

Language-universal constraints on speech segmentation

Dennis Norris

MRC Cognition and Brain Sciences Unit, Cambridge, UK

James M. McQueen, and Anne Cutler

Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands

Sally Butterfield and Ruth Kearns

MRC Cognition and Brain Sciences Unit, Cambridge, UK

Two word-spotting experiments are reported that examine whether the Possible-Word Constraint (PWC) is a language-specific or language-universal strategy for the segmentation of continuous speech. The PWC disfavors parses which leave an impossible residue between the end of a candidate word and any likely location of a word boundary, as cued in the speech signal. The experiments examined cases where the residue was either a CVC syllable with a schwa, or a CV syllable with a lax vowel. Although neither of these syllable contexts is a possible lexical word in English, word-spotting in both contexts was easier than in a context consisting of a single consonant. Two control lexical-decision experiments showed that the word-spotting results reflected the relative segmentation difficulty of the words in different contexts. The PWC appears to be language-universal rather than language-specific.

The segmentation of a written text such as this one into its component words is a trivial task for the reader, because the writers have helpfully left empty spaces between the individual words. Speakers do not help listeners

Requests for reprints should be addressed to Dennis Norris, MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge, CB2 2EF, UK.

E-mail: dennis.norris@mrc-cbu.cam.ac.uk

Ruth Kearns is now with Procter & Gamble Technical Centres Ltd. We would like to thank Maarten Jansonius and especially Keren Shatzman for their help with the preparation and running of Experiments 2 and 4, and Antje Meyer for giving us the opportunity to conduct these experiments at the University of Birmingham.

in this way. Spoken utterances are continuous, and there are no completely reliable cues to word boundaries (Lehiste, 1972; Nakatani & Dukes, 1977). Therefore one of the tasks which the listener has to accomplish, in order to understand what the speaker is trying to say, is segmentation: dividing the continuous signal into its constituent words.

There are two main classes of solution to the problem of segmenting continuous speech. The first, and most general, solution is that adopted by the connectionist models TRACE (McClelland & Elman, 1986) and Shortlist (Norris, 1994). In these models, all possible lexical candidates consistent with any sequence of phonemes in the input compete, so that overlapping candidates inhibit one another. The competition process allows these networks to parse the input into a sequence of non-overlapping words.

The second strategy is to adopt a more explicit segmentation procedure in an attempt to capitalise on the available phonetic or phonological segmentation cues in the input. One such strategy that has received considerable empirical support is the Metrical Segmentation Strategy (MSS) of Cutler and Norris (1988). As applied to English, the MSS is designed to take advantage of the fact that the majority of content words in English begin at the onsets of strong syllables (Cutler & Carter, 1987). However, the MSS does not operate in the same way for all languages. The strategy is sensitive to the rhythmic structure of the listener's native language. In English the rhythmic unit is the foot, in French the syllable, and in Japanese the mora. There is now extensive evidence that listeners do indeed draw on their knowledge about the rhythmic structure of their native language in segmenting speech. This produces language-specific effects in segmentation (Cutler, Mehler, Norris, & Seguí, 1986, 1992; Cutler & Otake, 1994; Otake, Hatano, Cutler, & Mehler, 1993; McQueen, Otake, & Cutler, 2001; Sebastián-Gallés, Dupoux, Seguí, & Mehler, 1992). The rhythmic structure of language helps segmentation in the native language, for instance, but can lead to inappropriate listening strategies when the input is in a non-native language which has a different rhythm (see Cutler, Dahan, & Donselaar, 1997, for a review). A similar story can be told for phonotactic constraints. Phonotactic sequence constraints can be effectively exploited to segment the native language (McQueen, 1998), but again can be misleading when the input is in a non-native language in which the constraints are different (Weber, 2000).

One very powerful weapon in the listener's armoury was discovered by Norris, McQueen, Cutler, and Butterfield (1997). This is a constraint—the Possible-Word Constraint (PWC)—which disfavors interpretations which would leave a residue of the input which could not itself be exhaustively parsed into one or more words. The initial evidence for the PWC came from an experiment in which listeners were required to spot real words in

short nonsense strings. Norris et al. found that words were harder to spot when the residue of the nonsense string was only a single consonant than when the residue was a syllable. Thus *sea* was harder to spot in *seash* than in *seashub*, and *apple* was harder to spot in *fapple* than in *vuffapple*. None of the residues—*sh*, *shub*, *f*, *vuff*—are in fact words of English. But *vuff* and *shub*, although they happen not to be members of the English vocabulary, satisfy all the phonological preconditions for membership and hence might well have been words. The two single consonants *sh* and *f*, on the other hand, could never themselves be viable candidate words. Norris et al. proposed that this constraint could provide a powerful method for inhibiting activation of words which are spuriously present in an utterance, and tested this by implementing the PWC in the computational model of spoken-word recognition Shortlist (Norris, 1994). In their implementation, the PWC works by penalising any candidate word which begins or ends at a point where there is no vowel between that point and a “known” boundary. Known boundaries are those locations in the speech signal where there is likely to be a word boundary, as determined by metrical information (Cutler & Norris, 1988), phonotactic information (McQueen, 1998) or silence (Norris et al., 1997).

Norris et al. (1997) presented simulations showing that Shortlist’s ability to parse the input improved when the PWC was incorporated into the model. Furthermore, with the PWC, Shortlist was able to simulate a range of data including the metrical segmentation effects reported in word-spotting by Cutler and Norris (1988), by McQueen, Norris, and Cutler (1994), and by Norris, McQueen, and Cutler (1995) and in identity priming by Vroomen and de Gelder (1995).

An illustration of the potential value of the PWC can be seen by considering the utterance “they met a fourth time”. This contains possible spurious occurrences of, for instance, *aim* (**they met**), *for* (**fourth**), *I’m* (**time**) and *metaphor* (**met a fourth**). But if these words were activated, each could be rejected by the PWC on the grounds that a single-consonant residue would inevitably be left unaccounted for. Thus *aim* leaves the initial sound of *they*, *I’m* leaves the initial *t* of *time*, and *for* and *metaphor* leave the final *th* of *fourth*; none of these residues could be parsed into words.

The PWC provides an integrated account of the use of both rhythmic and phonotactic cues. Rhythmic and phonotactic information, along with silence, all provide cues to the possible location of word boundaries. Furthermore, the PWC can make more effective use of these cues than theories which postulate that the input is explicitly segmented at possible segmentation boundaries. In contrast to the initial formulation of the MSS, or theories claiming that the syllable is a unit for lexical access (Mehler, 1981; Mehler, Dommergues, Frauenfelder, & Seguí, 1981; Seguí, 1984), the

segmentation process does not need to begin by assuming that there actually are word boundaries at all of these locations. For example, the second syllable of the word *delightful* is a strong syllable, and there is a phonotactically determined syllable boundary between the second and third syllables. A mechanism that attempted to force word boundaries at these points might degrade word recognition by disrupting the processing of the intended word. However, the PWC operates only by eliminating word candidates which are misaligned with these boundaries. For instance, it penalises *lie* in *delightful* because a *t* would be unaccounted for between *lie* and the phonotactic boundary between the end of *delight* and the beginning of *ful*. But the PWC will not hinder recognition of aligned words, even where they have word-internal boundaries.

Although rhythmic and phonotactic information can be used in all languages, the precise implementation of these strategies is determined by the specific properties of individual languages. The PWC operates on the output of those language-specific segmentation procedures. As currently implemented, however, the PWC itself operates in exactly the same way for all languages. The PWC simply determines whether there is a vowel between the end of a word candidate and the nearest boundary. In other words, it makes use of the language-universal constraint that all words must have vowels. In nearly all languages of the world it is the case that a single consonant cannot form a lexical word. The current implementation of the PWC therefore reflects the fact that every component in the lexical parse of a stretch of continuous speech must contain at least one vowel. A given candidate word is only a plausible part of that parse if the residue of speech between the word's edge and a likely word boundary contains a vowel.

However, languages differ not only in their rhythmic and phonotactic characteristics, but also in what counts as an acceptable minimal word. In English, although function words can consist of a single reduced syllable (reduced forms of *the*, *a*, *to*, etc.), content words cannot. If the PWC is sensitive to this phonological property of English words, listeners should find it as hard to spot words embedded in weak-syllable contexts as in consonant contexts. If the PWC is a simple language-universal constraint, however, word-spotting should be much easier when the residue constitutes a weak syllable than when it is a consonant.

Furthermore, no English word can consist of an open syllable with a short full vowel. The word *se*, for example, with the vowel of *sell*, is not a possible word. There are six such lax vowels in English (those in the words *hat*, *bet*, *hit*, *book*, *hot*, and *hut*). Open syllables with long vowels are acceptable (e.g., *sea*) and closed syllables with short vowels (e.g., *sell*) are also fine. However, although *se* is not a possible English word, it would be a perfectly good word in French. French allows open syllables with lax

vowels, as evidenced by words such as *thé* and *va*. In yet other languages, however, monosyllabic words of any kind are not well-formed words. In the Australian language Lardil and in the Bantu language Sesotho, for example, lexical words must have at least two syllables. The PWC could therefore be a language-specific constraint, one that reflects this linguistic variation. It is possible that the PWC is sensitive to the fact that in English, for example, an open syllable with a lax vowel is not a well-formed English word. If so, we would expect that if there was only a syllable like /sɛ/ as a residue between a candidate word and a likely word boundary in an ongoing parse of speech, that candidate would not be a plausible part of the parse.

The four experiments we report here address the question of the universality versus language-specificity of the PWC. The two critical experiments use the word-spotting paradigm, while two lexical decision experiments provide control data. If the PWC is language-universal, then a word will be hard to spot only if it is embedded in a context which could not be a word in any language. If the PWC is language-specific, however, then only contexts which could not be a word in the listeners' native language will make spotting embedded words difficult (even though the same contextual residues might be acceptable words in other languages). Both word-spotting experiments examine whether residues which contain vowels, but are not phonologically well-formed content words in English, behave like the syllables with full vowels studied by Norris et al. (1997), or whether they behave like non-viable consonant residues. The first word-spotting experiment examines syllables with the reduced vowel schwa; the second examines open syllables with lax vowels. Each of these experiments has a companion control experiment in which listeners were required to make lexical decisions to the words excised from the word-spotting stimuli. These control experiments enable us to determine whether any effects observed in the word-spotting experiments might be attributable to differences in the phonetic realisation of the stimuli in different contexts.

EXPERIMENT 1

If the PWC really is determined by a universal rather than language-specific notion of possible word, then we should expect word spotting to be easier for words in weak-syllable contexts than in consonant contexts. Weak-syllable contexts should behave just like the full syllable contexts in Norris et al. (1997). That is, spotting *sea* in /sɪfəb/ should be easier than spotting *sea* in /sɪf/. But if the PWC is in some fashion language-specific, then English listeners should not treat syllables with weak vowels the same way as those with full vowels.

Method

Subjects. Twenty-four native speakers of English from the MRC Cognition and Brain Sciences Unit volunteer panel were paid for their participation.

Materials. The stimuli were derived from the following-context materials in Norris et al. (1997) by changing the vowels in their full-vowel syllabic contexts to schwa. So, for example, the target word *sea* could appear with either a following consonant context (e.g., /sɪf/) or a following weak syllable context (e.g., /sɪfəb/). In the case of the following syllable contexts 11 of the 48 items retained exactly the same consonants (C*C) as in Norris et al. The remaining items were altered to avoid creating phonotactically illegal strings or strings that could be misheard as words, and to increase the variety of contexts. The experimental items are listed in Appendix A. There were 110 filler items; many of these had weak final syllables so that a final weak syllable was not a cue to the presence of an embedded word. There were also eight filler target words with following full syllables. As in Norris et al., half the target words were monosyllabic and half were bisyllabic. Target words only appeared with following contexts. Two experimental lists were constructed, each containing all the fillers, and all the targets, but with type of context counterbalanced over lists such that half the words on each list were in consonant contexts and half in weak syllable contexts.

Procedure. The materials were recorded onto Digital Audio Tape by a native speaker of English in a sound-attenuated booth. The stimuli were transferred to a computer and converted to stereo files where the speech was on one channel and inaudible timing marks were on the other. The two lists of stimuli were then transferred to a compact disc.

Listeners were tested individually in a quiet room; they heard the lists over headphones driven by a portable compact disc player. They were told they would hear nonsense words, some of which would contain real English words. They were asked to press a button as fast as possible whenever they spotted a real word, and to say aloud the word that they had spotted. Reaction Times (RTs) were measured from the offset of target-bearing items. Each listener heard a practice list, followed by one of the two experimental lists.

Results and discussion

Analyses of variance (ANOVAs) were performed on the latency and accuracy data. Four words were excluded from the analysis because they were missed by more than two-thirds of the subjects who heard them in

either consonant or syllable contexts: *eager*, *tell*, *chain*, and *bother*. The results are summarised in Table 1.

The main effect of context was significant by subjects in the RT analysis ($F_1(1, 22) = 5.71, p < .03$; $F_2(1, 40) = 2.07, p = .16$) and by both subjects and items in the error analysis ($F_1(1, 22) = 16.21, p < .001$; $F_2(1, 40) = 17.20, p < .001$). The effect of number of syllables was significant in the RT analysis ($F_1(1, 22) = 9.35, p < .01$; $F_2(1, 40) = 5.72, p < .05$) but not in the error analysis ($F_s < 1$). There was also a significant interaction between context and number of syllables in the error analysis ($F_1(1, 22) = 18.24, p < .001$; $F_2(1, 40) = 11.27, p < .002$) but not in the RT analysis ($F_s < 1$). Although the error rates are lower than in Norris et al., the overall pattern, including the fact that the context effect in errors was larger in bisyllables, is very similar to the corresponding consonant and full-syllable conditions in their Experiment 1. In that experiment the overall context effect was 45 ms in RTs and 15% in errors, compared with 94 ms and 13% here.

The results of Experiment 1 are very straightforward: word-spotting is easier in weak syllable contexts than in consonant contexts. Furthermore, this difference is, if anything, marginally greater than the difference between the corresponding full-syllable and consonant contexts in Norris et al. (1997). There is therefore no suggestion that weak syllables violate the PWC. This suggests that the PWC operates according to language-universal principles.

EXPERIMENT 2

The differences observed in Experiment 1 involve comparisons between different versions of the same target words produced in different contexts.

TABLE 1

Mean reaction times (RTs, in ms, measured from target offset) and mean percentage error rates, in word spotting (Experiment 1) and in lexical decision to the same target words, Excised from their contexts (Experiment 2)

Target:	Monosyllabic		Bisyllabic	
	CəC	C	CəC	C
Word spotting				
Mean RT	890	1001	789	866
Mean error	19%	22%	9%	32%
Example	seashəb	seash	sugarməl	sugarm
Lexical decision				
Mean RT	630	477	436	409
Mean error	14%	8%	4%	4%
Example	sea(shəb)	sea(sh)	sugar(məl)	sugar(m)

There is a possibility that the results could therefore reflect differences between the target words spoken in weak-syllable contexts and those spoken in consonant contexts, rather than an effect of the contexts per se. Experiment 2 was a control lexical decision experiment which addressed this concern. If the effect observed in Experiment 1 were due to the relative difficulty listeners had in segmenting the target words from their contexts, then there should be no difference in lexical decision performance on these target words when they are excised from their contexts and presented in isolation.

Method

Subjects. Twenty-four native speakers of English, students at the University of Birmingham, were paid for their participation.

Materials. Forty-four of the 48 targets words used in Experiment 1 were excised from their contexts using the Xwaves speech editor. Cuts were made at zero-crossings at the offsets of the final phonemes of the words. The word *eager* was not included since it had been excluded from the analysis of Experiment 1. Due to an error, the targets *usher*, *fool*, and *fun* were also excluded, instead of the targets *tell*, *chain*, and *bother* (which had also been excluded from Experiment 1). There were therefore 22 monosyllabic words and 22 bisyllabic words. These were counterbalanced over two lists, such that half the words on each list had been taken from consonant contexts and half from weak-syllable contexts. The words were presented in the same pseudorandom order in each list, mixed with 40 nonwords (20 monosyllabic and 20 bisyllabic), which had been excised from the Experiment 1 fillers.

Procedure. Separate speech files were prepared for each item, and transferred to a portable computer running NESU experimental control software. These files were played over headphones directly from the computer's hard disk. Listeners were tested individually in a quiet room. They were asked to press a button as fast as possible whenever they heard a real English word, and (by analogy to the word-spotting procedure) to then say aloud the word they had heard. RTs were again measured from word offset. Each listener heard a short practice list of words and nonwords (also excised from Experiment 1 materials) and one of the two experimental lists.

Results and discussion

Two monosyllabic and two bisyllabic words were not included in the analysis (*tell*, *chain*, and *bother*, because they had been excluded in

Experiment 1, and *nether*, which was missed by two-thirds of the subjects who heard the version extracted from its consonant context). The mean latency and accuracy values for the remaining 40 items are given in Table 1.

The main effect of context in RTs was significant by subjects and items (words taken from consonant contexts were detected, on average, 90 ms faster than words taken from weak-syllable contexts; $F_1(1, 22) = 97.64$, $p < .001$; $F_2(1, 38) = 30.80$, $p < .001$), as was the main effect of word length (monosyllables were detected, on average, 131 ms more slowly than bisyllables; $F_1(1, 22) = 88.10$, $p < .001$; $F_2(1, 38) = 17.82$, $p < .001$). These two effects interacted ($F_1(1, 22) = 39.62$, $p < .001$; $F_2(1, 38) = 14.08$, $p < .001$); while the context effect was larger for monosyllabic than for bisyllabic words, pairwise comparisons showed that it was reliable in each case. The only effect which was significant in the accuracy analysis was the effect of target length: monosyllabic words were missed more often than bisyllabic words ($F_1(1, 22) = 17.21$, $p < .001$; $F_2(1, 38) = 4.80$, $p < .05$).

Words excised from consonant contexts (like *sea* from /sɪf/) were detected more rapidly than words excised from weak-syllable contexts (like *sea* from /sɪfəb/). This is the reverse of the pattern observed in word spotting. The phonetic realization of the words in consonant contexts (particularly the monosyllables) was therefore better than that of the words in weak-syllable contexts. In spite of this difference in the quality of the words themselves, the targets in consonant contexts were harder to spot in Experiment 1. The difference in phonetic realisation of the words in the two contexts could therefore only have acted to weaken any effect of segmentation difficulty. Indeed, in a by-item Analysis of Covariance (ANCOVA) on the word-spotting RTs from Experiment 1 taking the lexical decision RTs as a covariate, the effect of context (which was not reliable in the original ANOVA) became significant ($F_2(1, 37) = 8.21$, $p < .01$). In a parallel analysis of the error data, the context effect in word spotting remained significant ($F_2(1, 37) = 13.18$, $p < .005$). Experiment 2 therefore confirmed that the advantage for words in weak-syllable contexts over those in consonant contexts in Experiment 1 reflects the relative segmentation difficulty of the words in these contexts.

EXPERIMENT 3

Experiment 1 revealed that weak-syllable contexts behave just like full-syllable contexts, despite the fact that weak syllables are not well-formed content words in English. This suggests that the PWC is language-universal rather than language-specific. Weak syllables, however, are

well-formed function words. It therefore remains possible that the PWC is language specific, but that it is only concerned with whether the context could be any possible word in that language, irrespective of its status as a content or function word. A stronger test would be to determine if the PWC is sensitive to whether or not the context is a syllable which could not possibly be a word of any sort in the language. We can test this in English by using contexts consisting of an open syllable with a lax vowel. Such syllables can never be well-formed words in English.

Experiment 3 examined English listeners' ability to detect bisyllabic words with Weak-Strong (WS) or Strong-Weak (SW) stress patterns, in nonsense contexts which could or could not themselves form possible English words. For WS words, *perturb* for example, the contexts consisted of a single consonant (*sperturb*, /spət b/), a Consonant-Vowel (CV) syllable with a tense vowel (*dahperturb*, /dɑpət b/), or a CV syllable with a lax vowel (*dEperturb*, /dɛpət b/). If the PWC is language-specific, *perturb* should be harder to spot after /s/ and /dɛ/ than after /dɑ/, since only the latter residue is a possible word of English. If, on the other hand, the PWC is language-universal, *perturb* should be hard to spot after /s/ but easier after both /dɛ/ and /dɑ/, which could be words in some language.

For SW words (e.g., *echo*) the contexts were single consonants (*shecho*, /ʃɛko/) or CVC syllables, again one with a tense vowel (*fooshecho*, /fuʃɛko/) and one with a lax vowel (*foshhecho*, /fɒʃɛko/). Single-consonant contexts should, again, be difficult. In this case, however, the two syllable contexts did not test whether the PWC is language-specific. Rather, they tested whether the difference between tense and lax vowels influences the location of perceived syllable boundaries. Lax vowels demand a closed syllable (*fosh*), which might lead to the segmentation *fosh-echo*. Detecting *echo* should therefore be easy in this condition, since the word is aligned with the syllable boundary after the /ʃ/ and the entire first syllable is a possible word. Tense vowels, however, allow an open syllable (*foo*), and, combined with the tendency of English to prefer maximal syllable onsets (Pulgram, 1970), this might lead to the segmentation *foo-shecho*. The target *echo* could therefore be as hard to spot in the tense vowel context as in the consonant context, since in both cases there is a single consonant between the beginning of the target and a likely word boundary (cued by the syllable boundary in *foo-shecho* and by the silence in *shecho*).

Method

Subjects. Thirty-six native speakers of English were paid for their participation. Participants were students at Girton College, Cambridge.

Materials and procedure. Forty-eight bisyllabic WS words (*perturb*) and 30 bisyllabic SW words (*echo*) were selected; none had other words embedded within them. The first syllables of the WS words consisted of a single consonant followed by schwa; the SW words all began with vowels. Twenty-four of the WS words were placed in three preceding contexts: a single consonant (*sperturb*); an open CV syllable with a lax vowel (*dEperturb*); and an open CV syllable with a tense vowel (*dahperturb*). It was not possible to find consonant contexts for the other 24 WS words (no phonotactically legal clusters could be formed with words beginning with voiced consonants, like *giraffe*, or those beginning with /s/, like *cigar*). These words were therefore only paired with tense and lax CV contexts. The SW words were also placed in three preceding contexts: a single consonant (*shecho*); a closed CVC syllable with a lax vowel (*foshecho*); and a closed CVC syllable with a tense vowel (*fooshecho*). In all strings, the only embedded real word was the intended target word. The items are listed in Appendix B.

The target-bearing items were divided over three lists, such that all of the SW words and the 24 WS words which had three contexts appeared on all three lists, with type of context counterbalanced over lists. The remaining target-bearing items (WS words with only two contexts) were also divided over the three lists; 16 of these words appeared in each list, each word appearing in only one context in a given list, with type of context counterbalanced over lists. Each list therefore contained 70 target-bearing items. A further 140 filler items containing no real English words were constructed. The fillers matched the target-bearing items in length and stress patterns; there were twice as many fillers with a particular number of syllables and stress pattern as there were target-bearing items with that structure. Each list contained all fillers, with target-bearing and filler items in pseudorandom order, such that there was always at least one filler between any two target-bearing items.

The materials were again recorded onto DAT by a native speaker of English in a sound-attenuated booth. The speaker attempted to minimise syllabification cues in the recording; medial consonants (/ʃ/ in *foshecho* and *fooshecho*; /p/ in *dEperturb* and *dahperturb*) were ambisyllabic, that is, were neither clearly syllabified in the first syllable nor in the second syllable. The procedure was otherwise identical to that of Experiment 1.

Results and discussion

ANOVAs were again performed on the RT and error data. An item was excluded from an analysis if, in any one condition in that analysis, it was missed by more than two-thirds of the subjects who heard it. The items excluded from the analyses involving the WS and SW items with all three

types of context were: *canal*, *convert*, *lapel*, *reserve*, *angel*, *option*, *ever*, *ulcer*, and *usher*. Those excluded from the analyses of WS items with only Tense and Lax contexts were: *convert*, *lapel*, *reserve*, and *behave* (*canal* was included here because only the consonant context for this item failed the criterion). Mean RTs and error rates are summarised in Table 2.

Overall analyses. The first analysis involved the items which were presented in all three contexts: Tense, Lax, and Consonant. The means for the WS words which could be included in this analysis were: Tense, 388 ms and 13% errors; Lax, 446 ms and 10% errors; and Consonant, 501 ms and 27% errors (i.e., in this last case as in Table 2). The means for the SW items are those listed in Table 2. In these analyses, the effect of context was significant both in RTs ($F_1(2, 60) = 19.38, p < .001; F_2(2, 86) = 10.89, p < .001$) and in errors ($F_1(2, 60) = 15.31, p < .001; F_2(2, 86) = 12.42, p < .001$).

TABLE 2

Mean reaction times (RTs, in ms, measured from target offset) and mean percentage error rates, in word spotting (Experiment 3) and in lexical decision to the same target words, excised from their contexts (Experiment 4)

	<i>Weak-Strong Target Contexts</i>		
	<i>Tense Vowel CV Syllable</i>	<i>Lax Vowel CV Syllable</i>	<i>Single Consonant</i>
Word spotting			
Mean RT	423	448	501
Mean error	14%	15%	27%
Example	dahperturb	dEperturb	sperturb
<i>Lexical decision</i>			
Mean RT	271	262	
Mean error	7%	5%	
Example	(dah)perturb	(dE)perturb	
	<i>Strong-Weak Target Contexts</i>		
	<i>Tense Vowel CVC Syllable</i>	<i>Lax Vowel CVC Syllable</i>	<i>Single Consonant</i>
Word spotting			
Mean RT	511	466	607
Mean error	14%	10%	20%
Example	fooshecho	foshecho	shecho
<i>Lexical decision</i>			
Mean RT	313	301	321
Mean error	6%	6%	2%
Example	(foosh)echo	(fosh)echo	(sh)echo

Note. Weak-Strong words from consonant contexts were not included in Experiment 4.

No other effects in either RT or error analysis reached significance across both subjects and items. Planned comparisons between the three contexts for each type of word were then carried out.

WS words. Responses to WS words like *perturb* were faster ($t_1(35) = 2.12, p < .05; t_2(19) = 2.37, p < .05$) and more accurate ($t_1(35) = 5.36, p < .001; t_2(19) = 4.18, p < .005$) in the lax-vowel syllable contexts than in the consonant contexts. This result suggests that the PWC is a language-universal mechanism: CV syllables with lax vowels are not treated as impossible residues in English segmentation, like single consonants are, in spite of the fact that such syllables are not possible English words.

Responses to words like *perturb* were also faster ($t_1(35) = 3.79, p < .005; t_2(19) = 3.85, p < .005$) and more accurate ($t_1(35) = 3.51, p < .005; t_2(19) = 2.51, p < .05$) in the tense-vowel syllable contexts than in the consonant contexts. This result replicates the finding that words are easier to spot in syllabic contexts than in consonantal contexts, as predicted by the PWC. Listeners were also faster to detect WS words in syllable contexts with tense vowels than in syllable contexts with lax vowels ($t_1(35) = 2.61, p < .05; t_2(19) = 2.34, p < .05$). Note however that listeners were slightly more accurate in detecting WS words in syllable contexts with lax vowels than in syllable contexts with tense vowels. This difference suggests that there was a small speed-accuracy trade-off, though this difference was not significant ($t_1(35) = 1.17, p > .2; t_2(19) = 1.12, p > .2$).

We also conducted a second analysis, of all the WS words which appeared in tense and lax syllable contexts (i.e., the words in the previous analysis plus those words like *giraffe* which appeared only in syllabic contexts). In this analysis, the difference between tense and lax syllable contexts (of 25 ms, on average; see Table 2), was not significant (all $F_s < 1$). There was also no difference in error rates between these two conditions in this analysis (all $F_s < 1$). It therefore appears that there was no robust difference between these conditions, while performance in both was reliably better than that in the consonant condition.

SW words. Responses to SW words like *echo* were faster ($t_1(35) = 3.56, p < .005; t_2(24) = 3.68, p < .005$) and more accurate ($t_1(35) = 2.88, p < .01; t_2(24) = 2.33, p < .05$) in the lax-vowel syllable contexts than in the consonant contexts. This difference is again as predicted by the PWC, and replicates Norris et al. (1997). No other differences within the SW words were fully reliable. This means that while listeners were not reliably slower or less accurate in detecting SW words in tense-vowel syllable contexts than in lax-vowel syllable contexts, they were also not reliably faster or more accurate in this condition than in the consonant context condition. The fact that listeners were numerically slightly slower and less accurate in

tense than in lax contexts suggests that there was some tendency for listeners to segment strings like *fooshecho* as *foo-shecho*, thus tending to make detection of *echo* as hard as in *shecho*. But, since the tense-vowel condition was also not reliably different from the lax-vowel condition, this tendency was not very strong. Contexts like *fosh* and *foosh* are both possible words, and there is no clear difference between these two conditions.

The principal result of Experiment 3 is clear. Listeners were able to spot words like *perturb* faster in CV syllable contexts with lax vowels than in single consonant contexts. This supports the claim that the PWC operates according to language-universal principles. Contexts which are possible words in some languages (CVs with lax vowels) should therefore be treated as acceptable residues in on-line speech segmentation in any language. For the same reasons which motivated Experiment 2, however, a control lexical decision experiment was run using the target words excised from the Experiment 3 materials.

EXPERIMENT 4

Method

Subjects. Twenty-five native speakers of English, students at the University of Birmingham, were paid for their participation. None took part in Experiment 2.

Materials and procedure. All 48 WS words used in Experiment 3 were excised from their syllabic contexts in the same way as in Experiment 2. The versions of the 24 WS words which had also been presented in consonant contexts were not used. In many cases it was not possible to find a splicing point which would produce a recognisable token of the word (for example, the words beginning with unvoiced stops, like *canal*, sounded as if they had voiced initial stops). The SW words were also excised from each of their three contexts (three of these words were excluded from the experiment: *ulcer* and *ever*, because they had particularly high error rates in word spotting, and *anchor*, to simplify counterbalancing). There were thus 48 WS words taken from tense and lax CV contexts, and 27 SW words taken from tense, lax and consonant contexts. The SW words were counterbalanced over three lists, such that each list contained nine words taken from each of the three contexts. The WS words were then counterbalanced over these three lists. Two of the lists each contained 24 WS words taken from each context. Two versions of the third list were made using the same counterbalancing of WS words (there were therefore four different lists, which could be treated either as three lists in the SW analysis or as two lists in the WS analysis). The words were presented in

the same pseudorandom order in each list, mixed with 40 WS nonwords and 24 SW nonwords which had been excised from the Experiment 3 fillers. The only difference between the lists was the contexts from which the target words had been excised. The procedure was identical to that in Experiment 2. Each listener heard a short practice list followed by one of the four experimental lists.

Results and discussion

Three WS words were excluded from the analysis because they were missed by two-thirds of the subjects who heard them in one condition (*carafe*, *banal*, and *verbose*). The results for the remaining items are given in Table 2. Separate ANOVAs were performed for the WS and SW words. There were no effects either in speed or accuracy which were significant by both subjects and items, neither for the WS words nor for the SW words. There was therefore no reliable difference between words taken from the different contexts. SW words taken from consonant contexts were just as easy to recognise as those taken from either tense- or lax-vowel contexts. Most importantly, there was no difference between the WS words taken from lax-vowel CV contexts (e.g., *perturb* taken from *dEperturb*) and WS words taken from tense-vowel CV contexts (e.g., *perturb* taken from *dahperturb*). By item ANCOVAs on the word-spotting data for these two conditions from Experiment 3, taking the lexical decision data as covariate, confirmed this: neither the difference in RTs nor the difference in errors was significant. Any potential difference in word spotting performance between these two conditions was therefore not masked by any difference due to the phonetic realisation of the words in each of these contexts.

GENERAL DISCUSSION

The PWC treats CV syllables with lax vowels and syllables with schwa as nucleus in exactly the same way as syllables with full vowels. In other words, although only syllables with full vowels can stand alone and thus serve as well-formed content words in English, any type of syllable is a well-formed possible word in the sense that it constitutes an acceptable residue when appended to a candidate word. It appears therefore that the acceptability of such residue strings is not determined by the demands of a particular language. There are phonological constraints on the form of words acceptable in English, but syllables which violate these constraints are still acceptable residues for English listeners. Thus the definitive criterion for a viable residue in listeners' construction of an acceptable parse of continuous speech seems to be language-universal. A viable residue must be a syllable; but any syllable will do.

The present experiments were conducted in English. Psycholinguistics has a long and embarrassing tradition of claims for language-universality based on data from English alone. However, in this case the English experiments contribute the crucial cornerstone to an edifice built of converging data from many languages. For instance, further evidence that the PWC is indeed a language-universal strategy comes from a word-spotting experiment in Sesotho, a Bantu language spoken in Southern Africa. In Sesotho, any surface realisation of a content word must have at least two syllables. Cutler, Demuth, and McQueen (submitted) asked Sesotho listeners to spot words like *alafa* (to prescribe) in contexts such as *halafa* and *roalafa*. In the former, the single consonant context *h* is an impossible word and an impossible syllable; in the latter, the monosyllabic context *ro* is a possible syllable of Sesotho but is not a well-formed Sesotho content word. Listeners spotted words significantly less rapidly and less accurately in the consonantal contexts than in the monosyllabic contexts. Thus even though *ro* is not a possible content word in Sesotho, this does not make it an unacceptable residue in Sesotho speech segmentation.

Similarly, McQueen and Cutler (1998) observed that Dutch listeners find it harder to spot words in preceding consonantal contexts (e.g., *lepel*, spoon, in /blepəl/) than in preceding CV contexts with schwa (*lepel* in /səlepəl/) or preceding CV contexts with a full vowel (*lepel* in /kulepəl/). As in English, weak syllables are not possible content words in Dutch; they are, however (again as in English), possible function words, and in this experiment were in the position (preceding a content word) which might have been filled by a function word. McQueen and Cutler's experiment thus suggests that a residue does not have to be a possible content word to be acceptable, but it leaves open the question of whether an acceptable residue must fulfil some potential role (e.g., as a function word) in the language in question.

These two experiments in Sesotho and Dutch motivate the same conclusion as the present English experiments: that the PWC is a language-universal strategy. But neither of them alone can fully make the case for this conclusion. From the Sesotho experiment we can conclude that it is not necessary for a syllabic residue to be a potential content word of the language; a monosyllable (which could in Sesotho be a morpheme, but not a lexical item) satisfies the relevant constraint. From the Dutch study, again, we conclude that potential content-word status is not required, since a syllable with schwa (which could in Dutch be a function word but not a content word) behaved just like a full syllable. The English data from the present experiments add the further important evidence that it is not even necessary that the syllabic residue be a viable stand-alone *syllable* of the language. Open syllables with lax vowels, such as /dɛ/, perform no function

in English. They cannot be content words, they cannot be morphemes, they cannot be function words. Nevertheless they constitute residues which are acceptable to the PWC.

The stimuli used in the Dutch and English experiments with reduced syllable residues differ in one interesting respect. Whereas the Dutch stimuli had the residue before the target word, in the English stimuli the residue followed the word. With the following contexts used in Experiment 1, the onset of the weak syllable in /sɪʃəb/ would not constitute a “known boundary” from the perspective of the PWC (neither phonotactics nor the onset of a full syllable signal a boundary). The PWC will therefore not trigger until the “known boundary” determined by the silence at the end of the stimulus. One might, therefore, be tempted to believe that listeners could sometimes respond before the PWC was violated. If this were so one might not expect to see any evidence of a language-specific PWC if listeners responded too quickly.

However, the results from the Dutch word-spotting experiments with reduced syllable residues before the word make this seem unlikely. McQueen and Cutler included stimuli like *səbegin* (/səbəxɪn/) where the target word *begin* has a WS stress pattern. In these materials there was no “known boundary” at word onset. If the PWC were language specific, and schwa were not treated like a full vowel, then the word onset, and indeed every position up to the onset of the strong syllable, would violate the PWC. The word would violate the PWC well before it could be recognised. However, *begin* was spotted as easily in *səbegin* (/səbəxɪn/) as in *geebegin* (/xɛbəxɪn/; with a CV context with a full vowel) and *zasəbegin* (/zasəbəxɪn/; with a bisyllabic context including a full vowel and a schwa). The results for English stimuli with contexts following the words are thus the same as the results in Dutch with contexts preceding the words.

There is a second reason why it is unlikely that the possibility of listeners responding too rapidly in Experiment 1 led us to miss a language-specific effect. This follows from the differences in response time to monosyllabic and bisyllabic words. Monosyllabic words, presumably because they have less perceptual support, are responded to more slowly than bisyllabic words. If there were a language-specific PWC effect that was emerging later in time, then it should be more apparent in the more slowly identified monosyllabic words. However, although there was a marginally larger PWC effect of context for monosyllabic words in RTs, there was a substantially bigger effect in the opposite direction in errors. The fact that these effects go in opposite directions means there is no support for the idea that the PWC effect emerges only for slower responses.

It appears that the PWC operates only on the basis of very simple phonological information (the presence or absence of a vowel), and not on any higher-order knowledge about what constitutes well-formed words or

even syllables in any particular language. Open syllables with lax vowels may constitute independent words in some languages. But they constitute viable residues for lexical activation even in languages in which they do not come into question as potential members of the vocabulary.

Some further recent results in Dutch are consistent with this view. Mauth (1999) has shown that single consonants which are inflectional morphemes in Dutch (the verbal third person singular *-t* and the nominal plural *-s*) are treated in segmentation in the same way as other single consonants which have no such morphemic status in the language. Dutch listeners found it as hard to spot words in morphemic consonantal contexts as in non-morphemic consonantal contexts. These results suggest that the PWC is not sensitive to the fact that in Dutch, the phonemes *t* and *s* are meaningful units; instead, it is a purely phonological mechanism which treats all vowel-less sequences as equivalent (and as equivalently unacceptable for the PWC's purposes).

The only remaining question left open by the current body of evidence on the PWC concerns a very small minority of the world's languages. It is claimed that some languages allow vowel-less syllables and words (see, e.g., Dell and El Medlaoui's [1985] analysis of Tashlhiyt Berber). If the PWC depends simply on the presence or absence of vocalic information in the speech signal, then it is unclear how it might apply in the segmentation of languages such as Tashlhiyt Berber. This question is currently being investigated in our laboratory.

If the PWC is indeed universal, then it could be of considerable value in language acquisition. Infants would not need to learn the relevant constraints for their own native language. If the PWC has the same form across languages, it could assist infants in their initial attempts to segment continuous speech and acquire their first words (Johnson, Jusczyk, Cutler & Norris, 2000; cf. Brent & Cartwright, 1996). Johnson et al. (2000) have recently shown that 12-month-old infants do indeed appear sensitive to PWC factors while listening to continuous speech. The infants were familiarised with lists of monosyllabic words (e.g., *rest*), and then heard passages of continuous speech containing words which had the familiarised words embedded within them. These matrix words were either monosyllabic (e.g., *crest*) or bisyllabic (e.g., *caressed*). The infants looked longer in the direction of a loudspeaker playing the passages if the passages contained the bisyllabic matrix words (where there is a syllabic residue left after the embedded word is found) than if they contained the monosyllabic matrix words (where there is only a consonantal residue). That is, it seemed that they were able to recognise *rest* in the context of *caressed* but not in the context of *crest*. These results suggest that the PWC is a powerful weapon for speech segmentation in the infant's as well as the adult listener's armoury. Note that the continuous speech used in this

experiment provides a more general test of the PWC than the isolated fragments of speech used in word spotting.

We began by asking whether the PWC is a language-specific or language-universal constraint on speech segmentation. The original demonstration of the PWC by Norris et al. (1997) compared consonant residues with syllable residues. This left open the possibility that the critical unit determining the viability of a parse might be either the minimal phonological word of a particular language, or the syllable as a phonological unit in general. The experiments reported here (and especially in combination with the findings of Cutler et al., submitted, and of McQueen & Cutler, 1998) provide a clear answer to this question. Segmentation is impaired when the residue between the end of a candidate word and the nearest known boundary is a consonant, but not when it is a syllable, regardless of whether the syllable is a possible word, of any kind, in the vocabulary of that particular language. The PWC is satisfied by any syllable, irrespective of whether or not the syllable in question might have a place in the language-specific vocabulary.

The simulations reported in Norris et al. (1997) used a modified version of the Shortlist model (Norris, 1994). The algorithm implemented in the model consisted simply of a penalty applied to any candidate word where the stretch of speech input between the edge of that candidate and the nearest known boundary did not contain a vowel. Our results suggest that this simple algorithm, with no special provisions for any particular vocabulary but applicable to all languages in the same way, is likely to be the correct characterization of the PWC. Of course, the boundaries which allow the PWC computations to be performed include phonotactic constraints on syllable structure and reflections of rhythmic structure, so that they themselves are dependent on the rhythmic and phonological characteristics of each language and thus are by no means language-universal. But although these boundaries can be determined in a highly language-specific way, the PWC which operates upon them does so in every language in the same way; it is language-universal.

REFERENCES

- Brent, M.R., & Cartwright, T.A. (1996). Distributional regularity and phonotactic constraints are useful for segmentation. *Cognition*, *61*, 93–125.
- Cutler, A., & Carter, D.M. (1987). The predominance of strong initial syllables in the English vocabulary. *Computer Speech and Language*, *2*, 133–142.
- Cutler, A., Dahan, D., & Donselaar, W. van (1997). Prosody in the comprehension of spoken language: A literature review. *Language and Speech*, *40*, 141–201.
- Cutler, A., Demuth, K., & McQueen, J.M. (submitted). Universality versus language-specificity in listening to running speech.
- Cutler, A., Mehler, J., Norris, D.G., & Seguí, J. (1986). The syllable's differing role in the segmentation of French and English. *Journal of Memory and Language*, *25*, 385–400.

- Cutler, A., Mehler, J., Norris, D.G., & Seguí, J. (1992). The monolingual nature of speech segmentation by bilinguals. *Cognitive Psychology*, *24*, 381–410.
- Cutler, A., & Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, *14*, 113–121.
- Cutler, A., & Otake, T. (1994). Mora or phoneme? Further evidence for language-specific listening. *Journal of Memory and Language*, *33*, 824–844.
- Dell, F., & El Medlaoui, M. (1985). Syllabic consonants and syllabification in Imdlawn Tashlhiyt Berber. *Journal of African Languages and Linguistics*, *7*, 105–130.
- Johnson, E., Jusczyk, P.W., Cutler, A., & Norris, D. (2000). *12-month-olds show evidence of a Possible-word Constraint*. Paper presented to the 140th Meeting, Acoustical Society of America, November.
- Lehiste, I. (1972). The timing of utterances and linguistic boundaries. *Journal of the Acoustical Society of America*, *51*, 2018–2024.
- Mauth, K. (1999). Morphology and the segmentation of spoken language. Poster presented at *Tutorials in Behavioural and Brain Sciences*, Ohlstadt, Germany, July 1999.
- McClelland, J.L., & Elman, J.L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, *18*, 1–86.
- McQueen, J.M. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language*, *39*, 21–46.
- McQueen, J.M., & Cutler, A. (1998). Spotting (different types of) words in (different types of) context. *Proceedings of ICSLP 98, Sydney, Australia* (pp. 2791–2794). Rundle Mall: Causal Productions.
- McQueen, J.M., Norris, D.G., & Cutler, A. (1994). Competition in spoken word recognition: Spotting words in other words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 621–638.
- McQueen, J.M., Otake, T., & Cutler, A. (2001). Rhythmic cues and possible-word constraints in Japanese speech segmentation. *Journal of Memory and Language*, *44*, 103–132.
- Mehler, J. (1981). The role of syllables in speech processing: Infant and adult data. *Philosophical Transactions of the Royal Society, Series B*, *295*, 333–352.
- Mehler, J., Dommergues, J.-Y., Frauenfelder, U.H., & Seguí, J. (1981). The syllable's role in speech segmentation. *Journal of Verbal Learning and Verbal Behavior*, *20*, 298–305.
- Nakatani, L.H., & Dukes, K.D. (1977). Locus of segmental cues for word juncture. *Journal of the Acoustical Society of America*, *62*, 714–719.
- Norris, D.G. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition*, *52*, 189–234.
- Norris, D., McQueen, J.M., & Cutler, A. (1995). Competition and segmentation in spoken word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1209–1228.
- Norris, D., McQueen, J.M., Cutler, A., & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology*, *34*, 193–243.
- Otake, T., Hatano, G., Cutler, A., & Mehler, J. (1993). Mora or syllable? Speech segmentation in Japanese. *Journal of Memory and Language*, *32*, 258–278.
- Pulgram, E. (1970). *Syllable, Word, Nexus, Cursus*. Mouton: The Hague.
- Sebastián-Gallés, N., Dupoux, E., Seguí, J., & Mehler, J. (1992). Contrasting syllabic effects in Catalan and Spanish. *Journal of Memory and Language*, *31*, 18–32.
- Seguí, J. (1984). The syllable: A basic perceptual unit in speech processing? In H. Bouma & D.G. Bouwhuis (Eds.), *Attention and performance X: Control of language processes* (pp. 165–181). Hove, UK: Lawrence Erlbaum Associates Ltd.
- Vroomen, J., & de Gelder, B. (1995). Metrical segmentation and lexical inhibition in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 98–108.
- Weber, A. (2000). The role of phonotactics in the segmentation of native and non-native continuous speech. In A. Cutler, J.M. McQueen & R. Zondervan (Eds.), *Proceedings of the Workshop on Spoken Word Access Processes* (pp. 143–146). Nijmegen: Max Planck Institute for Psycholinguistics.

APPENDIX A

Target-bearing items used in Experiment 1. The targets are spelled in standard English orthography; the contexts are transcribed in the International Phonetic Alphabet (IPA).

<i>Monosyllables</i>		<i>Bisyllables</i>	
<i>Consonant</i>	<i>Schwa</i>	<i>Consonant</i>	<i>Schwa</i>
tea/v/	tea/vəm/	mirror/f/	mirror/fəs/
plough/n/	plough/nəb/	feather/n/	feather/nəθ/
run/dʒ/	run/dʒəm/	ulcer/f/	ulcer/fəp/
pole/f/	pole/fəb/	lager/f/	lager/fəd/
tell/θ/	tell/θəl/	heather/f/	heather/fəm/
fool/tʃ/	fool/tʃəl/	suffer/θ/	suffer/θək/
sun/tʃ/	sun/tʃən/	bother/n/	bother/nəm/
toy/n/	toy/nəθ/	cover/k/	cover/kəf/
chain/θ/	chain/θəb/	leather/n/	leather/nəd/
fun/tʃ/	fun/tʃəθ/	shiver/θ/	shiver/θəf/
doll/v/	doll/vəd/	quiver/n/	quiver/nəl/
shoe/m/	shoe/məf/	powder/m/	powder/məb/
spell/f/	spell/fəb/	gather/m/	gather/məf/
dull/f/	dull/fəd/	eager/θ/	eager/θəb/
sea/f/	sea/fəb/	ever/θ/	ever/θək/
oil/tʃ/	oil/tʃən/	anchor/θ/	anchor/θən/
vein/tʃ/	vein/tʃəl/	foster/n/	foster/nəθ/
boy/s/	boy/səl/	usher/f/	usher/fəm/
key/v/	key/vəm/	nether/f/	nether/fəθ/
smell/f/	smell/fəs/	measure/m/	measure/məb/
knee/ð/	knee/ðəf/	tremor/f/	tremor/fəb/
bell/f/	bell/fəf/	horror/θ/	horror/θəl/
zoo/θ/	zoo/θəb/	sugar/m/	sugar/məl/
gun/tʃ/	gun/tʃəθ/	weather/f/	weather/fəf/

APPENDIX B

Target-bearing items used in Experiment 2. The targets are spelled in standard English orthography; the contexts are given in IPA transcription

Weak-Strong triples (24)

<i>Tense</i>	<i>Lax</i>	<i>Consonant</i>
/zɪ/canal	/zɛ/canal	/s/canal
/dʒaɪ/canoe	/dʒɛ/canoe	/s/canoe
/zəʊ/carafe	/zɪ/carafe	/s/carafe
/gəʊ/caress	/gɪ/caress	/s/caress
/vɔ/cavort	/vɛ/cavort	/s/cavort
/vaɪ/collect	/vɒ/collect	/s/collect
/θ /command	/θɛ/command	/s/command
/vɜ/c ₃ ontain	/vʌ/c ₃ ontain	/s/c ₃ ontain
/lɑʊ/convert	/lɛ/convert	/s/convert
/tʃaɪ/convict	/tʃæ/convict	/s/convict
/lɑʊ/kebab	/lɛ/kebab	/s/kebab
/t /lapel	/tʌ/lapel	/f/lapel
/lɜ/c ₃ liaise	/læ/c ₃ liaise	/b/c ₃ liaise
/fʊ/mature	/fɒ/mature	/s/mature
/v /morose	/vɛ/morose	/s/morose
/rɜ/c ₃ neglect	/rɪ/c ₃ neglect	/s/c ₃ neglect
/dɑ/perturb	/dɛ/perturb	/s/perturb
/g /possess	/gɪ/possess	/s/possess
/kɑ/c ₃ rebuke	/kɪ/c ₃ rebuke	/f/c ₃ rebuke
/lɑ/c ₃ regret	/lʌ/c ₃ regret	/t/c ₃ regret
/vəʊ/reserve	/vɪ/reserve	/θ/reserve
/nɑ/resign	/nɛ/resign	/t/resign
/pɔ/resist	/pʌ/resist	/f/resist
/tʃaɪ/result	/tʃʌ/result	/θ/result

Tense/Lax Weak-Strong items with no matching consonant item (24)

<i>Tense</i>	<i>Lax</i>
/vəʊ/banal	/vɒ/banal
/tʃaɪ/behave	/tʃɒ/behave
/lə/berserk	/le/berserk
/rə/buffoon	/ræ/buffoon
/g /cigar	/gɒ/cigar
/gəʊ/convey	/gɪ/convey
/dɔɪ/degree	/dʊ/degree
/gəʊ/demand	/gʌ/demand
/tʃɔɪ/dessert	/tʃɪ/dessert
/nə/deserve	/nɪ/deserve
/dʒəʊ/design	/dʒʌ/design
/nəɪ/detect	/nɪ/detect
/pəɪ/dissolve	/pɪ/dissolve
/vɪ/disturb	/væ/disturb
/həɪ/divorce	/hʌ/divorce
/tʃəʊ/ façade	/tʃɪ/ facade
/g /fatigue	/gɛ/ fatigue
/kəʊ/gazelle	/kæ/gazelle
/tʃəʊ/gazette	/tʃɪ/gazette
/tʃaɪ/giraffe	/tʃæ/giraffe
/dʒəʊ/select	/dʒɪ/select
/d /survey	/de/survey
/nɪŋ/syringe	/nɒ/syringe
/g /verbose	/gɛ/verbose

Strong-Weak triples (30)

<i>Tense</i>	<i>Lax</i>	<i>Consonant</i>
/zəʊm/abbey	/zɛm/abbey	/m/abbey
/mɔɪz/absent	/mɛz/absent	/z/absent
/tʃəʊg/action	/tʃɛg/action	/g/action
/gɑz/ancient	/gɛz/ancient	/z/ancient
/gɑʊz/angel	/gɛz/angel	/z/angel
/lɑθ/angle	/lɛθ/angle	/θ/angle
/nʌl/angry	/nɛl/angry	/l/angry
/fɔɪk/ankle	/fɛk/ankle	/k/ankle
/vəʊm/anxious	/vɛm/anxious	/m/anxious
/vʌf/apple	/vɛf/apple	/f/apple
/nɔθ/eagle	/nɛθ/eagle	/θ/eagle
/fuʃ/echo	/fɛʃ/echo	/ʃ/echo
/d m/effort	/dɛm/effort	/m/effort
/kɪv/equal	/kɛv/equal	/v/equal
/pɔɪf/extra	/pɛf/extra	/f/extra
/gɪn/image	/gɛn/image	/n/image
/fuʃ/object	/fɛʃ/object	/ʃ/object
/n m/ointment	/nɛm/ointment	/m/ointment
/nɔdʒ/onion	/nɛdʒ/onion	/dʒ/onion
/sɔθ/option	/sɛθ/option	/θ/option
/fʊm/oven	/fɛm/oven	/m/oven
/dɔl/oyster	/dɛl/oyster	/l/oyster
/kəʊs/ugly	/kɛs/ugly	/s/ugly
/vɑθ/uncle	/vɛθ/uncle	/θ/uncle
/vʊl/empty	/vɛl/empty	/l/empty
/fɑʊdʒ/anchor	/fɛdʒ/anchor	/dʒ/anchor
/ziθ/eager	/zɛθ/eager	/θ/eager
/d p/ever	/dɛp/ever	/p/ever
/l ɪ/ulcer	/lɛb/ulcer	/b/ulcer
/dʒ/usher	/dɛb/usher	/b/usher