syllable. For example, recognition of words like *perturb* was not significantly different with preceding /dɛ-/ vs. preceding /da-/, although both were faster than recognition after a single consonant such as /s-/ (Norris et al. 2001). If listeners were making use of the language-specific information that an English word cannot end in a lax vowel, then they should infer that a word boundary does not follow the syllable /dɛ-/, and it should be more difficult to recognize a word embedded after such a syllable.

Furthermore, Cutler et al. (2002) showed that in Sesotho, which has a bisyllabic prosodic minimality requirement, recognition of words embedded after a single syllable vs. after two syllables was equally easy, whereas recognition of words embedded after a single consonant was hard. For example, *alafa* 'to prescribe' was embedded in the contexts *h**alafa***, *ro**alafa***, *pafo**alafa***, and only the single consonant condition differed from the other two (with the single consonant condition being harder). This again indicates that listeners make use of a universal criterion that there must be at least a syllable left over when recognizing words, but they do not make use of language-specific phonological information about what constitutes a minimal word beyond that.

Thus, there is reason to think that listeners exploit information from a variety of language-specific phonological patterns, but not all such patterns, in the task of segmenting speech into words. This fact has been incorporated into the Shortlist model of spoken word recognition (Norris et al. 1997), which models listeners' use of multiple boundary cues, including language-specific ones. The current experiments investigate whether listeners use a new

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

type of phonological information: surface regularities created by the interaction of several phonological processes within one language (Korean). They also investigate how listeners use phonological information in recognizing morphemes as well as words.

1.2 *Aspects of Korean phonology relevant to the current experiments*

Overall, Korean has a great many phonological alternations to consider, as a result of which surface forms can look very different from underlying forms. (See Cho & Sells (1995) and Cho (1998) for summaries of many features of Korean phonology.) For the purposes of the current experiments, we will focus on aspects of the phonology which affect coronal-/i/ sequences. Most of the information and examples below are derived from Rogers, You & Richards (1992), Cho & Sells (1995), and Cho (1998).[3] Phonological phenomena will be described here in derivational terms, but this is a descriptive convenience, not a theoretical claim about rule-based vs. constraint-based phonology.

A. INTERVOCALIC VOICING. Voiceless obstruents become voiced intervocalically (exemplified below in (1)–(4) and (6)–(7)).

B. PALATALIZATION. /t, t^h/ become palatalized /c, c^h/ before /i/ or /y/ if and only if the /i/ or /y/ is in a suffix or the copula, cf. (1). That is, the boundary between the consonant and the palatalizing sound must be a suffix boundary or the boundary between content word and copula in order for the palatalization to take place. Palatalization does not apply within a morpheme, cf. (2) (that is, it shows non-derived environment blocking), nor does it apply across a word boundary, cf. (3).

(1)	/mat-i/ [maji]	'firstborn-SUBJ'
	/mit^h-ita/ [mic^{h}ida]	'below-COP'
	/kat^h-i/ [kac^{h}i]	'same-ADV' ('together')
	/tat-hyʌsʌ/ [tac^{h}yʌsʌ][4]	'close-PASS-because'
(2)	/ʌti/ [ʌdi]	'where'
	/canti/ [candi]	'lawn'
	/t^{h}i/ [t^{h}i]	'dust particle'
(3)	/kot iʌsʌ/ [kodiʌsʌ]	'soon to follow'

[3] For transcribing Korean, [c, j] are used for voiceless and voiced palatal affricates, [y] is used for the palatal glide, and C' denotes a tense consonant. Otherwise, transcriptions are in IPA. The (former) phonemes /ɛ/ and /æ/ (recently neutralized or in the process of neutralization) are both transcribed as [ɛ]. The boundary between a stem and a suffix or the copula is marked with a hyphen, the boundary within a compound with a plus sign, and the boundary between words with a space.

[4] An additional alternation here merges the preceding consonant and /h/ into a single aspirated consonant.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

C. NASAL INSERTION AND NASAL ASSIMILATION. A nasal /n/ is optionally inserted before compound- or word-initial /i/. Furthermore, final /t/ becomes /n/ by assimilation to a following nasal, as is shown in (4).

(4) /pat^{h} + ilaŋ/ [panniɾaŋ, padiɾaŋ] 'field ridge'

(See point D below for the issue of /t^{h}/ vs. /t/.) Thus, when /t-i/ appears across a compound boundary, nasal insertion and nasal assimilation may apply, yielding a geminate nasal, but if nasal insertion does not apply, palatalization does not either because it does not apply across compound boundaries. (The conditions on nasal insertion are quite complicated, and are explained in detail by Cho (1998).)

D. CODA NEUTRALIZATION. Coronal obstruents are all neutralized to /t/ in coda position, cf. (5). (Like all coda stops in Korean, final /t/ is unreleased. This will not be marked elsewhere in the paper.)

(5)	/pyʌt^{h}/ [pyʌt$^{\neg}$]	'sunlight'
	/nac/ [nat$^{\neg}$]	'daytime'
	/pic^{h}/ [pit$^{\neg}$]	'light'
	/pis/ [pit$^{\neg}$]	'comb'
	/iss-ta/ [itt'a][5]	'exist'

Resyllabification applies across suffix boundaries, so if such consonants are word-final and a vowel-initial suffix follows, they are resyllabified, and coda neutralization does not apply, as in (6). However, resyllabification does not apply across word or compound boundaries, so coda neutralization does apply if the following morpheme is a separate word or the second part of a compound rather than a suffix, as in (7).

(6)	/nac-ɛ/	[najɛ]	'daytime-LOC'
	/pic^{h}-i/	[pic^{h}i]	'light-SUBJ'
(7)	/k'oc^{h} + iɬim/	[k'odiɾim, k'onniɾim][6]	'flower name'
	/pat^{h} ilku-ko/	[padilgugo]	'till the field and'

The examples in (7) demonstrate that because of coda neutralization, the surface sequences [di, nni] can also be derived from sequences other than /ti/, such as /t^{h}i, c^{h}i/, etc.

E. RARITY OF THE SEQUENCE /nni/ MORPHEME-INTERNALLY. A final point about the phonology of Korean which is relevant to the current work is that the sequence /nni/ is rare within a morpheme. (8) shows one high-frequency word with morpheme-internal /nni/, but there do not seem to be many such

[5] An additional process of the phonology, obstruent tensing, is evident here, but is not relevant to the current study.

[6] As in (4), the two alternatives reflect the optionality of nasal insertion.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

words.[7] There is a question suffix /-ni/, and (9) shows that a /nni/ sequence can occur because of nasal assimilation before this suffix, but this is the only inflectional suffix beginning with /-ni/ listed by Cho & Sells (1995), so /nni/ clusters across a suffix boundary are unlikely, except when this is the suffix.

(8) /ʌnni/ [ʌnni] 'older sister'
(9) /tat-ni/ [tanni] 'close-QUEST' 'are you closing?'

Thus, the /nni/ sequence is rare except across compound or word boundaries, but it does occur easily at compound boundaries, since the initial consonant of the second part of a compound can be geminated, as in (10).[8]

(10) /alɛ + ni/ [arɛnni] 'lower teeth'
/k'ɛ + nip^h/ [k'ɛnnip] 'sesame leaf'

The sequence /nni/ is unlikely to occur across word boundaries, however, because initial /n/ before /i/ has been deleted in a historical sound change, so a sequence of a word ending in /n/ (or a consonant that neutralizes to /n/) followed by a word beginning with /ni/ is unlikely, with the exception of a few recent loanwords. Since Korean has maximally CVC syllables, /nn/ cannot be either a coda or an onset. That is, the /nn/ sequence must be split across a syllable boundary.

1.3 *Outline of the paper*

In this paper, we will present two experiments designed to investigate the influence of Korean phonological patterns on listeners' processing of speech. Experiment 1, a word-spotting experiment, provides data on whether Korean listeners use phonotactic constraints in segmenting words out of connected speech. Experiment 2 examines listeners' use of affricate vs. stop forms which result from the palatalization rule (B above) more directly through a morpheme monitoring task. In that experiment, listeners are asked to decide whether they have heard the subject particle morpheme, which triggers palatalization in a preceding /t/ (see (1) above).

2. EXPERIMENT 1: SPOTTING /i/-INITIAL WORDS

2.1 *Rationale*

The phonological processes which create the surface distributions of coronal-/i/ sequences in Korean (section 1.2 above) are far more complex than those

[7] This may be because sequences of two identical consonants within the morpheme in general are rare in Korean, and such sequences may be viewed as two instances of the same consonant occurring together by coincidence rather than as a distinction between singleton and geminate consonants. However, what is important for the current experiment is only that /nni/ is rare within the morpheme, not whether or not /nn/ is a true geminate.

[8] Rather than gemination, this may be insertion of a genitive marker (spelled with the /s/ symbol), which then assimilates to the following nasal, as in (4).

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

which produce the phonotactic constraints whose effect on word-spotting has been tested in Dutch (McQueen 1998), English (Weber 2001), and French (Dumay et al. 2002). For Dutch, the Sonority Sequencing Principle is involved in the fact that clusters such as /mr/ cannot form onsets, and syllable-final devoicing leads to the fact that /dr/ must be an onset rather than being split across a syllable boundary. For English, the Sonority Sequencing Principle and more specific syllable structure constraints are involved in determining that neither /nl/ nor /ʃl/ can be syllable onset clusters, while /sl/ and /kl/ can. The situation is similar for French.

By contrast, the Korean patterns in coronal-/i/ sequences derive from the interaction of palatalization, nasal insertion, nasal assimilation, coda neutralization, and a constraint on the occurrence of identical consonant sequences, as well as the general maximal CVC syllable structure of Korean. Furthermore, the same coronal-/i/ sequences that are created by these various phonological processes can also simply occur underlyingly, with none of the above processes except intervocalic voicing affecting them.

Thus, the relevant surface phonotactic patterns in Korean differ considerably from those which have been tested in other languages. In Korean, there are multiple interacting segmental alternations involved in creating the surface patterns. There can be a great many differences between underlying and surface forms in coronal-/i/ sequences. In the previous work on phonotactic patterns, the Dutch final devoicing, which forces the syllabification of /fid.rɔk/, is such a case, but most of the other factors examined are syllable structure constraints rather than alternations that change the form.

Furthermore, most of the phonotactic constraints that have been tested previously for their influence on speech segmentation were constraints on syllable structure. McQueen (1998) argues that such constraints give the listener information about the location of syllable boundaries, and that listeners use syllable boundaries to locate word boundaries. Specifically, listeners are said to hypothesize that a new word begins at a syllable boundary. Most of the phonological patterns discussed here for Korean do not involve syllable boundaries. They make reference to morpheme boundaries, word boundaries, the boundary between parts of a compound, or the lack of any boundary at all, but only coda neutralization and the impossibility of /nn/ within a syllable make obvious reference to syllable structure. Therefore, if the Korean phonotactic patterns provide information about word boundaries to the listener, it cannot be entirely by means of providing information about syllable boundaries. This has the further implication that if speech segmentation takes place primarily by locating syllable boundaries, Korean listeners should not be much helped by the phonotactic information available in their language. However, if they can make use of information about a variety of boundary types, that information should indeed be useful.

Because so many phonological processes influence coronal-/i/ sequences in Korean, the surface distribution of such sequences is not immediately clear

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

from a listing of the processes involved. Furthermore, the fact that many of the processes apply or fail to apply at specific types of boundaries makes a summary of the relevant patterns necessary. (11) shows what type of boundary, if any, the surface sequences [ji, di, nni] must occur at, classified by the underlying form from which the sequence derives. The phonological patterns above (A–E) that are involved in each case are listed as superscripts for reference.

(11)	SURFACE	UNDERLYING	BOUNDARY ENVIRONMENT
	[ji]	/ti/	suffix boundary or copula^{AB}
		/ci/	suffix or no boundary^{A}
	[di]	/ti/	word/compound or no boundary^{A}
		/ci, c^{h}i, t^{h}i/…	word/compound boundary^{AD}
	[nni]	/ti, ci, c^{h}i, t^{h}i/…	compound boundary^{CD}
		/tni, cni, c^{h}ni, t^{h}ni/…	suffix boundary with *-ni* suffix^{CDE}
		/nni, ni/	compound boundary^{E}
		/nni/	no boundary (rare)^{E}

We can summarize the distributions relative to types of boundaries, regardless of underlying source:

- Surface [ji] can only occur either crossing a suffix boundary or morpheme-internally (if the copula is counted as equivalent to a suffix). It cannot occur spanning a word or compound boundary, because coda neutralization would then apply.
- [di] can occur either at a word or compound boundary or morpheme-internally. It cannot occur spanning a morpheme boundary, because palatalization would apply.
- [nni] usually occurs at a compound boundary. Apart from compound boundaries, it can occur at a suffix boundary only if the suffix is the [-ni] question suffix, and it can occur morpheme-internally in the word /ʌnni/ 'elder sister' and any other such cases, thought to be rare.
- The distribution of surface [c^{h}i] parallels that of [ji] (if the underlying form includes aspiration), and [t^{h}i] largely parallels [di].[9]

Although the origin of these patterns is complex, the surface distributions can be viewed as follows: [ji] does not occur at a word/compound boundary, [nni] occurs predominantly at a compound boundary, and [di] can occur either at a compound boundary or within the morpheme.

In this experiment, we investigate whether Korean listeners can exploit these phonotactic distributions to help them locate word boundaries in speech. We use a word-spotting task (as reviewed in section 1.1 above), in which real words are embedded within a larger nonsense string, and listeners

[9] [t^{h}i] cannot occur across a word or compound boundary because of coda neutralization. It otherwise parallels [di].

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

are asked to say the word they have heard whenever they hear a real word within an item. In this experiment, real words beginning with /i/ were embedded after [j]/[cʰ], [d]/[tʰ], or [nn]. Examples of materials appear in (12).

(12) /imin/ 'migration' in [pʰyojimin, pʰyodimin, pʰyonnimin]
/ipyʌl/ 'separation' in [pʌswɛcʰibyʌl, pʌswɛtʰibyʌl, pʌswɛnnibyʌl]

Based on the distributions of sequences discussed above, one would expect that listeners would have the most difficulty detecting the embedded words in the [ji]/[cʰi] sequences, because these sequences cannot occur at a word or compound boundary. Detection of embedded words should be easiest for [nni], because this sequence usually occurs at compound boundaries, and also signals a syllable boundary.[10] Detection of embedded words should be intermediately difficult for [di], because this sequence can, but does not always, occur at a word or compound boundary.[11] In word-spotting, easier detection of words is reflected by faster reaction times and lower error rates, and more difficult detection by slower reaction times and higher error rates (McQueen 1996).

2.2 *Methods*

2.2.1 *Materials*

A native speaker of Korean created the materials by embedding 42 /i/-initial Korean words of two or three syllables after nonsense material. The nonsense material was phonotactically possible in Korean (in that it consisted of Korean phonemes and had maximally consonant–glide–vowel–consonant syllable structure). The nonsense material consisted of either one or two syllables. For each nonsense plus word combination, the final consonant (cluster) of the nonsense was [j] or [cʰ] in one condition, [d] or [tʰ] in another, and [nn] in the third, as shown in (12). The full set of experimental stimulus materials appears in appendix 1. (Hereafter, we will refer to both [j] and [cʰ] as the [j] condition, and to both [d] and [tʰ] as the [d] condition.) The aspirated consonants were used as well as the voiced ones in order to make the materials more varied and more representative of the language.

The materials contained no embedded words longer than a CVC syllable other than the targets. It is preferable in word-spotting that there be no embedded words other than the target items, so that listeners do not detect words other than the intended targets. However, in Korean, it is impossible

[10] There is an additional potential reason why the [nni] condition might make detection of an embedded /i/-initial word relatively easy. Because of the historical deletion of /n/ before /i/ in words that historically began with /ni/, there are alternations between /i/ and /ni/, and no words beginning with /ni/. An embedded [ni] sequence may therefore be a cue to an /i/-initial word. This is not in conflict with the explanation presented above.

[11] Although [tʰi] cannot occur at a word/compound boundary.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Number	Characteristics	Example
9	/i/-initial real words embedded after non-coronal	[ihɛ] 'comprehension' in [sɛnwʌmihɛ]
44	Non-/i/-initial real words embedded after non-coronal	[kwagʌ] 'the past' in [sʌgwagʌ]
14	/i/-initial non-words embedded after [j]	[ikʰam] in [kwɛjikʰam]
14	/i/-initial non-words embedded after [d]	[ijop] in [cʰudijop]
14	/i/-initial non-words embedded after [nn]	[ibil] in [hinnibil]
127	Non-words of similar length to other stimuli, /i/ rare	[pɛhyobʌ]

Table 1
Types and numbers of fillers, with examples, for Experiment 1

to avoid other embeddings entirely without resorting to phonotactically highly improbable (if not impossible) nonsense syllables, since most CV and CVC syllables of Korean constitute words. It is therefore possible that listeners would sometimes respond to a word other than the target. For this reason, it is standard in word-spotting that listeners must say the word they have detected out loud, and any responses to words other than the intended one are excluded as errors. It is worth noting that the problem of avoiding unintended embeddings is also quite difficult for Japanese, and word-spotting experiments have been successfully carried out in that language, without excessively high false alarm rates (McQueen, Otake & Cutler 2001).

The materials included a relatively large number of recent foreign borrowings, as well as a few function words. Inclusion of such words was necessary in order to find an acceptable number of items with initial /i/ and minimal other embedded words. However, each item is embedded after each of the three phonological contexts of interest ([j, d, nn]), and it is the comparison across these conditions that is important, rather than across particular items. Therefore, effects specific to particular items should not be problematic.

Several types of fillers were created, as shown in table 1. The non-word fillers did not contain any embedded words longer than CVC. There were an unusually large number of fillers, and number of types of fillers, but this was necessary because all of the targets began with the same phoneme, /i/. Therefore, a wide variety of real-word and non-word fillers was used to prevent listeners from developing strategies such as listening for a real word every time they heard the phoneme /i/. A further 15 similar items, including both test items and representative fillers, were constructed for use as practice stimuli.

A female native speaker of Korean was recorded producing all the items and fillers. The speech was recorded to digital audio tape in a sound-treated

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

booth using high quality equipment. It was then sampled at 16,000 Hz and transferred to computer for further analysis using the Xwaves software. Each stimulus was spliced out of the recording. The 42 test items were divided into three groups for counterbalancing, such that each group of listeners heard 14 test items embedded after [j], 14 embedded after [d], and the remaining 14 embedded after [nn]. The 42 test items and 222 fillers were arranged into two different pseudo-random orders (one the reverse of the other), such that items containing real words (either test items or fillers) never occurred immediately after each other.

2.2.2 *Subjects*

Forty-eight native Korean speakers, primarily students of the Psychology Department at Korea University in Seoul, participated in the study. They were paid a small fee for their participation. All had Korean as their first language and had attended high school in Korea. The subjects were divided into three groups of 16 for counterbalancing purposes.

2.2.3 *Procedures*

Subjects were tested at Korea University in Seoul. Testing was done individually in a sound-treated room. A native speaker of Korean carried out all testing. Subjects were instructed in Korean that whenever they heard a real Korean word within an item, they should say the word they had heard out loud. Reaction times were recorded by means of a voice key: the reaction time is the onset of the subject's speech.[12] When a subject said a detected word out loud, the experimenter decided whether the word the subject had detected was the target word, and pressed a button that recorded this information. If a subject did not respond at all, there was either a three-second or a five-second time-out, after which the next stimulus would be played.[13] Stimulus presentation and response collection were controlled by computer (PIII, 800 Mhz) with the DMDX software program (Forster & Forster 2003).

Subjects were randomly assigned to one of six versions (the three counterbalancing conditions by the two stimulus orders). All subjects were presented

[12] The use of a voice key response rather than a button-press response is not problematic in this experiment because all targets begin with the phoneme /i/, and therefore the time from initiation of the response to activation of the voice key mechanism should not differ systematically across conditions.

[13] Approximately 3/4 of the subjects had the five-second time-out, and 1/4 had the three-second time-out. The change was a matter of procedural convenience, and the cut-off for slow responses chosen (see Results, section 2.3) makes the difference irrelevant for data analysis purposes.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

first with the practice test, given a chance to ask questions, given a chance to repeat the practice test if they wished, and then presented with the experimental stimuli. The entire experiment lasted approximately 25 minutes, and subjects could take a break in the middle of the experiment. Reaction times from item onset (by means of speech onset) and error rates were recorded.

2.3 *Results*

Reaction times (RTs) were converted to RTs from offset of the item. Responses to which listeners responded by indicating that they had detected a word other than the target (1.8% of all trials with target-bearing stimuli) were treated as errors, along with target-bearing items to which listeners failed to respond. RTs faster than 200 ms or slower than 1890 ms from item offset were counted as errors. Responses faster than 200 ms are very likely to be responses begun before the subject heard the entire item (thus before the subject could know what the word was), because it takes some time to plan and initiate a response. Responses slower than a few seconds are likely to reflect a different processing mechanism, perhaps use of a conscious strategy rather than online spoken word recognition. Inclusion of outliers also lowers the power of statistical analyses. (Ratcliff (1993) discusses reasons for exclusion of outliers, distinctions between types of outliers, and various methods of excluding outliers.) The exact cut-off of 1890 ms was chosen based on the three-second time-out from stimulus onset and the duration of the longest stimulus, so that the same response window was analysed for all stimuli. (This is within the range of slow response cut-offs for other word-spotting experiments, e.g. Suomi et al. (1997) and McQueen (1998), and the use of such a cut-off is standard for this task.) These fast and slow responses represented approximately 8.1% of all responses. Seven subjects were excluded from the by-subjects analysis because they either failed to respond to any item in one preceding consonant condition, or had error rates of at least 75% for all three conditions. Seven items were excluded from the by-items analysis because they either received no responses at all in one preceding consonant condition or had error rates of at least 75% in all three conditions. (These items are marked in the item list in appendix 1.) Remaining RTs and error rates were averaged across items for the by-subjects analysis and across subjects for the by-items analysis.

The results appear in figure 1. For each analysis, a single-factor ANOVA was used, with preceding consonant ([j, d, nn]) as the within-subjects factor. The overall effect of preceding consonant was significant both for RTs ($F1(2,80) = 11.81$, $p < 0.001$, $F2(2,68) = 17.93$, $p < 0.001$) and for the error rate ($F1(2,80) = 45.44$, $p < 0.001$, $F2(2,68) = 33.96$, $p < 0.001$). Comparing the [ji] and [di] conditions, the [ji] condition had slower RTs, although this effect was not quite significant in the by-subjects analysis ($F1(1,40) = 3.41$, $p = 0.072$,

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

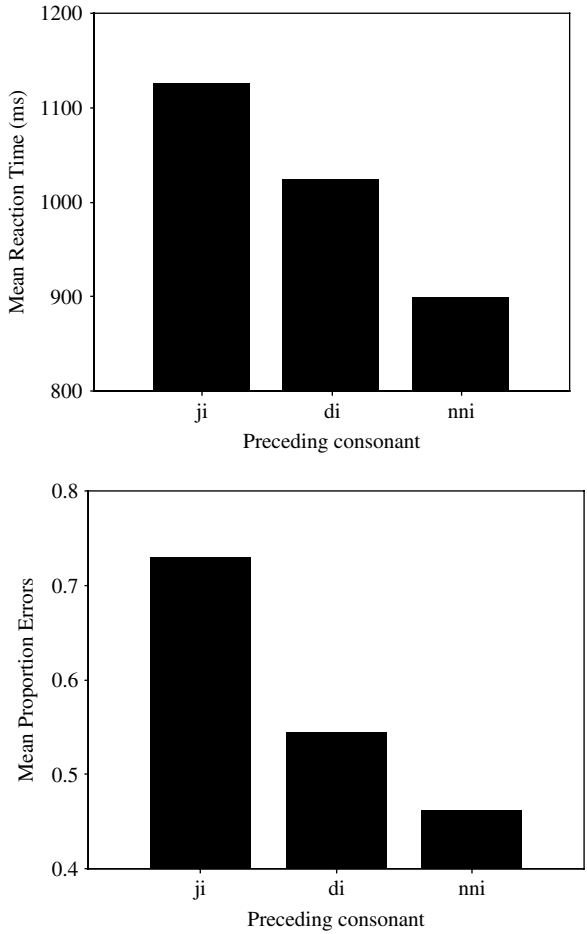


Figure 1
Results for Experiment 1 (word-spotting). A. Reaction times. B. Error rates.

$F2(1,34) = 7.40$, $p < 0.02$). The [ji] condition had significantly higher error rates ($F1(1,40) = 62.14$, $p < 0.001$, $F2(1,34) = 26.34$, $p < 0.001$). Comparing the [di] and [nni] conditions, the [di] condition had significantly slower RTs ($F1(1,40) = 13.77$, $p < 0.005$, $F2(1,34) = 9.48$, $p < 0.005$) and higher error rates ($F1(1,40) = 6.93$, $p < 0.02$, $F2(1,34) = 7.88$, $p < 0.01$).[14]

[14] Because [tʰi] cannot occur at a word boundary, the results were also checked without the seven sets of items with aspiration. The pattern of results was similar to the overall results, except that the [di] condition is not significantly different from the [nni] condition. This is what one would expect if the [tʰi] items are more difficult because this sequence cannot cross a word boundary.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

2.4 *Discussion*

The results show that Korean listeners detect embedded /i/-initial words more quickly and accurately after [d] than after [j], and even more quickly and easily after [nn] than after [d]. This indicates that Korean listeners are able to use the regularities in the distribution of the sequences [ji, di, nni] relative to various boundaries to help them in locating word onsets, and thus in recognizing words. Thus, even when surface phonotactic regularities are created by several interacting phonological alternations, rather than by simpler phonotactic constraints, listeners are able to use those regularities to help them segment speech into words.

Although the error rates are high overall, word-spotting is a relatively difficult task, and these error rates are not substantially higher than those in some other word-spotting experiments (Cutler et al. 2002). Particularly for the [ji] condition, where a word boundary should be impossible, the high error rate is not surprising. Since error rates and RTs show the same pattern, the results are reliable.

These results show that Korean listeners are sensitive to the overall patterns of the phonology. We would also like to examine listeners' sensitivity to particular phonological alternations. Furthermore, since the phonological processes discussed above make reference to within-word morpheme boundaries as well as word boundaries, we would like to investigate the influence of phonological alternations on processing of morphemes as well as on word segmentation. The word-spotting task investigates the segmentation of words from connected speech, but it does not assess listeners' recognition of morphemes within words. In Experiment 2 we turn to the issue of morpheme recognition.

3. Experiment 2: morpheme monitoring

3.1 *Rationale*

In this experiment, we investigate more directly whether a particular phonological alternation, the palatalization described in section 1.2 point B above, affects listeners' online processing of speech. Palatalization is of particular interest because it is a so-called 'non-derived environment blocking effect' – palatalization affects /t/ before a suffix, but not /t/ within a morpheme (see (1)–(2) above). Although the term 'non-derived environment' comes from derivational phonology, even in a constraint-based theory such as Optimality Theory, one clearly must make reference to particular types of boundaries in order to map the relation between underlying and surface forms where palatalization is involved. The palatalization alternation is specified as applying at the boundary between a stem and a suffix or the copula, but not at stronger boundaries (compound or word), and also not in the absence of any boundary. This fact cannot be explained through syllable

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

structure. None of the other phonological alternations discussed above make reference to morpheme-internal vs. suffix boundary position. This condition on palatalization makes it more abstract than many phonological patterns.

Even though listeners are sensitive to the distribution of phonemes relative to each other, as shown in Experiment 1, it is possible that the abstract restriction on palatalization applying only in derived environments is not active in online processing. The presence of a pattern in the phonology is not a guarantee that listeners attempt to make use of that pattern for synchronic lexical processing. Furthermore, the presence or absence of palatalization in a /ti/ sequence gives information specifically about the boundary between a stem and a suffix or the copula. The word-spotting task used in Experiment 1 investigates listeners' ability to locate word boundaries, but not their processing of sub-word units such as suffixes. Experiment 2 addresses directly whether listeners can use the presence or absence of palatalization to recognize a suffix morpheme, the subject particle /-i/. This particle does provide an environment for palatalization, as shown in (1) above.

In this experiment, we address processing of morphemes directly by asking subjects to listen for a particular morpheme, the subject particle, rather than to listen for real words in general. Subjects were presented with words and non-words, some of which ended in /i/ (which constitutes one allomorph of the subject particle). Subjects were asked to press a button whenever the item ended with the subject particle. This can be termed a morpheme-monitoring task. As far as we know, morpheme-monitoring is a novel task. However, it is closely related to phoneme-monitoring, in which listeners are asked to press a button whenever they hear a particular sound (Connine & Titone 1996). Fragment-monitoring, in which subjects are asked to press a button whenever they hear a particular string of phonemes (e.g. /bal/, cf. Mehler, Dommergues, Frauenfelder & Segui 1981; Frauenfelder & Kearns 1996), is another similar task. In this experiment, we present listeners with items ending in [ji] and others ending in [di] in order to determine whether listeners will use the information about morpheme boundaries provided by the palatalization alternation in order to decide whether a given item contains the subject particle.

3.2 *Methods*

3.2.1 *Materials*

The materials included both real words and non-words in several conditions, along with several types of fillers. The materials were recorded and processed in the same way as for Experiment 1 above. One real-word condition consisted of real words in which the root itself ends in [ji] or [c^hi]. That is, the final /i/ is not the subject particle. (Hereafter we will refer to this as the real-word [ji] condition, abstracting away from aspiration.) The real-word [di]

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Number	Characteristics	Example
19	Real words ending in [di]	[ʌdi] 'where'
19	Real words ending in [ji]	[ɨiji] 'intention'
38	Non-words ending in [di]	[kwadi] –
38	Non-words ending in [ji]	[kwaji] –

Table 2
Types of experimental items, with examples, in Experiment 2

condition consisted of real words with the root ending in [di] or [t^hi] (also not containing the subject particle).

There were two important non-word conditions: one in which the non-word items ended in [ji/c^hi] ('non-word [ji]'); and another in which the items ended in [di/t^hi] ('non-word [di]'). Since these items are non-words, one cannot state whether the [i] in the [ji/c^hi] condition should be interpreted as the subject particle or not. The non-word [di/t^hi] condition, however, should not be interpretable as containing the subject particle, because if the final [i] were a subject particle, it would trigger palatalization.

The experimental items consisted of 19 real-word [ji] items, 19 real-word [di] items, 38 non-word [ji] items, and 38 non-word [di] items. (There were twice as many non-word items so that they could be divided into two groups for counterbalancing.) Table 2 lists and exemplifies the categories of experimental items. A full list of experimental test items appears in appendix 2.

The experiment also included a variety of filler items (table 3). The most important group of fillers were real words with the root ending in /t^h/, /c/, or /c^h/[15] followed by the subject particle /i/. These items contained a real word and ended in [ji/c^hi], just as in the real-word [ji] condition, but in this condition the final [i] was indeed the subject particle. These items can serve as a check on the subjects' ability to perform the task, as well as providing subjects with some items for which the correct response is clear. A further 14 items similar to the experimental items and fillers were created for use as a practice test.

If listeners are able to make use of the fact that the palatalization alternation fails to apply within the morpheme in speech processing, when asked to decide whether a nonsense word contains the subject particle or not, they should be less likely to parse the final /i/ as the subject particle if it follows [d] than if it follows [j]. Therefore, we predict that listeners will respond that they have heard the subject particle less often in the non-word [di] condition than

[15] Very few Korean words end in underlying /t/ without aspiration. The majority of the items in this filler group had aspiration because of the lexical frequency of the relevant sequences.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Number	Characteristics	Example
19	Real words + subj. particle	[naji] 'day-SUBJ'
131	Real words with no /i/	[kagu] 'furniture'
10	Real words ending in a consonant other than [j, d, nn], then /i/	[koŋgi] 'air'
10	Real words with medial /i/	[kyosil] 'classroom'
86	Non-words with no /i/	[nok^{h}ʌ]
10	Non-words ending in a consonant other than [j, d, nn], then /i/	[yup^{h}i]

Table 3
Types and numbers of fillers, with examples, for Experiment 2

in the non-word [ji] condition. We also expect this to affect reaction times: when listeners respond that they have heard the subject particle, we predict listeners to be faster for the non-word [ji] condition than the non-word [di] condition. When listeners respond that they have not heard the subject particle, we predict the opposite pattern for reaction times. For real-word items, listeners have lexical information as well as phonotactic information available, so their responses are expected to be quite accurate. Listeners are expected to respond overwhelmingly that the real-word [ji] and real-word [di] items do not contain the subject particle. The availability of lexical information might make the phonotactic information unnecessary, and therefore responses and reaction times might not differ across the real-word [ji] and real-word [di] conditions. Alternatively, the additional phonotactic information might lead listeners to respond more quickly, or perhaps even more accurately, in the real-word [di] condition than the real-word [ji] condition.

3.2.2 *Subjects*

Twenty native Korean speakers who had not participated in Experiment 1 participated in this experiment. The subjects were native Korean speakers who were students at the University of Melbourne, Australia, students at English as a Second Language schools in Melbourne, and visitors to Melbourne. The subjects all had Korean as their first language, and had attended high school in Korea. They were all speakers of the standard Seoul dialect. Subjects were paid a small fee for their participation. The subjects were randomly divided into two groups for counterbalancing purposes.

3.2.3 *Procedures*

The 76 experimental items and 266 filler items were arranged in a pseudo-random order, such that there were at least two non-word items before each

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

real-word item. The two counterbalanced groups of subjects received the materials in the same order, with the only difference being the substitution of the items in the counterbalanced non-word [di] and [ji] conditions. Subjects were first presented with the practice list, and subsequently with the experimental list.

Subjects were instructed to decide whether each stimulus contained the subject particle /-i/, and to press a button marked 'yes' if it did and a button marked 'no' if it did not. Subjects were told that there would be non-words in the experiment, and that for non-words ending in the sound /i/, they should choose whether they thought that /i/ seemed to be the subject particle or not. Examples of real words where the final /i/ was or was not the subject particle were given. These examples did not have [d] or [j] before the [i].

Testing procedures were the same as for Experiment 1, except that subjects were asked to press a button, either 'yes' or 'no', for each stimulus, rather than responding only when they detected the target. This two-response method was used so that reaction times for both 'yes' and 'no' decisions could be evaluated. As in Experiment 1, if a subject failed to respond (with either 'yes' or 'no') within five seconds from the beginning of the stimulus, the experiment continued with the next stimulus. Testing was performed in a sound-protected booth at the University of Melbourne Psychology Department. The experiment lasted approximately 25 minutes, and subjects were given the opportunity to take a break. The proportion of 'yes' responses and reaction times (for both 'yes' and 'no' responses) were evaluated.

3.3 *Results*

Responses with reaction times faster than 200 ms from the end of the item or slower than 2000 ms from the end of the item were excluded. Since the number of syllables in the stimuli was varied, and /i/ could also occur word-medially, subjects could not know until they had heard the entire stimulus whether an /i/ was the subject particle or not, so responses faster than 200 ms are likely to be errors. Responses slower than 2000 ms are likely to represent a different processing mechanism, or perhaps a conscious strategy. These cut-offs resulted in 5.1% of all responses to the experimental conditions being excluded as too fast and 6.1% being excluded as too slow. The difference in the cut-offs used as compared to Experiment 1 is a result of the different time-outs used for the differing experimental setups, and is not large enough to affect the results substantially.

No subjects or items were excluded for failing to meet a criterion, as was done for Experiment 1. In the current experiment, 'yes' response rates to the non-word test items were low for most subjects and items, and for the non-word [ji] items, it cannot be said that there is a correct answer, so the non-word items cannot be used to set a criterion. Furthermore, performance on the real-word conditions indicates that all subjects performed adequately.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

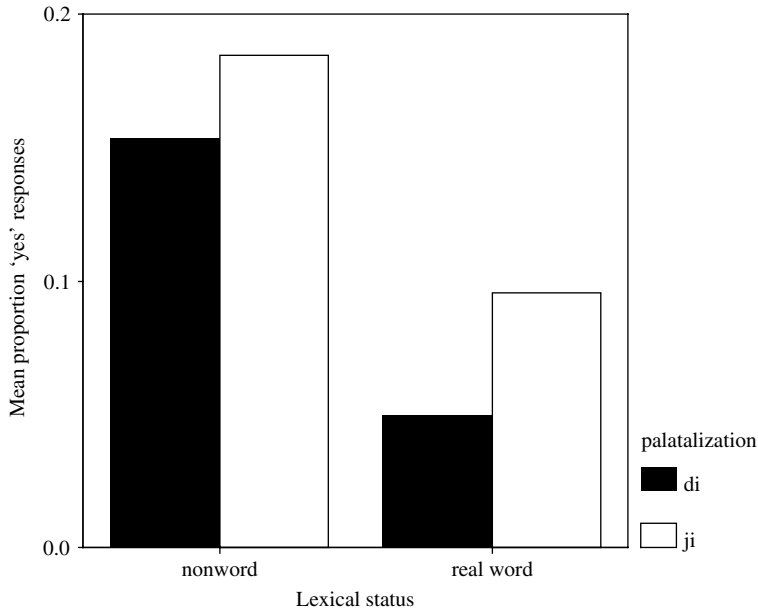


Figure 2
Proportion of 'yes' (subject particle) responses in Experiment 2 (morpheme-monitoring)

Although two subjects incorrectly responded 'yes' to more than 25% of the real-word [ji] items (27 and 31%), no subject performed poorly on both real-word conditions and the real word plus particle fillers.

The remaining data after exclusion of very fast and slow responses was averaged across items for the by-subjects analysis and across subjects for the by-items analysis. The proportion of 'yes' results for all experimental conditions is presented in figure 2. The data for the real-word (not particle) and the non-word conditions is analysed as a two-factor design, with lexical status (real-word vs. non-word) and palatalization ([di] vs. [ji]) as the factors. For the by-subjects analysis, both factors are repeated measures, but for the by-items analysis, neither factor is repeated measures, because the items cannot be matched across the word and non-word conditions and the real words cannot be matched across the palatalization condition.

Listeners responded that they had heard the subject particle ('yes' responses) significantly more often with non-words than with real words ($F1(1,19) = 8.24$, $p < 0.02$, $F2(1,108) = 12.30$, $p < 0.005$). There were more 'yes' responses to palatalized ([ji]) items than to non-palatalized ([di]) items as well, although this effect was only fully significant in the by-subjects analysis ($F1(1,19) = 5.19$, $p < 0.04$, $F2(1,108) = 3.53$, $p = 0.063$). The interaction of lexical status and palatalization was not significant (F1 and F2 both < 1). In order to confirm that listeners were able to do the task, one can examine the

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

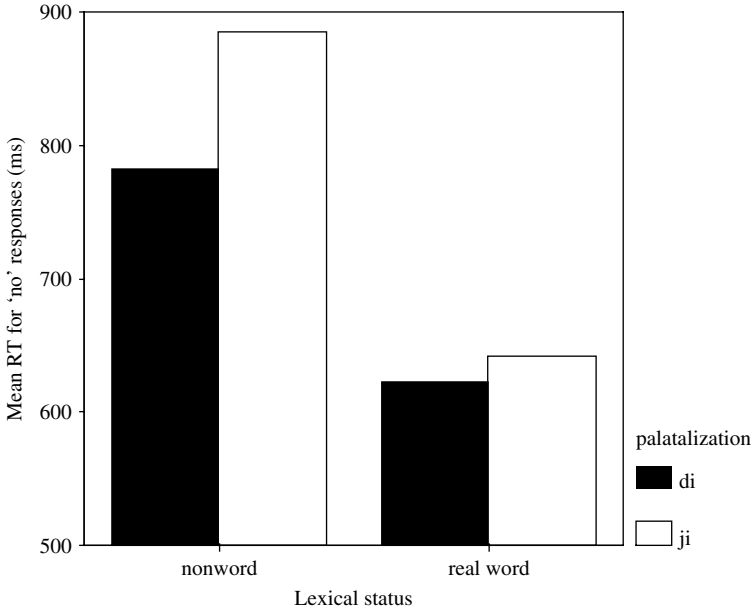


Figure 3
Reaction times for 'no' (not subject particle) responses for Experiment 2 (morpheme-monitoring)

proportion of 'yes' responses to the real-word fillers which actually did end in the subject particle. Listeners correctly responded that they had heard the subject particle with 91% of such stimuli, indicating that they were able to perform the task.

Because the experiment had a time-out and both 'yes' and 'no' responses were recorded, the proportion of 'yes' responses is not exactly the inverse of the proportion of 'no' responses. However, during the entire experiment, in only 11 tokens within the experimental test conditions did a subject fail to respond either 'yes' or 'no' to a stimulus within the five seconds allowed. This represents 0.6% of all stimulus presentations in the experimental test conditions. Therefore, no separate analysis of the proportion of 'no' responses was undertaken.

Because subjects responded 'yes' to a small proportion of items for all conditions, we chose to focus on the reaction times for 'no' responses. Figure 3 shows the relevant means. 'No' responses were slower with non-words than with words, and were slower with [ji] conditions than with [di] conditions. The effect of lexical status was significant ($F1(1,19) = 20.38$, $p < 0.001$, $F2(1,108) = 21.28$, $p < 0.001$), and the effect of palatalization was significant in the by-subjects analysis ($F1(1,19) = 4.99$, $p < 0.04$, $F2(1,108) = 3.53$, $p = 0.063$). The interaction of lexical status and palatalization was not significant ($F1(1,19) = 2.78$, $p > 0.1$, $F2 < 1$).

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

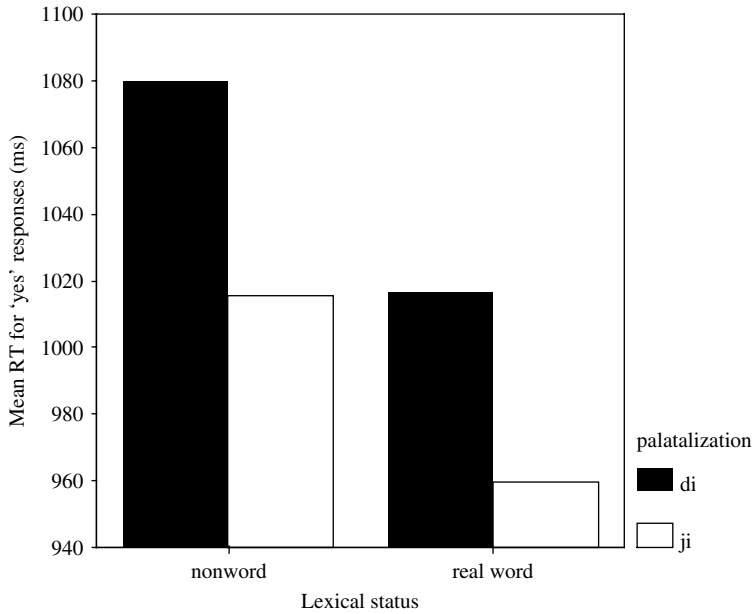


Figure 4
Reaction times for 'yes' (subject particle) responses for Experiment 2 (morpheme-monitoring)

Since the non-word conditions are the ones of particular interest, we also tested the effect of palatalization for just the non-word items. In this single-factor analysis, palatalization is a repeated measures factor for both the by-subjects and the by-items analysis. The effect of palatalization is significant both by subjects and by items in this analysis ($F1(1,19) = 4.53$, $p < 0.05$, $F2(1,37) = 5.70$, $p < 0.03$). This provides further support for the reliability of the effect of palatalization on RTs.

Although there were not many 'yes' responses to the four conditions considered above, the RTs for 'yes' responses may still be informative. Figure 4 shows these results. However, some subjects did not respond 'yes' to any items at all, and some items received no 'yes' responses at all, so there are fewer and unequal numbers of subjects and items, lowering the power of any statistical tests. Neither main effect nor the interaction was significant for this data (lexical status: $F1 < 1$, $F2(1,73) = 1.54$, $p > 0.1$; palatalization: $F1(1,7) = 1.24$, $p > 0.1$, $F2(1,73) = 2.53$, $p > 0.1$; interaction: F1 and $F2 < 1$).

3.4 *Discussion*

These results show that Korean listeners can use the information provided by the palatalization alternation in parsing a string into morphemes, and

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

specifically in deciding whether a final phoneme /i/ is the subject particle morpheme or simply part of a stem. The crucial fact about the palatalization alternation is that it applies to /ti/ sequences with a suffix boundary before the /i/, but not to /ti/ sequences within a morpheme. Therefore, the unpalatalized surface sequence [di] provides the listener with the information that there is no suffix boundary before the [i], and therefore that the [i] is not the subject particle. Despite the abstractness of this distribution, listeners are able to apply it in detecting the subject particle in real words and in deciding whether an /i/ at the end of a non-word should be considered as the subject particle. Thus, listeners not only use phonotactic information to help them parse the speech signal into words, they also use phonotactic information, even extending to the fact that a particular alternation does not apply within morphemes, to help them recognize morphemes within words. This finding is novel in two ways. First, it shows that phonotactic information is useful in segmentation within words, as well as for finding word boundaries. Second, it shows an effect of the patterns created by a very abstract phonological concept, non-derived environment blocking, on listeners' speech processing behavior.

Listeners responded that final /i/ was the subject particle significantly less often for real words (where it was indeed not the subject particle) than for non-words (where there was no correct answer). They also responded more quickly to real words than non-words for both 'yes' and 'no' responses (although this is only significant for 'no'). All of these effects are consistent with the idea that when the stimulus is a real word, listeners can use both lexical information and phonotactic information to parse the morphemes. The fact that listeners did not respond 'yes' very often to non-words, particularly in the [ji] condition, might seem surprising. It appears that listeners depended on lexical information to some extent, and were biased toward saying that /i/ was not the subject particle unless they could parse out the preceding root. However, despite this apparent use of lexical information, the phonotactic constraint clearly also influenced listeners' behavior.

It might be argued that the effect of palatalization on processing in this experiment is actually an effect of transitional probabilities rather than an effect of a phonological phenomenon. [di] might be a low-frequency sound sequence in Korean, and if the sequence [di] has a low transitional probability, listeners might be faster to parse [i] off from the preceding material, leading to faster reaction times. Alternatively, listeners might find it difficult to process low transitional probability sequences, which could lead to slower reaction times whenever a low probability sequence appears. (We make no prediction regarding how transitional probabilities would affect the percentage of 'yes' or 'no' responses.)

We are not aware of any direct evidence on the question of how transitional probabilities across boundaries affect speed of segmentation (although see van der Lugt (2001) for the effect of transitional probabilities

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

within the word on segmentation, and Vitevitch & Luce (1999) for effects of transitional probabilities on processing of entire items). However, whether a low transitional probability sequence facilitates or inhibits the process of recognizing the second phoneme, one would expect any effect of transitional probabilities on reaction times to be in a consistent direction. In the current experiment, the effect of palatalization on reaction times is in the opposite direction for the 'no' responses and the 'yes' responses. (It is not significant for the 'yes' responses, but it is clearly not in the same direction as for the 'no' responses.) Listeners are faster to conclude that the final /i/ is not the subject particle in the [di] sequence than in the [ji] sequence, but they are (non-significantly) faster to conclude that it is the subject particle in the [ji] sequence than in the [di] sequence. This is consistent with the interpretation that the unpalatalized [di] sequence facilitates detection of the morpheme boundary, while the palatalized [ji] sequence does not. It is not consistent with the interpretation that parsing is simply slower or faster in low transitional probability sequences. Thus, the effect is one of an abstract phonological pattern on parsing, not an effect of transitional probabilities.

4. General discussion

4.1 *Perceptual use of complex phonological patterns*

These two experiments together show, first of all, that even in a language with very complex phonological phenomena acting together to create the surface patterns, listeners can exploit surface phonological regularities for parsing speech. The phonotactic patterns which have previously been shown to influence speech segmentation include particular constraints on what can form a syllable onset or coda (sometimes related to sonority sequencing, sometimes to language-specific patterns such as final devoicing), patterns with a choice between two categories (vowel harmony), and a general constraint against leaving consonants unparsed (the Possible Word Constraint). Experiment 1 of the current work shows that much more complex patterns, caused by a variety of interacting processes, influence parsing just as a single phonotactic pattern does. Experiment 2 shows that even the fact that a particular phonological alternation makes reference to certain types of boundaries is used in parsing. It also shows that this particular alternation, palatalization, affects parsing even when it is isolated from the other alternations investigated in Experiment 1.

Korean differs from the languages in which the influence of phonotactics on segmentation has previously been investigated: in Korean, many phonological alternations affect the surface form, and these alternations are extremely productive in the modern language. Although English, for example, has many phonological alternations, most of them are restricted to a particular lexical stratum, apply only with certain affixes, or may not be

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

entirely active in speakers' synchronic knowledge of the language. In Korean, simply adding a highly productive suffix such as the subject particle, or combining two lexical items into a compound, triggers multiple phonological alternations, such as palatalization, nasal insertion, and nasal assimilation. Not adding suffixes, on the other hand, brings about coda neutralization. Hence, a wide variety of interacting phonological alternations are both productive and frequent.

In a language with such complicated productive phonology, it could be that listeners simply could not make any use of the surface phonological patterns for parsing, because they are so variable and they interact with each other. As shown in (11) above, these interacting phenomena lead to a great deal of potential for confusion in the phonotactics: each possible surface sequence can have more than one underlying sequence, but which underlying sequences are possible depends on the type of intervening boundary. In such a situation, the mapping between surface and underlying patterns might be too complex for listeners to use the surface patterns in parsing the speech. The fact that a pattern is evident in the formal analysis of the phonology of a language does not necessarily mean that listeners can or do use that pattern in online processing. However, our two experiments show that even in a language with complex productive phonology, listeners do make use of whatever information the phonology affords.

4.2 *Perceptual use of non-derived environment blocking*

Among the various phonological alternations involved in the sequences examined here, the palatalization alternation has a special status as an alternation which fails to apply within the morpheme (non-derived environment blocking). Such an alternation has a greater level of phonological abstraction than other alternations, because it requires reference to the presence or absence of particular strengths of boundaries. Among the other alternations relevant to this study, nasal insertion requires reference to word boundaries, since it occurs only word-initially or initial to the second part of a compound. Coda neutralization requires reference to syllable boundaries. But none of the other alternations distinguish between presence and absence of a suffix boundary, applying at a morpheme boundary but not within a morpheme. That is, palatalization is the phonological alternation which requires reference to the morphological environment.

Abstract patterns such as non-derived environment blocking might perhaps be relegated to the phonological analysis that linguists perform, without having any effect on listeners' lexical processing behavior.[16] However, the

[16] John Kingston (personal communication) points out that non-derived environment blocking may not be as abstract for listeners as other types of opacity, such as counterbleeding and counterfeeding. The fact that non-derived environments apply at boundaries

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

two experiments presented here show that even non-derived environment blocking does indeed affect lexical processing, along with other parts of the phonology in Experiment 1, and alone in Experiment 2. Listeners use this aspect of the phonology both in word segmentation and in recognition of morphemes. The blocking in a non-derived environment, rather than the application of palatalization at a boundary, is what provides information: lack of palatalization tells the listener that there cannot be a morpheme boundary, and listeners are faster to rule out a subject particle interpretation when palatalization has not applied.

4.3 *How much phonology do listeners perform online?*

These findings raise the question of how much phonological analysis native speakers and listeners of a language do in order to make use of the information phonological patterns can provide about boundaries. If listeners' parsing of the speech stream is affected by several interacting phonological alternations, to what extent does that mean the listeners are performing the phonological analysis online as they segment speech? Given the complex mapping between underlying forms and surface forms at various levels of boundaries, as shown in (11), do the current results imply that listeners are aware of this mapping and are working backward to underlying forms as they parse speech?

Our experiments do not investigate these questions directly, but we would like to suggest that the results do not argue strongly for online phonological analysis in parsing. Although (11) shows which surface forms can be derived from which underlying forms at which types of boundaries, it is probably not necessary for listeners to calculate which underlying forms could be mapped to the surface string they are hearing in order to use the phonological patterns to find boundaries. Instead, the listener could simply be aware that [ji] occurs only at a suffix boundary or within a morpheme, without keeping track of the fact that it must have an underlying form /ci/ rather than /ti/ if it is within a morpheme, while it can map to either underlying /ci/ or underlying /ti/ if it is at a suffix boundary. The listener would need to know that [di] occurs only at a word or compound boundary, or within a morpheme (that is, not at a suffix boundary), but would not need to know which underlying form [di] could be derived from in which case. That is, in order to use phonotactic patterns as an aid to locating boundaries, the listener only needs to be aware of the surface distribution of sounds relative to boundaries, not of what in the underlying representations or phonological alternations created those surface distributions.

---

but not morpheme-internally may make them particularly useful to listeners, as we have found.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Of course, in order to perform spoken word recognition, the listener must eventually recover the lexical representation of the lexical item, whatever it might be. If each lexical item is stored as a single invariant lexical representation, rather than having all its allomorphs stored, then the listener will eventually recover the underlying form. However, that step is not necessary in order to use phonotactic information to locate boundaries. Furthermore, if all allomorphs are stored, then it is not necessary for the listener to 'undo' the phonology at all. Our experiments do not address the question of whether lexical representations include only a single form or multiple forms. For purposes of segmentation, we think it is quite possible that listeners are aware of the distribution of surface patterns without making reference to underlying patterns or to the phonology that connects the two levels.

4.4 *Usefulness of boundaries other than syllables for segmentation*

Previous research on phonotactic information and speech segmentation has largely involved information about either syllable boundaries or utterance boundaries. The research by McQueen (1998) on Dutch, Weber (2001) on English, and Dumay et al. (2002) on French shows that listeners can use the fact that sequences such as /mr/ cannot be syllable onsets, while sequences like /dr/ either can be onsets (in English) or must be onsets (in Dutch). McQueen proposes that listeners use whatever information is available to help them locate syllable boundaries, and that they use syllable boundaries to help locate word boundaries, since words even in connected speech very often begin at syllable boundaries. The findings on Finnish vowel harmony and parsing (Suomi et al. 1997) show that listeners also make use of information about words rather than syllables, but vowel harmony is a different type of phonological alternation from those investigated here.

The current experiments show that listeners make use of information about the distribution of surface segments relative to many types of boundaries: morpheme-internal position, the boundary between a root and a suffix, and the boundary between parts of a compound or between words. As noted above, most of the relevant phonological patterns do not involve syllable structure. This suggests that listeners can use phonotactic information to help them with segmentation directly, not only as a means to locate syllable boundaries. It appears that Korean listeners are sensitive to which surface sound sequences occur and do not occur at which types of boundaries, and can use this information in segmentation whether the phonotactic patterns provide any information about syllable boundaries or not. Thus, the current findings extend the previous work not only by testing a more complicated phonological situation, but also by testing phonotactic patterns which involve boundaries at levels other than the syllable.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

4.5 *Word vs. compound boundaries*

In (11) above, we did not distinguish the boundary between words of a sentence from the boundary between the two parts of a compound. For example, we refer to [di] as occurring at word or compound boundaries, with no attempt to distinguish these two types of boundaries. This decision is in part due to the difficulty of determining exactly what constitutes a compound vs. a string of two words, and what the exact environment of nasal insertion is (section 1.2 point C above; see Cho (1998) on the complications of this alternation). Furthermore, compounds vary in the strength of the association between the two parts. Korean has many compounds in which the two morphemes were borrowed from Chinese, and many native speakers of Korean may not be aware that some of these are in fact etymologically compounds. Many other compounds are formed quite productively, though.

The choice to collapse word and compound boundaries is also partly based on the task we used: word-spotting requires the listener to recognize a real word with a small quantity of non-word material next to it. In our Experiment 1, the nonsense material consisted of either one or two syllables and an onset consonant, placed before the real word (e.g. /imin/ 'immigration' embedded in /p^{h}yoj**imin**/). To listeners, such items could easily sound like constructed compounds, with the first member of the compound a non-word and the second a real word. In fact, since there are no grammatical particles between the nonsense and the real word, the stimuli may sound more like a compound than like two independent words concatenated at the syntactic level. /p^{h}yoj**imin**/ may sound more like a compound with /imin/ as the second part of it than like a phrase or sentence containing the independent word /imin/. Therefore, in Experiment 1, it is reasonable to assume that phonotactic information about compound boundaries would be just as helpful to listeners recognizing the embedded real word as information specifically about word boundaries would be. Since the task may not distinguish between detection of part of a productively formed compound and detection of an independent word, we do not distinguish between phonotactic information about these two types of boundaries.

5. Conclusions

We have shown that listeners of a language with a very complex and yet very productive phonology can make use of phonotactic information to help them segment the speech stream into words. Our results also demonstrate that abstract aspects of the phonological system of a language, such as non-derived environment blocking, affect speech parsing. Listeners' use of such information fits in with a model of spoken word recognition such as Shortlist (Norris et al. 1997), in which listeners use a great many types of information, both language-specific and more general, in the process of parsing continuous speech into words.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

APPENDIX 1

Full list of experimental test items for Experiment 1

(Items which were excluded from the items analysis because of poor recognition are marked with an asterisk.)

REAL WORD	TRANSCRIP-TION	GLOSS	STIMULI	TRANSCRIPTION
이제	[ijɛ]	now	효지제	[hyojijɛ]
			효디제	[hyodijɛ]
			횬니제	[hyonnijɛ]
이념	[inyʌm]	ideology	카지념	[k^{h}ajinyʌm]
			카디념	[k^{h}adinyʌm]
			칸니념	[k^{h}anninyʌm]
이삭	[isak]	ears of grain	커지삭	[k^{h}ʌjisak]
			커디삭	[k^{h}ʌdisak]
			컨니삭	[k^{h}ʌnnisak]
이온	[ion]	ion	히지온	[hijion]
			히디온	[hidion]
			힌니온	[hinnion]
이직	[ijik]	loss of employment	쿄드지직	[k^{h}yodɨjijik]
			쿄드디직	[k^{h}yodɨdijik]
			쿄든니직	[k^{h}yodɨnnijik]
이치	*[ic^{h}i]	(good) reason	느뮈지치	[nɨmwijic^{h}i]
			느뮈디치	[nɨmwidic^{h}i]
			느뮌니치	[nɨmwinnic^{h}i]
이행	[ihɛŋ]	execution	기커지행	[kik^{h}ʌjihɛŋ]
			기커디행	[kik^{h}ʌdihɛŋ]
			기컨니행	[kik^{h}ʌnnihɛŋ]
이튿날	[it^{h}ɨnnal]	the day after	르지튿날	[lɨjit^{h}ɨnnal]
			르디튿날	[lɨdit^{h}ɨnnal]
			른니튿날	[lɨnnit^{h}ɨnnal]
이다지	*[idaji]	to this extent	므지다지	[mɨjidaji]
			므디다지	[mɨdidaji]
			믄니다지	[mɨnnidaji]

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Appendix 1 (cont.)

REAL WORD	TRANSCRIP-TION	GLOSS	STIMULI	TRANSCRIPTION
이슬람	[isɨllam]	Islam	누아지슬람	[nuajisɨllam]
			누아디슬람	[nuadisɨllam]
			누안니슬람	[nuannisɨllam]
이른바	*[iɾɨnba]	so-called	새쿠치른바	[sɛkʰucʰiɾɨnba]
			새쿠티른바	[sɛkʰutʰiɾɨnba]
			새쿤니른바	[sɛkʰunniɾɨnba]
이만	[iman]	this far	니지만	[nijiman]
			니디만	[nidiman]
			닌니만	[ninniman]
이쪽	*[ic'ok]	this side	프지쪽	[pʰɨjic'ok]
			프디쪽	[pʰɨdic'ok]
			픈니쪽	[pʰɨnnic'ok]
이별	[ibyʌl]	separation	버쇠치별	[pʌswɛcʰibyʌl]
			버쇠티별	[pʌswɛtʰibyʌl]
			버쇤니별	[pʌswɛnnibyʌl]
이빨	[ip'al]	tooth (teeth)	퍼지빨	[pʌjip'al]
			퍼디빨	[pʌdip'al]
			펀니빨	[pʌnnip'al]
이유	[iyu]	cause	츠지유	[cʰɨjiyu]
			츠디유	[cʰɨdiyu]
			츤니유	[cʰɨnniyu]
이런	[iɾʌn]	my goodness!	으지런	[ɨjiɾʌn]
			으디런	[ɨdiɾʌn]
			은니런	[ɨniɾʌn]
이의	*[iɨi]	objection	퍄지의	[pʰyajiɨi]
			퍄디의	[pʰyadiɨi]
			퍈니의	[pʰyanniɨi]
이익	*[iik]	profit	며브지익	[myʌbɨjiik]
			며브디익	[myʌbɨdiik]
			며븐니익	[myʌbɨnniik]
이율	*[iyul]	interest rate	만푸지율	[manpʰujiyul]
			만푸디율	[manpʰudiyul]
			만푼니율	[manpʰunniyul]

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Appendix 1 (cont.)

REAL WORD	TRANSCRIP-TION	GLOSS	STIMULI	TRANSCRIPTION
익살	[iksal]	witticism	고니칙살	[konic^{h}iksal]
			고니틱살	[konit^{h}iksal]
			고닌닉살	[koninniksal]
이탤릭	[it^{h}ɛllik]	italic	며지탤릭	[myʌjit^{h}ɛllik]
			며디탤릭	[myʌdit^{h}ɛllik]
			면니탤릭	[myʌnnit^{h}ɛllik]
이브닝	[ibɨniŋ]	evening	먀지브닝	[myajibɨniŋ]
			먀디브닝	[myadibɨniŋ]
			먄니브닝	[myannibɨniŋ]
이번에	[ibʌnɛ]	this time	도프치번에	[top^{h}ɨc^{h}ibʌnɛ]
			도프티번에	[top^{h}ɨt^{h}ibʌnɛ]
			도픈니번에	[top^{h}ɨnnibʌnɛ]
이중	[ijuŋ]	duplication	뮈지중	[mwijijuŋ]
			뮈디중	[mwidijuŋ]
			뭔니중	[mwinnijuŋ]
이끼	[ik'i]	moss	괴지끼	[kwɛjik'i]
			괴디끼	[kwɛdik'i]
			괸니끼	[kwɛnnik'i]
이탈	[it^{h}al]	secession	브지탈	[pɨjit^{h}al]
			브디탈	[pɨdit^{h}al]
			븐니탈	[pɨnnit^{h}al]
이혼	[ihon]	divorce	투러지혼	[t^{h}urʌjihon]
			투러디혼	[t^{h}urʌdihon]
			투런니혼	[t^{h}urʌnnihon]
이승	[isɨŋ]	this world	캬지승	[k^{h}yajisɨŋ]
			캬디승	[k^{h}yadisɨŋ]
			캰니승	[k^{h}yannisɨŋ]
이륙	[iɾyuk]	landing	쇠지륙	[swɛjiɾyuk]
			쇠디륙	[swɛdiɾyuk]
			쇤니륙	[swɛniɾyuk]
이미	[imi]	already	과지미	[kwajimi]
			과디미	[kwadimi]
			관니미	[kwannimi]

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Appendix 1 (cont.)

REAL WORD	TRANSCRIP-TION	GLOSS	STIMULI	TRANSCRIPTION
이민	[imin]	immigration	뵤지민	[pyojimin]
			뵤디민	[pyodimin]
			뵨니민	[pyonnimin]
이견	[igyʌn]	different view	부너치견	[punʌcʰigyʌn]
			부너티견	[punʌtʰigyʌn]
			부넌니견	[punʌnnigyʌn]
이웃	[iut]	neighbor	티르지웃	[tʰirɨjiut]
			티르디웃	[tʰirɨdiut]
			티른니웃	[tʰirɨnniut]
이스트	[isɨtʰɨ]	yeast	위지스트	[wijisɨtʰɨ]
			위디스트	[widisɨtʰɨ]
			윈니스트	[winnisɨtʰɨ]
이윽고	[iɨkko]	after a while	눠지윽고	[nwʌjiɨkko]
			눠디윽고	[nwʌdiɨkko]
			눤니윽고	[nwʌnniɨkko]
이렇게	[iɾʌkʰɛ]	in this way	퍼노지렇게	[pʰʌnojiɾʌkʰɛ]
			퍼노디렇게	[pʰʌnodiɾʌkʰɛ]
			퍼논니렇게	[pʰʌnonniɾʌkʰɛ]
이라크	[iɾakʰɨ]	Iraq	마호치라크	[mahocʰiɾakʰɨ]
			마호티라크	[mahotʰiɾakʰɨ]
			마혼니라크	[mahonniɾakʰɨ]
이득	[idɨk]	gains	라지득	[lajidɨk]
			라디득	[ladidɨk]
			란니득	[lannidɨk]
이용	[iyoŋ]	utilization	로지용	[lojiyoŋ]
			로디용	[lodiyoŋ]
			론니용	[lonniyoŋ]
이북	[ibuk]	North Korea	냐지북	[nyajibuk]
			냐디북	[nyadibuk]
			냔니북	[nyannibuk]
이동	[idoŋ]	movement	저드치동	[jʌdɨcʰidoŋ]
			저드티동	[jʌdɨtʰidoŋ]
			저든니동	[jʌdɨnnidoŋ]

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

APPENDIX 2

Full list of experimental test items for Experiment 2

REAL WORD [di]		REAL WORD [ji]	
어디	[ʌdi] where	의지	[ɨiji] intention
마디	[madi] joint	미지	[miji] unknown
부디	[pudi] without fail	부지	[puji] site
귀티	[kwit^{h}i] an air of nobility	마치	[mac^{h}i] as if
부티	[put^{h}i] an air of wealth	배치	[pɛc^{h}i] arrangement
본디	[pondi] by nature	토지	[t^{h}oji] land
불티	[pult^{h}i] sparks	노다지	[nodaji] bonanza
금잔디	[kɨmjandi] golden [autumnal] turf	재배지	[jɛbɛji] cultivation place
피디	[p^{h}idi] P.D. producer	표지	[p^{h}yoji] (book) cover
씨디	[s'idi] C.D. compact disc	사치	[sac^{h}i] luxury
바디	[padi] body	바가지	[pagaji] nagging
캔디	[k^{h}ɛndi] candy	거지	[kʌji] mendicant
간디	[kandi] Gandhi	처지	[c^{h}ʌji] condition
팬티	[p^{h}ɛnt^{h}i] panties	도라지	[doraji] broad bellflowers
코매디	[k^{h}omɛdi] comedy	에너지	[ɛnʌji] energy
레이디	[lɛidi] lady	메시지	[mɛsiji] message
리버티	[libʌt^{h}i] liberty	스위치	[sɨwic^{h}i] switch
게런티	[kɛɾʌnt^{h}i] guarantee	스케치	[sɨk^{h}ɛc^{h}i] sketch
로얄티	[loyalt^{h}i] loyalty	쏘세지	[s'oseji] sausage

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Appendix 2 (cont.)

NON-WORDS (each with both [di] and [ji])			
과지/과디	[kwaji kwadi]	괴지/괴디	[kwɛji kwɛdi]
니지/니디	[niji nidi]	냐지/냐디	[nyaji nyadi]
뵤지/뵤디	[pyoji pyodi]	눠지/눠디	[nwʌji nwʌdi]
먀지/먀디	[myaji myadi]	므지/므디	[mɨji mɨdi]
브지/브디	[pɨji pɨdi]	궈지/궈디	[kwʌji kwʌdi]
뮤지/뮤디	[myuji myudi]	로지/로디	[loji lodi]
르지/르디	[lɨji lɨdi]	위지/위디	[wiji widi]
쇠지/쇠디	[swɛji swɛdi]	츠지/츠디	[cʰɨji cʰɨdi]
캬지/캬디	[kʰyaji kʰyadi]	텨지/텨디	[thyʌji thyʌdi]
퍄지/퍄디	[pʰyaji pʰyadi]	퍼지/퍼디	[pʰʌji pʰʌdi]
히지/히디	[hiji hidi]	효지/효디	[hyoji hyodi]
고니치/고니티	[konicʰi konitʰi]	기커지/기커디	[kikʰʌji kikʰʌdi]
나개지/나개디	[nagɛji nagɛdi]	누아지/누아디	[nuaji nuadi]
느뮈지/느뮈디	[nɨmwiji nɨmwidi]	도프치/도프티	[tophucʰi tophutʰi]
만푸지/만푸디	[manpʰuji manpʰudi]	마호치/마호티	[mahocʰi mahotʰi]
버쇠치/버쇠티	[peswɛcʰi peswɛtʰi]	새쿠치/새쿠티	[sɛkʰucʰi sɛkʰutʰi]
저드치/저드티	[cʌtɨcʰi cʌtɨtʰi]	카츠지/카츠디	[kʰacʲɨji kʰacʰɨdi]
쿄드지/쿄드디	[kʰyodɨji kʰyodɨdi]	투러지/투러디	[tʰuɾʌji tʰuɾʌdi]
티르지/티르디	[tʰiɾɨji tʰiɾɨdi]	퍼노지/퍼노디	[pʰʌnoji pʰʌnodi]

REFERENCES

Cho, Y.-M. (1998). Language change as reranking of constraints. In Hogg, R. M. & van Bergen, L. (eds.), *Historical linguistics 1995: selected papers from the 12th International Conference on Historical Linguistics, Manchester, August 1995* (vol. 2). Amsterdam: John Benjamins. 45–62.

Cho, Y.-M. & Sells, P. (1995). A lexical account of inflectional suffixes in Korean. *Journal of East Asian Linguistics* **4**. 119–174.

Connine, C. & Titone, D. (1996). Phoneme monitoring. *Language and Cognitive Processes* **11**. 635–645.

Cutler, A., Demuth, C. & McQueen, J. (2002). Universality versus language-specificity in listening to running speech. *Psychological Science* **13**. 258–262.

Dumay, N., Frauenfelder, U. & Content, A. (2002). The role of the syllable in lexical segmentation in French: word-spotting data. *Brain and Language* **81**. 144–161.

Forster, K. & Forster, J. (2003). DMDX: a windows display program with millisecond accuracy. *Behavioral Research Methods, Instruments & Computers* **35**. 116–124.

Frauenfelder, U. & Kearns, R. (1996). Sequence monitoring. *Language and Cognitive Processes* **11**. 665–673.

McQueen, J. (1996). Word spotting. *Language and Cognitive Processes* **11**. 695–699.

McQueen, J. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language* **39**. 21–46.

McQueen, J., Otake, T. & Cutler, A. (2001). Rhythmic cues and possible-word constraints in Japanese speech segmentation. *Journal of Memory and Language* **45**. 103–132.

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press

Mehler, J., Dommergues, J., Frauenfelder, U. & Segui, J. (1981). The syllable's role in speech segmentation. *Journal of Verbal Learning and Verbal Behavior* **20**. 298–305.
Norris, D., McQueen, J., Cutler, A. & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology* **34**. 191–243.
Norris, D., McQueen, J., Cutler, A., Butterfield, S. & Kearns, R. (2001). Language-universal constraints on speech segmentation. *Language and Cognitive Processes* **16**. 637–660.
Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin* **114**. 510–532.
Rogers, M., You, C. & Richards, K. (1992). *College Korean*. Berkeley: University of California Press.
Suomi, K., McQueen, J. & Cutler, A. (1997). Vowel harmony and speech segmentation in Finnish. *Journal of Memory and Language* **36**. 422–444.
van der Lugt, A. (2001). The use of sequential probabilities in the segmentation of speech. *Perception and Psychophysics* **63**. 811–823.
Vitevitch, M. & Luce, P. (1999). Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory and Language* **40**. 374–408.
Weber, A. (2001). *Language-specific listening: the case of phonetic sequences*. Ph.D. dissertation, University of Nijmegen. [Max Planck Series in Psycholinguistics **16**.]

Authors' address: (*Warner*)
Department of Linguistics, University of Arizona,
PO Box 210028, Tucson, AZ 85721-0028, U.S.A.
E-mail: nwarner@u.arizona.edu

https://doi.org/10.1017/S0022226705003294 Published online by Cambridge University Press