

Prosodic cues for morphological complexity: The case of Dutch plural nouns

RACHÈL J. J. K. KEMPS and MIRJAM ERNESTUS
Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands

ROBERT SCHREUDER
University of Nijmegen, Nijmegen, The Netherlands

and

R. HARALD BAAYEN
*Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands
and University of Nijmegen, Nijmegen, The Netherlands*

It has recently been shown that listeners use systematic differences in vowel length and intonation to resolve ambiguities between onset-matched simple words (Davis, Marslen-Wilson, & Gaskell, 2002; Salverda, Dahan, & McQueen, 2003). The present study shows that listeners also use prosodic information in the speech signal to optimize morphological processing. The precise acoustic realization of the stem provides crucial information to the listener about the morphological context in which the stem appears and attenuates the competition between stored inflectional variants. We argue that listeners are able to make use of prosodic information, even though the speech signal is highly variable within and between speakers, by virtue of the relative invariance of the duration of the onset. This provides listeners with a baseline against which the durational cues in a vowel and a coda can be evaluated. Furthermore, our experiments provide evidence for item-specific prosodic effects.

Several studies in the visual modality have shown surface frequency effects in the comprehension of fully regular inflections, thus providing evidence for storage of the inflected form as a whole at some level of representation. These effects have been shown for both nouns and verbs, and in several languages. For regularly inflected verbs, evidence for full form storage has been found for Dutch (Baayen, Schreuder, De Jong, & Krott, 2002; Schreuder, De Jong, Krott, & Baayen, 1999) and for English (Alegre & Gordon, 1999). For regularly inflected nouns, evidence for full form storage has been found for Dutch (Baayen, Dijkstra, & Schreuder, 1997), Finnish (Bertram, Laine, Baayen, Schreuder, & Hyönä, 1999), English (Alegre & Gordon, 1999; Sereno & Jongman, 1997), and Italian (Baayen, Burani, & Schreuder, 1997).

Recently, experiments in the *auditory* modality have also shown effects of full form frequency for both nominal and verbal regular inflections in Dutch, suggesting the existence of full form representations of regularly inflected forms in the auditory modality as well (Baayen, McQueen, Dijkstra, & Schreuder, 2003). This finding is

surprising in the light of models of spoken-word recognition that incorporate some form of lexical competition, such as the revised cohort model (Marslen-Wilson, 1990; Marslen-Wilson, Moss, & Van Halen, 1996), TRACE (McClelland & Elman, 1986), and Shortlist (Norris, 1994). In these models, stored regularly inflected forms would be cohort competitors of their corresponding uninflected forms: In many languages (e.g., Dutch, German, and English), the uninflected form is onset embedded in the longer inflected form, and thus, at the phonemic level, the signal is ambiguous until the offset of the last phoneme of the stem (e.g., uninflected [singular] form, *book*; inflected [plural] form, *books*). In other words, the two candidates will keep on competing for recognition (i.e., in some models, inhibiting one another) until after offset of the uninflected form. Storage of regularly inflected forms creates a recognition problem in the domain of inflection, similar to the recognition problem that exists outside the domain of inflection—as, for example, in the perception of onset-embedded words that have longer, morphologically unrelated competitors, such as *ham* in *hamster*.

Using the frequency counts in the CELEX lexical database (Baayen, Piepenbrock, & Van Rijn, 1993), we estimated how often both types of embedding (inflectional embedding vs. morphologically unrelated embedding) occur in Dutch. We selected all 5,129 monomorphemic lemmas that had a lemma frequency greater than zero. Subsequently, we selected all phonological word forms (uninflected and inflected forms) that corre-

Part of this work was made possible by the support of a Major Collaborative Research Initiative (MCRI) grant from the Social Sciences and Humanities Research Council of Canada awarded to Gary Libben. Correspondence concerning this article should be addressed to R. J. J. K. Kemps, Utrecht University, Faculty of Social Sciences, Special Education, P.O. Box 80.140, 3508 Utrecht, The Netherlands (e-mail: r.j.j.k.kemps@fss.uu.nl).

sponded to these lemmas. When we encountered a phonological form that contained an uninflected form at its onset and that shared its stress pattern, we determined whether the stem of that form was the uninflected form. If so, we counted the phonological form as an inflectional continuation (e.g., [buk]–[bukə(n)], “book”–“books”). If the stem was not shared, we counted the phonological form as a morphologically unrelated continuation (e.g., [fiam]–[fiamstəR], “ham”–“hamster”). This procedure resulted in the following counts: 2,188,144 tokens (307 types) were morphologically unrelated continuation forms, and 2,243,990 tokens (3,015 types) were inflectional continuations.¹ These numbers show that inflectional embedding is a highly frequently occurring phenomenon: token-wise, approximately equally as frequent as morphologically unrelated embedding; type-wise, approximately 10 times as frequent as morphologically unrelated embedding. The word recognition system would, therefore, benefit considerably from the presence and the functionality of acoustic cues that would distinguish the segmentally ambiguous portions of uninflected and inflected forms.

In fact, evidence is accumulating that subtle subsegmental acoustic cues can reduce the ambiguity between onset-embedded words and their longer competitors, thus assisting the perceptual system in distinguishing them before the point in the acoustic signal at which disambiguating phonemic information comes in. Salverda, Dahan, and McQueen (2003) recorded participants' eye movements while they listened to Dutch sentences including a word with an onset-embedded word (e.g., *hamster* containing *ham*). The subjects saw four pictures of objects on a computer screen and were instructed to use the computer mouse to move the picture of the object that was mentioned in the sentence. There were more fixations to a picture representing the embedded word (*ham*) when the first syllable of the target word (*hamster*) had been replaced by a recording of the embedded word than when it came from a different recording of the target word. Subtle acoustic information in the speech signal—namely, the duration of the embedded word relative to the duration of its corresponding syllable in the target word—appeared to lead the word recognition system to favor the correct interpretation of lexically ambiguous spoken input.

Experiments by Davis, Marslen-Wilson, and Gaskell (2002) similarly suggest that both durational and intonational differences assist the perceptual system in distinguishing short words from longer morphologically unrelated words beginning with these shorter words. In a gating task, subjects were presented with sentence fragments. In one condition (long-word condition), the sentence fragments ended in a long carrier word of which the initial syllable formed an onset-embedded word (e.g., *captain* containing *cap*). In the other condition (short-word condition), the sentence fragments ended in the short word corresponding to the initial syllable of the carrier word, followed by a word with an onset that

matched the continuation of the longer carrier word (e.g., *cap tucked*; compare *captain*). The first syllable in the short-word condition was significantly longer than the first syllable in the long-word condition, and there was a marginally significant difference in average fundamental frequency (average fundamental frequency was higher in the long-word condition than in the short-word condition). Significantly more short-word responses were made to gates from short-word stimuli than to gates from long-word stimuli, suggesting that listeners are influenced by the acoustic differences that exist between short and long word sequences in responding to the initial syllables of the target word. Similar results were obtained in a cross-modal priming task. The stimuli from the gating task were presented up to the offset of the first syllable of the target word (e.g., *cap* from either *cap* or *captain*) as auditory primes and were followed by a visual target that was either the short word (*cap*) or the long word (*captain*). Greater facilitation occurred when prime syllables came from the same word as the target.

Outside the domain of inflection, listeners thus appear to be sensitive to durational and intonational differences between short words and longer lexical competitors. It is not self-evident that such prosodic differences are also sufficiently present to be functional for inflected words. Consider the Dutch singular and plural forms of “book”: *boek* [buk] and *boeken* [bukə(n)].

First, two phonetic processes exert their influence in parallel: a *shortening* process and a *lengthening* process. For Dutch, the shortening process has been described by Nooteboom (1972). In a stress-timed language such as Dutch, the duration of a stressed vowel decreases as a function of the number of unstressed syllables that follow (see also Lehiste, 1972, and Fowler, 1977, for English; Lindblom & Rapp, 1973, for Swedish). Therefore, the duration of the vowel in the first syllable in *hamster* is expected to be shorter than the duration of the same vowel in *ham*. The same holds for the vowel in the first syllable in *boeken*, as compared with the same vowel in *boek*. However, since the second syllable in *boeken* is less complex than the second syllable in *hamster*, it is conceivable that the amount of shortening in words such as *boeken* versus *boek* is smaller, as compared with the amount of shortening in words such as *hamster* versus *ham*. The amount of shortening might not be enough to be functional for the listener.

Simultaneously, a prosodic lengthening process applies: The last syllable before a prosodic boundary (e.g., a prosodic word boundary or a phonological phrase boundary) is lengthened. Therefore, the form *ham* (which is followed by a word boundary) is expected to be longer than the first syllable in *hamster* (which is not followed by a word boundary). Cambier-Langeveld (2000) points out that when the rhyme of the last syllable consists of a schwa—as, for example, in words such as *boeken* [bukə(n)]—prosodic lengthening also applies to the penultimate syllable. In other words, in *hamster* only the last syllable is subject to prosodic lengthening, whereas

in *boeken* both syllables are lengthened. Thus, it is likely that the difference between *boek* and the first syllable of *boeken* is smaller than the difference between *ham* and the first syllable of *hamster*.

Bearing these phonetic considerations in mind, it is not self-evident that durational modification of the first syllable occurs in inflected forms to the same extent as it does in words carrying onset-embedded, morphologically unrelated words. The durational modification in inflected forms might not be sufficiently present to be functional.

Linguistic considerations lead to the same conclusion. Various linguists have argued that it is preferable for the phonological form of the stem to remain unaltered after affixation. For instance, Aronoff (1976) has pointed out that affixes that leave their base words unchanged—that is, that are phonologically transparent—are more productive than affixes that lead to phonological opacity (see also Dressler, Mayerthaler, Panagl, & Wurzel, 1987, for morphophonological processes in general). In optimality theory, this idea is implemented by means of alignment constraints between prosodic and morphological constituents (e.g., McCarthy & Prince, 1993). These linguistic considerations lead one to expect that it would be dysfunctional for the stem in isolation to differ from the stem followed by an inflectional ending.

Considered jointly, these phonetic and linguistic considerations show that it is not obvious that systematic subsegmental differences between inflected forms and their base words might exist and be functional for the listener.

On the other hand, the existence of functional prosodic differences in the domain of inflection would reduce the competition problem created by the storage of regular inflected forms in the auditory modality. Such differences would distinguish uninflected forms from their longer inflectional counterparts well before the offset of the uninflected form; their uniqueness point would then occur considerably earlier than suggested by their phonemic representation.

Interestingly, an indication that subsegmental differences may exist between uninflected forms and their longer inflectional counterparts has been provided in a pilot study by Baayen et al. (2003). The singular and plural forms of five nouns were realized five times by four native speakers of Dutch. The mean duration of the singulars was longer (98 msec on average) than that of the stems embedded in the plurals.

The question arises as to whether such prosodic cues in the domain of inflection can be functional for the listener, given the enormous variability of speech within and across speakers. In the present article, we address this question by means of an experimental study of regular plural nouns in Dutch. In Dutch, the regular plural form of many nouns consists of the noun stem and the plural suffix *-en* (usually realized as just a schwa; e.g., *boek* [buk] “book”—*boeken* [bukə] “books”). We studied both the combined and the independent effects of dura-

tional and intonational information in the speech signal on the processing of singular and plural forms, using a number decision task, as well as a lexical decision task.² We furthermore investigated whether item-specific prosodic information might affect lexical processing.

EXPERIMENT 1

The question addressed in Experiment 1, in which a number decision task was employed, was whether listeners would be sensitive to prosodic differences between singular forms and the stems of plural forms. If so, listeners would be expected to be slowed down in their responses when there was a mismatch between the prosodic (durational and intonational) information in the acoustic signal of a word, on the one hand, and the word's number as it was conveyed by the presence or absence of the plural suffix, on the other hand. Moreover, the magnitude of the delay in response latencies would be expected to covary with the degree of prosodic mismatch. We tested the covariance between degree of prosodic mismatch and magnitude of the delay in response latencies in a regression design. If listeners were not sensitive to prosodic differences between singular and plural forms—in other words, if listeners relied on segmental information only—mismatching prosodic information should not affect response latencies.

Method

Subjects. Forty-six subjects, mostly students at the University of Nijmegen, were paid to participate in the experiment. All were native speakers of Dutch.

Materials. From the CELEX lexical database (Baayen et al., 1993), we selected all Dutch monomorphemic nouns that met the following criteria. Their initial phoneme was not a vowel, their plural was formed by adding the suffix *-en* [ə(n)] to the stem, and they did not also function as verbal forms. Furthermore, they ended in an underlyingly voiceless plosive. In Dutch, the rule of final devoicing applies: Underlyingly voiced obstruents in syllable-final position are devoiced. The plural suffix *-en* [ə(n)] induces resyllabification of the stem-final obstruent as onset of the next syllable, and hence, an underlyingly voiced stem-final obstruent will remain voiced (Booij, 1995). As a consequence, only stems ending in underlyingly voiceless obstruents phonemically have the same base in the singular as in the plural form. We therefore selected only nouns with stems ending in an underlyingly voiceless plosive, so that there was no change of the voicing characteristics of the plosive when the stems occurred in isolation. Finally, the singular surface frequencies and plural surface frequencies of the nouns were larger than zero. (Singular surface frequency and plural surface frequency are token counts. Token counts in CELEX are based on a corpus of 42 million words of written Dutch.) From the resulting group of 135 nouns, we selected 48 experimental nouns that contained a simplex coda. These nouns are listed in Appendix A. In addition, we randomly selected 48 filler nouns from the group of 133 Dutch monomorphemic nouns that met all the above criteria, except that these nouns could also function as verbal forms.

We excluded nouns containing a complex coda for the following reason. As was mentioned above, for stress-timed languages, the vowel duration in a stressed syllable decreases as a function of the number of unstressed syllables that follow (Lehiste, 1972, and

Fowler, 1977, for English; Lindblom & Rapp, 1973, for Swedish; Nootboom, 1972, for Dutch). This effect of the number of following syllables is smaller with smaller vowel duration in the stressed syllable (Lehiste, 1972; Nootboom, 1972). In other words, the smaller the vowel duration in the monosyllabic singular form, the smaller the difference that is to be expected between the vowel duration in the singular form and the vowel duration in the bisyllabic plural form. Since vowels have a smaller duration when they are followed by a complex coda than when they are followed by a single consonant (Munhall, Fowler, Hawkins, & Saltzman, 1992, for English; Waals, 1999, for Dutch), the difference between singular and plural forms would be expected to be smaller for words ending in a complex coda than for words ending in a single consonant. We decided to exclude nouns with a complex coda, so that the durational difference to be expected between the vowel in the singular form and the vowel in the plural form was maximal.

Three reading lists were created: a list containing the singular forms of the experimental nouns, a list containing the plural forms of the experimental nouns, and a list containing the plural forms of the filler nouns. The order of the nouns within lists was randomized three times, resulting in nine reading lists. In order to maximize durational differences between singular and plural forms, the noun forms were read in isolation. The lists were recorded in a sound-proof recording booth by a native female speaker of Dutch, who was naive regarding the purpose of the experiment. The recordings were digitized at 16 kHz.

For each noun form, the best realization (of three) was selected and spliced out of its list, using the PRAAT speech-editing software (Boersma & Weenink, 1996). Subsequently, out of the experimental noun forms, we created two types of singular forms: *normal* singular forms and *constructed* singular forms. The normal singular form consisted of the singular form exactly as it was uttered by the speaker. The constructed singular form consisted of the stem of the plural form—in other words, it was the plural form with the plural suffix *-en* [ə(n)] spliced off. The point of splicing was located at the onset of the voicing of the schwa following the stem-final consonant. As a result, the normal singular form's prosodic information matched its number information (as conveyed by the absence of the plural suffix), whereas the constructed singular form's prosodic information mismatched its number information: Its prosodic characteristics signaled a plural form, whereas the absence of the plural suffix signaled a singular form. Total duration, vowel duration, closure duration, and release noise duration were measured for the two types of singular forms. Onset of the vowel was defined as onset of voicing if the preceding segment was voiceless and as the end of the release noise if the vowel followed a fully voiced stop. In all other cases (i.e., if the preceding segment was [l, r, m, n, v] or [v]), onset of the vowel was defined as the point of change in the periodicity pattern in the waveform. The end of the vowel and the beginning of closure were defined as the end of the second formant of the vowel. The end of closure was located at the onset of the sudden discontinuity in the waveform for the release noise. A paired *t* test showed that, on average, the constructed singular forms had a significantly smaller total duration than the normal singular forms did [$t(47) = 18.2, p < .0001$]. The mean difference in total duration was 87 msec. The mean difference in vowel duration was 17 msec [$t(47) = 14.8, p < .0001$], the mean difference in closure duration was 26 msec [$t(47) = 10.9, p < .0001$], and the mean difference in release noise duration was 37 msec [$t(47) = 13.8, p < .0001$]. An analysis of variance with total duration as the dependent variable showed that there was no significant interaction between type of singular form (normal vs. constructed singular form) and type of vowel (phonologically and phonetically long vs. short vowel): The difference in duration between normal and constructed singular forms was comparable for words with phonologically and phonetically long and short vowels [$F(1,92) = 0.4, p = .52$]. Table 1 lists

Table 1
Experiment 1: Mean Durations (in Milliseconds,
With Standard Deviations) for Normal and
Constructed Singular Forms

	Normal Singular Form		Constructed Singular Form		Duration Difference
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Whole form	388	73	301	73	87
Vowel	138	45	121	42	17
Closure	88	21	62	14	26
Release noise	76	24	39	15	37

the mean durations with their standard deviations for the two kinds of singulars.

Furthermore, we measured the average fundamental frequencies of the normal and the constructed singular forms. Recall that Davis et al. (2002) found that the average fundamental frequency was higher in the initial syllables of bisyllabic words than in monosyllabic words. We obtained a similar result: The constructed singular forms had a significantly higher average fundamental frequency than the normal singular forms did [$t(47) = -2.0, p < .05$]. The mean difference in average fundamental frequency was 7 Hz (185 Hz for the normal singular forms and 192 Hz for the constructed singular forms). Our explanation for this finding is that all (monosyllabic and bisyllabic) forms were pronounced with an intonational phrase final pitch accent H*L, which was aligned differently in monosyllabic than in bisyllabic words. In the case of the monosyllabic forms, both H and L were realized within one syllable. In the case of the bisyllabic forms, H was assigned to the first (stressed) syllable, and L was assigned to the second syllable. Consequently, average fundamental frequency was higher in the first syllables of the bisyllabic forms than in the monosyllables.

The normal and constructed singular forms functioned as experimental target words. Filler words consisted of the plural filler nouns, exactly as they were uttered by the speaker.

Three experimental trial lists and their complements were created in such a way that each list contained all 48 filler items (plural forms), 24 normal singular forms, and 24 constructed singular forms. One list never contained both the normal and the constructed singular forms of a single noun: If a given list contained the normal singular form of a noun, the constructed singular form of that noun was contained in its complementary list. The order of presentation of the stimuli was pseudorandomized within the three lists: No more than three singular forms of the same type occurred successively. Orders were identical in complementary lists. The subjects were randomly assigned to experimental trial lists. Practice trials were presented prior to the actual experiment. The practice set consisted of 16 trials: 8 plural forms, 4 normal singular forms, and 4 constructed singular forms (all taken from a different recording of the complete experimental and filler sets). None of the nouns in the practice set was presented in the actual experiment.

Procedure. The subjects were instructed to decide as quickly as possible whether the form they heard was a singular or a plural form. They responded by pressing one of two buttons on a button box. All experimental items required the *singular* response, whereas all filler items required the *plural* response (assuming that decision on number is based primarily on the presence vs. absence of a plural suffix). Each trial consisted of the presentation of a warning tone (377 Hz) for 500 msec, followed after an interval of 450 msec by the auditory stimulus. Stimuli were presented through Sennheiser headphones. Reaction times (RTs) were measured from stimulus offset. Each new trial was initiated 2,500 msec after offset of the previous stimulus. When a subject did not respond within 2,000 msec postoffset, a time-out response was recorded. Prior to

the actual experiment, the set of practice trials was presented, followed by a short pause. The total duration of the experimental session was approximately 10 min.

Results and Discussion

No subjects or items were excluded from the analyses, since they all showed error rates below 20%. In all the analyses, only trials eliciting correct responses were included. The mean RTs (measured from word offset and calculated over the correct trials only) and the percentages of incorrect trials for the two kinds of singulars are summarized in Table 2.

If listeners are sensitive to prosodic differences between singular and plural forms, our dependent variable RT should covary with the degree of prosodic mismatch between normal and constructed singular forms. Simply finding a delay in processing [109 msec ; $t_1(45) = -16.0$, $p < .0001$; $t_2(47) = -15.4$, $p < .0001$] is not sufficient evidence for the occurrence of a prosodic mismatch effect, since this delay might as well be a consequence of the splicing manipulation that has been applied to the constructed singular forms. What needs to be shown is a correlation between the magnitude of the prosodic mismatch and the delay in processing.

We therefore applied a covariance analysis along the lines of Lorch and Myers (1990) to the RT data corresponding to the constructed singular forms. We operationalized the amount of prosodic mismatch as the difference between the duration of the constructed singular form and the duration of the corresponding normal singular form. Since a mismatch in intonational contour is not straightforwardly quantifiable—average fundamental frequency does not capture contour information—we did not include intonational mismatch in the numeric operationalization of prosodic mismatch. It is conceivable, however, that the amount of intonational mismatch codedetermined RTs to the constructed singular forms, and we will return to this issue below. We fitted a linear model to the data for each subject separately, with log RT as the dependent variable and log singular surface frequency, duration of the form itself, and the durational difference score as predictors. Singular surface frequency was included as a predictor in order to ascertain that any observed effect of the durational difference score could not be a consequence of confounding durational differences with differences in frequencies between the items. The

t tests on the coefficients of the subjects for the predictors revealed that duration had a facilitatory effect [the longer the duration, the shorter the response latencies; $t(45) = -3.0$, $p < .01$] and that durational difference had an inhibitory effect [the larger the durational mismatch, the longer the response latencies; $t(45) = -3.0$, $p < .01$].

Using a multilevel extension of the Lorch and Myers technique (Pinheiro & Bates, 2000), we tested whether any effect of durational difference remained after partialing out the effects of the other predictors. This was indeed the case [$F(1,1035) = 6.0$, $p < .05$], indicating that durational difference had an independent effect on the RTs to the constructed singular forms.³

Apparently, when listeners segmentally perceive a singular form but prosodically (durational) a plural form is signaled, their number decision is adversely affected. What then happens in the opposite situation? What happens when, segmentally, a plural form is presented but prosodic cues in the stem signal a singular form? In Experiment 2, we investigated whether we could replicate this prosodic mismatch effect for plural forms. We created two types of plural forms: one form in which the prosodic (durational and intonational) cues *matched* the number of the form as conveyed by the presence of the suffix, and one form in which the prosodic cues *mismatched* the number of the form as conveyed by the presence of the suffix.

EXPERIMENT 2

Method

Subjects. Forty-three subjects, mostly students at the University of Nijmegen, were paid to participate in the experiment. All were native speakers of Dutch. None of them had participated in Experiment 1.

Materials. The target items in this experiment were normal and constructed *plural* forms. Contrary to Experiment 1, both types were now created by means of a splicing manipulation, which allows a factorial experimental design contrasting normal and constructed forms. The filler items were now singular forms.

We selected the same experimental and filler nouns as those in Experiment 1. The singular forms of the experimental nouns, the plural forms of the experimental nouns, and the singular forms of the filler nouns were assigned to separate reading lists. The order of the nouns within lists was randomized three times, resulting in nine reading lists. These lists were read by the same native female speaker as in Experiment 1. The lists were recorded in a soundproof recording booth. The recordings were digitized at 16 kHz.

Subsequently, we created the two types of plural forms: normal plural forms and constructed plural forms. Both types of plural forms were created using a splicing technique: The beginning of one speech token was combined with the ending of a different speech token. From both the singular and the plural forms of a noun, we selected the portion of signal from the first phoneme up to and including the closure of the final plosive of the stem. In other words, we selected the stem without the release noise of the final plosive. From another realization of the plural form of the same noun, we selected the portion from the release noise of the final plosive of the stem up to and including the last phoneme. To create the normal plural form, we concatenated the latter portion to the initial portion originating from the plural form. To create the constructed plural form, we concatenated it to the initial portion originating

Table 2
Experiment 1: Mean Response Latencies (RTs, in Milliseconds, With Standard Deviations) Measured From Word Offset (Calculated Over Correct Trials Only) and Percentages of Incorrect Trials for Normal Singular Forms and Constructed Singular Forms

Type of Singular Form	RT		% Incorrect
	<i>M</i>	<i>SD</i>	
Normal	335	44	2.0
Constructed	444	36	1.4

from the singular form. As a result, the normal plural form consisted of two portions of a signal, both originating from plural forms, whereas the constructed plural form consisted of an initial portion originating from a singular form and a final portion originating from a plural form. This splicing manipulation is illustrated in Figure 1.

Note that by applying this splicing procedure to both the normal and the constructed plural forms, we ensured that any observed difference in response latencies could not be a consequence of a difference in splicing manipulation: A delay in processing for the constructed plural forms would constitute sufficient evidence for the occurrence of a prosodic mismatch effect.

Since the initial portion of the constructed plural form originated from a singular form, it was expected to contain durational and intonational cues that mismatched the number of the word as it was conveyed by the presence of the plural suffix. A paired *t* test indeed showed a significant difference in total duration between the normal and the constructed plural forms: The constructed plural form was longer (29 msec on average) than the normal plural form [$t(47) = 5.6, p < .0001$]. The difference in vowel duration was 15 msec [$t(47) = 6.6, p < .0001$], and the difference in closure duration was 19 msec [$t(47) = 6.4, p < .0001$]. Table 3 lists the mean total durations with their standard deviations for the two types of plural forms.

In addition, intonational differences were present between the initial portions of the normal and the constructed forms: The average fundamental frequency of the initial portion of the constructed plural form was, on average, 11 Hz lower than the average fundamental frequency of the initial portion of the normal plural form

Table 3
Experiment 2: Mean Durations (in Milliseconds, With Standard Deviations) for Normal and Constructed Plural Forms

	Normal Plural Form		Constructed Plural Form		Duration Difference
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Whole form	487	77	516	75	29
Vowel	117	45	132	42	15
Closure	65	17	84	23	19

[$t(47) = -10.9, p < .0001$; 190 Hz for the normal plural forms and 179 Hz for the constructed plural forms].

In the case of the constructed plural forms, this splicing procedure gave rise to artificial plosives that combined the closure of a singular realization with the release noise of a plural realization. Or put differently, durational information contained in the original release noise of the singular realization was no longer present in the acoustic signal of the constructed plural form. Recall that we applied this splicing procedure in order to ensure that any observed difference in response latencies could not be a consequence of a difference in splicing manipulation. But would it have been more natural, and more analogous to the creation of the constructed singular forms in the previous experiment, to simply concatenate the plural suffix to the singular stem when forming constructed plural forms? Actually, it turned out that the latter procedure gave rise to very unnatural-




Figure 1. A normal plural form, consisting of two portions of a signal originating from plural forms (upper panel), and a constructed plural form, consisting of an initial portion originating from a singular form and a final portion originating from a plural form (lower panel).

sounding stimuli. In fact, this by itself already exactly answers our research question: A plural form is not simply a singular form with a plural suffix concatenated to it, neither in production nor in perception. The stem in the plural form contains acoustic cues that distinguish it from the same stem in the singular form. In order to prevent the subjects from showing unnatural behavior as a result of the presence of very unnatural-sounding stimuli in the experiment, and in order to determine whether prosodic cues other than the nature of the release noise play a role in the processing of plurals, we opted for the present, somewhat more complicated cross-splicing procedure.

Three trial lists and their complements were created in the same manner as that in Experiment 1: Each list contained all 48 filler items (singular forms), 24 normal plural forms, and 24 constructed plural forms. The subjects were randomly assigned to experimental trial lists. Practice trials were presented prior to the experiment. The practice set consisted of 16 trials: 8 singular forms, 4 normal plural forms, and 4 constructed plural forms. None of the nouns in the practice set was presented in the actual experiment.

Procedure. The same experimental procedure was used as that in Experiment 1, except that now all the experimental items required the *plural* response, and all the filler items required the *singular* response (again assuming that number decision is based primarily on the presence vs. absence of a plural suffix).

Results and Discussion

All items and subjects were included in the analyses, since they all showed error rates below 20%. Table 4 lists the mean RTs (calculated over the correct trials only) and the percentages of incorrect trials for the two types of plural forms.

Table 4
Experiment 2: Mean Response Latencies (RTs, in Milliseconds, With Standard Deviations) Measured From Word Offset (Calculated Over Correct Trials Only) and Percentages of Incorrect Trials for Normal Plural Forms and Constructed Plural Forms

Type of Plural Forms	RT		% Incorrect
	<i>M</i>	<i>SD</i>	
Normal	299	50	1.2
Constructed	323	53	1.0

Paired *t* tests showed a significant difference in response latencies: Response latencies to the constructed plural forms were longer (24 msec on average) than those to the normal plural forms [$t_1(42) = -3.6, p < .001$; $t_2(47) = -2.3, p < .05$]. The physically longer items thus produced the longer RTs. A simple processing explanation (i.e., longer signal to process, longer processing time), however, seems rather unlikely, since RTs were measured from word offset. Furthermore, the covariance analysis in Experiment 1 shows that duration, in fact, has a facilitatory effect: Longer item durations were associated with shorter RTs.

Instead, the prosodic mismatch effect originally observed for singular forms appears to have occurred for plural forms as well. Interestingly, the effect for the plural forms was considerably smaller than the effect for the




Figure 2. The waveform (upper panel) and the intonational contour (lower panel) of a singular stem.

singular forms (24 msec for the plurals in Experiment 2, as opposed to 109 msec for the singulars in Experiment 1). There are three possible explanations for this.

First, the magnitude of the prosodic mismatch was larger for the singulars in Experiment 1 than for the plurals in Experiment 2 [$t(47) = -9.1, p < .0001$]. Whereas in Experiment 1 all durational information carried by the stem of the plural form was contained in the constructed singular form, in Experiment 2 durational information contained in the release noise of the final plosive of the singular was no longer present in the constructed plural form, as a consequence of the splicing procedure. An explanation of the different delay magnitudes between experiments in terms of different mismatch magnitudes is supported by the fact that there was an inhibitory effect of durational difference in Experiment 1 [$F(1,1035) = 6.0, p < .05$], as well as in Experiment 2 [$F(1,999) = 26.1, p < .0001$].

Second, the nature of the expectancy violation in Experiment 1 was different from that in Experiment 2. In Experiment 1, presentation of the constructed singular form led the listener to expect a plural form on the basis of the durational (and possibly intonational) cues that were present in the signal, but then suddenly the auditory signal was broken off, leaving the listener with conflicting evidence. In Experiment 2, presentation of the constructed plural form initially led the subjects to expect a singular form, but then the auditory signal contin-

ued until the end of the plural suffix. Evidence pointing to the plural form thus kept accumulating after the stem, partly compensating for the subtle prosodic cues in the stem that pointed to the singular form. It is possible that this difference in the nature of the violation of the expectancy was reflected also in the different magnitudes of the prosodic mismatch effect in response latencies.

Finally, it is possible that the difference in delay magnitudes between Experiments 1 and 2 was a result of the fact that, in Experiment 1, the manipulation of interest had been systematically confounded with the splicing manipulation. Thus, the delay observed for the constructed forms in Experiment 1 may have partly been the result of the splicing manipulation applied to these forms. There was no such splicing confound in Experiment 2. We cannot rule out this possibility on the basis of our results, but we would like to stress here that the crucial finding in Experiment 1 was not the delay per se, but the relation between the magnitude of the durational differences and the response latencies. This relation shows that the delay observed in Experiment 1 cannot be attributed solely to the splicing manipulation.

The covariance analyses described in Experiments 1 and 2 showed that RTs to the constructed singular forms in that experiment were at least partly determined by the magnitude of the durational mismatch between the normal and the constructed forms. As was mentioned before, mismatch in intonational contour is not as easily quan-




Figure 3. The waveform (upper panel) and the intonational contour (lower panel) of a plural stem.

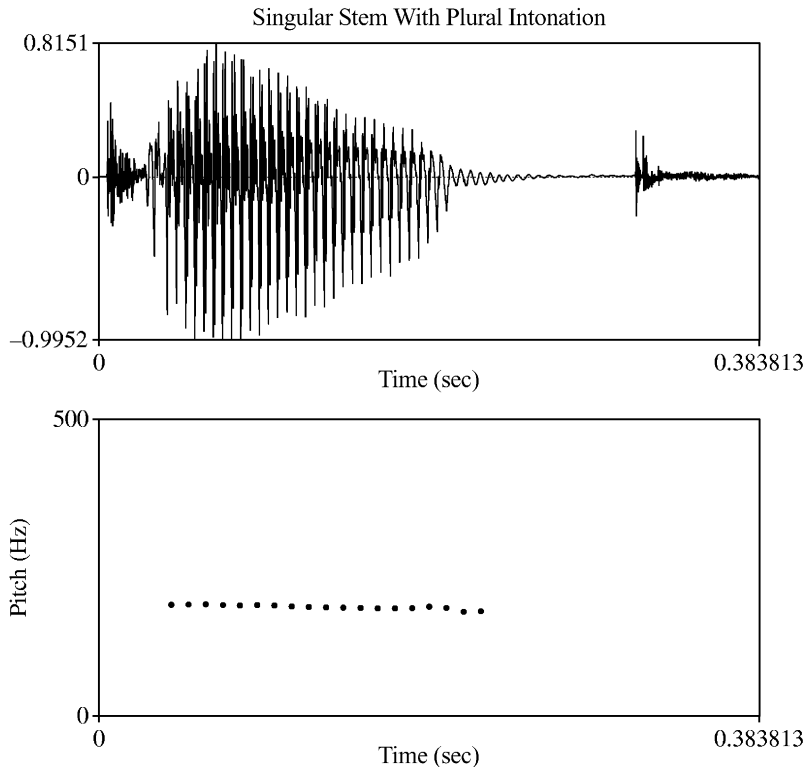


Figure 4. A singular stem with a plural intonational contour (upper panel), resulting from combining the waveform in Figure 2 with the intonational contour in Figure 3 (lower panel).

tifiable and cannot, therefore, similarly be included as a predictor in a linear model. We therefore investigated the individual contribution of intonational information to the prosodic mismatch effect in a separate experiment. In Experiment 3, again, normal and constructed singular forms were presented, but now these two types of singular forms differed only in intonational contour. If intonational cues contribute to the prosodic mismatch effect, we should observe longer response latencies to the forms with the mismatching intonational contour.

EXPERIMENT 3

Method

Subjects. Forty-nine subjects, mostly students at the University of Nijmegen, were paid to participate in the experiment. All were native speakers of Dutch. None of them had participated in Experiment 1 or 2.

Materials. The normal singular forms from Experiment 1 were used, with no further manipulation. In addition, new constructed singular forms were created by taking the normal singular forms and overlaying them with the intonational contours taken from the stems of the plural forms. This manipulation was carried out using the PSOLA (pitch-synchronous overlap and add) resynthesis method in the PRAAT speech-editing program (Boersma & Weenink, 1996). Figures 2, 3, and 4 illustrate the manipulation of the intonational contour. The waveform of the singular stem in the upper panel of Figure 2 was combined with the intonational contour of the plural stem in the lower panel of Figure 3, resulting in the singular stem with the plural intonational contour displayed in Figure 4.

As a result, the durations of the two types of singular forms were identical, but one type of singular form carried the intonational contour of the singular (*normal* singular form), whereas the other type of singular form carried the intonational contour of the plural (*constructed* singular form). The same filler words were used as those in Experiment 1.

Three trial lists and their complements were created in the same manner as that in the previous experiments: Each list contained all 48 filler items, 24 normal singular forms, and 24 constructed singular forms. The subjects were randomly assigned to experimental trial lists. Practice trials were presented prior to the actual experiment. The practice set consisted of 16 trials: 8 plural forms, 4 normal singular forms, and 4 constructed singular forms. None of the nouns in the practice set was presented in the actual experiment.

Procedure. The same experimental procedure was followed as that in the previous experiments.

Results and Discussion

We included all items and subjects in the analyses, since they all showed error rates below 20%. Table 5 lists the mean RTs (calculated over the correct trials only) and the percentages of incorrect trials for the two kinds of singular forms.

The subjects responded, on average, 10 msec more slowly to the constructed singular forms than to the normal singular forms. In a paired *t* test, this difference was significant by subjects [$t_1(48) = -2.2, p < .05$], but not by items [$t_2(47) = -1.5, p = .14$]. Since the item analysis may be too conservative for the type of experimental design used in this study (Raaijmakers, Schrijnemakers,

Table 5
Experiment 3: Mean Response Latencies (RTs, in Milliseconds, With Standard Deviations) Measured From Word Offset (Calculated Over Correct Trials Only) and Percentages of Incorrect Trials for Normal Singular Forms and Constructed Singular Forms When the Two Types of Singular Forms Differed Only in Intonational Contour

Type of Singular Form	RT		% Incorrect
	<i>M</i>	<i>SD</i>	
Normal	333	51	1.7
Constructed	343	47	0.9

& Gremmen, 1999), we additionally ran a covariance analysis (Lorch & Myers, 1990), in which the type of singular form (normal singular form vs. constructed singular form) and the covariate log singular surface frequency factors predicted log RTs. This analysis revealed significant effects of both type of singular form [$t(48) = -2.8$, $p < .01$] and singular surface frequency [$t(48) = -4.7$, $p < .0001$].

These results suggest that when intonational information mismatches number information (conveyed by the presence/absence of the plural suffix), number decision is hindered. Both duration and intonation thus appear to serve as cues in perceptually distinguishing between singular and plural forms. The processing delay for stimuli with mismatching intonational contour was only 10 msec. Note, however, that the stimuli in our experiments were presented in isolation. The subjects did not hear surrounding speech that could function as a frame of reference against which they could evaluate the fundamental frequency of the stimuli. It is conceivable that, when singulars and plurals are presented in their context, intonation serves as a considerably stronger cue than it did in this experiment. An alternative explanation for the relatively small effect of intonational mismatch on RTs is that the intonational difference is peculiar to the context in which the words were produced—contrary to the durational difference, which is probably quite systematically present between singulars and plurals produced in any context. In a list context, each word will have an intonational phrase final contour. This contour will be aligned differently for monosyllables than for bisyllabic forms, leading to differences in average fundamental frequency in the first syllable. However, singulars and plurals do not typically occur in phrase-final position and will, therefore, not show differences in average fundamental frequency that are as systematic as the durational differences. If intonational differences are indeed less systematic than durational differences, it is not surprising that listeners are less sensitive to intonational mismatch than to durational mismatch.

It may be argued that the delay observed for the constructed singular forms was not the result of intonational mismatch but, instead, of the fact that the signal for the constructed singulars had been manipulated, whereas the signal for the normal singulars had not been manipulated. We cannot rule out this possibility. However, the

fact that the constructed singular forms sounded extremely natural suggests to us that intonational mismatch does indeed have a role to play, even though the 10-msec effect observed here may constitute an upper limit for the effect of intonational mismatch for materials presented in isolation. Subsequent research is needed to elucidate the potential effects of intonational information in the speech signal.

In all the experiments so far, we employed a number decision task. In the next and last experiment, we replicated the basic finding, using another experimental paradigm: auditory lexical decision. We opted for lexical decision for two reasons. First, auditory lexical decision is a task in which the number of syllables is irrelevant: Whereas for number decision the number of syllables and, thus, the prosodic structure of the stem are informative with respect to the decision to be made, for lexical decision they are not. A first question addressed in Experiment 4, therefore, was whether listeners would also be sensitive to prosodic cues under these circumstances. Second, the responses to normal and constructed pseudoword singulars might shed light on whether the prosodic mismatch effect observed for existing words results purely from the representations stored in the mental lexicon or whether it is mediated at some prelexical level.

EXPERIMENT 4

Method

Subjects. Forty-two subjects, mostly students at the University of Nijmegen, were paid to participate in the experiment. All were native speakers of Dutch. None of them had participated in Experiments 1–3.

Materials. Four experimental item types were included in the experiment: normal and constructed singular *word* items and normal and constructed singular *pseudoword* items. The word items were the same experimental items as those used in Experiment 1 (i.e., 48 normal singular forms and 48 constructed singular forms).

Out of the singular word items, 48 singular pseudoword items were created by changing one to three phonemes in such a way that the phonotactic constraints of Dutch were not violated and the pseudowords' prosodic structure was identical to that of the words. Subsequently, the plural forms of these pseudowords were created by adding the plural suffix *-en* [$9(n)$], which is the appropriate allomorph, since the stems consisted of a single syllable. The 48 singular and 48 plural forms were assigned to separate reading lists. The orders within these lists were randomized twice, resulting in four reading lists. Due to an error, 1 pseudoword eventually had to be removed from the design. The remaining 47 pseudowords are listed in Appendix B.

In addition, 100 filler words were included in the experiment: 25 monomorphemic uninflected nouns, 25 inflected nouns (plural and diminutive inflections), 25 uninflected and inflected verbs, and 25 uninflected and inflected adjectives. The number of syllables of the filler words ranged from one to three. Out of these filler words, 100 filler pseudowords were created by changing one to three phonemes, again in such a way that the phonotactic constraints of Dutch were not violated and the pseudowords' prosodic structure was identical to that of the words. The filler words and the filler pseudowords were assigned to one reading list. The order within this list was randomized three times, resulting in three reading lists.

One more reading list was created, consisting of 10 words, 5 *singular* pseudowords, and 5 *plural* pseudowords. These items were

used to create practice trials. The order within this list was randomized twice, resulting in two reading lists.

All nine reading lists were recorded by the same native female speaker of Dutch as in the previous experiments. The recordings were made in a soundproof recording booth and subsequently digitized at 16 kHz.

From the reading lists containing the experimental pseudoword items, the best realizations (of two) of the singular and the plural forms were selected. The singular forms served as the normal singular pseudoword items. Constructed singular pseudoword items were created by splicing the *stems* out of the plural forms. From the reading lists containing the filler items, the best realizations (of three) of all the filler words and of all the filler pseudowords were selected. Finally, from the lists containing the practice items, the best realizations (of two) of all the words and pseudowords were selected. The 10 words and the five singular pseudowords were spliced out of the lists exactly as they were realized by the speaker. From the plural pseudowords, we selected only the stems, creating five constructed singular pseudowords.

Three experimental trial lists and their complements were created in such a way that a given list contained 100 filler words, 100 filler pseudowords, 24 normal singular word items, 24 constructed singular word items, 24 (or 23) normal singular pseudoword items, and 23 (or 24) constructed singular pseudoword items. One list never contained both the normal and the constructed singular form of a single noun (word or pseudoword): If one list contained the normal singular form of a noun, the constructed singular form of that noun was contained in its complementary list. The order of presentation of items was pseudorandomized: No more than three singular forms of the same type occurred successively. Orders were identical in complementary lists. The subjects were randomly assigned to experimental trial lists. The 20 practice trials were presented prior to the experiment.

The pseudoword items showed differences in duration and in intonation, similar to those observed in the word items (see Experiment 1): The normal singular forms were significantly longer (94 msec on average) than the constructed singular forms [$t(46) = 21.5, p < .0001$], and the constructed singular forms had a higher average fundamental frequency (5 Hz on average, 198 Hz for the normal forms and 203 Hz for the constructed forms) than the normal singular forms did [$t(46) = -1.9, p < .1$]. The mean difference in vowel duration was 15 msec [$t(46) = 8.1, p < .0001$], the mean difference in closure duration was 16 msec [$t(46) = 6.0, p < .0001$], and the mean difference in release noise duration was 71 msec [$t(46) = 26.1, p < .0001$]. Table 6 lists the mean durations for the normal and the constructed singular pseudowords. The magnitudes of the differences between the normal and the constructed forms were similar for words and pseudowords [duration, $F(1,93) = 1.5, p = .22$; fundamental frequency, $F(1,93) = 0.2, p = .64$].

Procedure. The subjects were instructed to decide as quickly as possible whether the form they heard was a word or a pseudoword. They responded by pressing one of two buttons on a button box. Each trial consisted of the presentation of a warning tone (377 Hz)

Table 6
Experiment 4: Mean Durations (in Milliseconds,
With Standard Deviations) for Normal and
Constructed Singular Pseudowords

	Normal		Constructed		Duration Difference
	Singular Form		Singular Form		
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Whole form	451	66	357	67	94
Vowel	124	42	109	42	15
Closure	93	21	77	15	16
Release noise	108	23	37	12	71

Table 7
Experiment 4: Mean Response Latencies (RTs, in Milliseconds,
With Standard Deviations) Measured From Word Offset
(Calculated Over Correct Trials Only) and Percentages
of Incorrect Trials for Normal and Constructed Singular
Word Items and for Normal and Constructed Singular
Pseudoword Items

Item Type	RT		% Incorrect
	<i>M</i>	<i>SD</i>	
Normal singular word item	442	79	2.3
Constructed singular word item	531	89	3.3
Normal singular pseudoword item	524	87	2.0
Constructed singular pseudoword item	583	79	4.3

for 500 msec, followed after an interval of 450 msec by the auditory stimulus. The stimuli were presented through Sennheiser headphones. RTs were measured from stimulus offset. Each new trial was initiated 2,500 msec after offset of the previous stimulus. When a subject did not respond within 2,000 msec postoffset, a time-out response was recorded. Prior to the actual experiment, the set of practice trials was presented, followed by a short pause. Two short pauses were included in the experiment, resulting in three experimental trial blocks of approximately equal size. The total duration of the experimental session was approximately 30 min.

Results and Discussion

The data for all the subjects were included in the analyses, since they all showed error rates below 20%. Nine word items and three pseudoword items elicited error rates above 20%. These items and their corresponding forms in the complementary condition were excluded from the analyses. Table 7 lists the mean RTs (calculated over the correct trials only) and the percentages of incorrect trials for the four experimental item types (after exclusion of the items with high error rates).

The RTs to the constructed experimental forms were significantly longer (89 msec on average for words, 59 msec on average for pseudowords) than the RTs to the normal experimental forms [$F_1(1,41) = 100.4, p < .0001$; $F_2(1,81) = 55.6, p < .0001$; no interaction of type of singular form and word status, $F_1(1,41) = 1.7, p = .22$, and $F_2(1,81) = 2.3, p = .14$]. In order to rule out the possibility that the observed delay to the constructed singular forms was solely the result of the splicing manipulation applied to these forms, we ran a covariance model on the RT data for the constructed singular forms. A linear model was fitted to the data for each subject separately (cf. Lorch & Myers, 1990), in which log RTs were predicted by the duration of the form itself, by the durational difference score (i.e., the difference in duration between the normal and the constructed forms), and by lexical status (word vs. pseudoword). The t tests on the coefficients of the subjects on the three predictor variables yielded a facilitatory main effect of duration [$t(41) = -8.6, p < .0001$] and a significant interaction of durational difference and lexical status [$t(41) = -4.9, p < .0001$]. A multilevel extension of the Lorch and Myers technique (Pinheiro & Bates, 2000) revealed that for words, durational difference had an inhibiting effect: the

larger the durational difference, the longer the RTs [$t(1815) = -3.2, p < .01$]. For pseudowords, however, we obtained the opposite effect: the larger the durational difference, the shorter the RTs [$t(1815) = 2.11, p < .05$]. In other words, large prosodic (durational) mismatch appears to make words less word-like and pseudowords more pseudoword-like. A comparison between words and pseudowords of the coefficients for the correlation between durational differences and RTs revealed that this correlation was significantly stronger for words than for pseudowords ($z = -2.3, p < .05$).

To conclude, the results of this experiment show that the prosodic mismatch effect is not restricted to the number decision task but is also visible in auditory lexical decision. It is clear that the subjects took the prosodic cues into account, even though these cues were irrelevant for making auditory lexical decisions. Interestingly, the correlational analysis revealed that the prosodic mismatch effect was stronger for words than for pseudowords, suggesting a word-specific component to the prosodic mismatch effect.

GENERAL DISCUSSION

In this study, we investigated whether uninflected and inflected forms have different prosodic characteristics and whether such characteristics are functional for the listener in distinguishing these forms, by reducing the ambiguity between them. We found that, indeed, such acoustic differences exist between uninflected and inflected forms and that listeners are sensitive to them. When prosodic information mismatches segmental information, subjects show a delay in processing (Experiments 1, 2, and 3 [number decision] and Experiment 4 [auditory lexical decision]). We refer to this phenomenon as the *prosodic mismatch effect*. In distinguishing singular forms from the stems of their corresponding plural forms, two sources of nonsegmental information, in particular, play an important part: duration (Experiments 1, 2, and 4) and intonation (Experiment 3). The acoustic mismatch effect occurs in both singulars and plurals (Experiment 2) and in words and pseudowords (Experiment 4). The prosodic differences between uninflected forms and the stems of their corresponding inflected forms reduce the ambiguity between these forms. Our results suggest that these acoustic cues help the perceptual system in determining early in a signal whether an inflected (bisyllabic) or an uninflected (monosyllabic) form is heard.

The existence of the prosodic mismatch effect has important consequences for theories of lexical processing and lexical representation. In classical models of lexical processing, the dominant view has been that all phonetic variation in the speech signal is abstracted from in the course of an acoustic-phonetic analysis, in which the speech signal is translated into a string of discrete phoneme-like units. This abstract string constitutes an intervening representational level through which the speech signal is mapped onto representations in the mental lexicon

(Pisoni & Luce, 1987). Since the abstract segmental representation of the singular form would be identical to that of the stem of the plural form, there is no reason why a delay in processing would occur when there is a mismatch between prosodic and segmental information: After acoustic-phonetic analysis, the processing system no longer has access to prosodic information, at either the prelexical or the lexical level. Thus, models of speech perception that propose a strictly phonemic account of lexical access are challenged by the acoustic mismatch effect observed in the present study.

An alternative account of lexical processing and representation, originally proposed as an answer to the inability of the conventional models to deal with phonological variation, abandons the notion of an intervening segmental level (Lahiri & Marslen-Wilson, 1991). Instead, it assumes that the input to the lexical level is featural. It furthermore assumes that there is a single underlying phonological representation for each lexical item, which abstracts from all surface detail and is compatible with all phonologically permissible variants in a given context. The lexical representations in this framework contain only distinctive and marked information. Predictable information is not specified. For instance, in English, a word-final /n/ can be realized as /n/, as /m/, or as /ŋ/, depending on the place of articulation of the following segment: *green berry* (/m/), *green glass* (/ŋ/), or *green dress* (/n/). Hence, the final nasal of *green* is unspecified for place of articulation. In other words, in this framework, phonemic variation is not represented lexically if it is predictable. This suggests that predictable variation that is prosodic in nature is not represented lexically either. If so, it is unclear how the prosodic mismatch effect might arise in this kind of approach.

An approach that can account for the prosodic mismatch effect is that of Johnson (1997). He trained a connectionist (exemplar-based) model on vector-quantized speech data, which contained, among other things, information regarding the durations of the segments. Johnson's model correctly anticipated whether or not the incoming syllable was followed by another (unstressed) syllable. Davis et al. (2002) also favored a subsymbolic model that is sensitive to subphonemic properties of the acoustic input.

Our explanation for the occurrence of the prosodic mismatch effect is framed in the exemplar-based or episodic approach of Goldinger (1998), but it can be incorporated into other theoretical approaches as well. We think that in parallel to the processing of the acoustic signal of the stem, an expectation regarding the number of unstressed syllables that will follow is built up on the basis of the durations of the segments. A delay in processing will occur when this expectation is violated by the segmental material that either does or does not follow the stem. The buildup of an expectation regarding the possible continuation of the signal would be advantageous at several levels.

First, it would provide information regarding the prosodic makeup of the utterance. Salverda et al. (2003) pointed out that subtle acoustic cues may signal the pres-

ence or the absence of a prosodic word boundary. They argued that a prosodic representation is computed, in part, on the basis of these acoustic cues and in parallel to the segmental encoding. This prosodic representation would contribute to lexical activation by favoring candidates whose boundaries are aligned with the hypothesized prosodic boundary.

Furthermore, the expectation about whether an unstressed syllable is to follow would also provide information regarding the morphological makeup of the incoming speech signal. The prosodic cues signal whether the acoustic signal at hand is that of an unmodified (monosyllabic) stem or that of the same stem, but now followed by an unstressed (inflectional or derivational) suffix or by an (unstressed) clitic. We showed that listeners probably determine whether a stem is part of a morphologically simplex form or not well before the segmental information comes in that signals the presence or absence of a suffix (or clitic).

If it is true that the prosodic mismatch effect arises from the violation of an expectation that is based on the durations of segments, the question arises as to how it is

possible that listeners are sensitive to these durations, given the enormous variability in the temporal structure of speech. Speech rate varies between speakers, within speakers, and within speakers even within one sentence. Hence, the absolute durations of segments will vary tremendously from utterance token to utterance token. We think that the solution of this riddle lies in the relative durations of the segments in the stem.

Consider Figure 5, which summarizes the distributions of durations by means of boxplots of the onset, the vowel, and the coda of the monosyllabic stems of the words from Experiment 1 (upper panel) and the pseudowords from Experiment 4 (lower panel). The boxes show the interquartile range, the horizontal line in the box denotes the median, and the “whiskers” extend to the observations within 1.5 times the interquartile range. Outliers beyond this range are represented by individual circles. Differences in duration that are significant in two-tailed pairwise *t* tests, as well as in two-tailed paired Wilcoxon tests ($p < .0001$), are marked with asterisks.

Figure 5 shows that there is no reliable difference in duration between the onset of the singular form and the

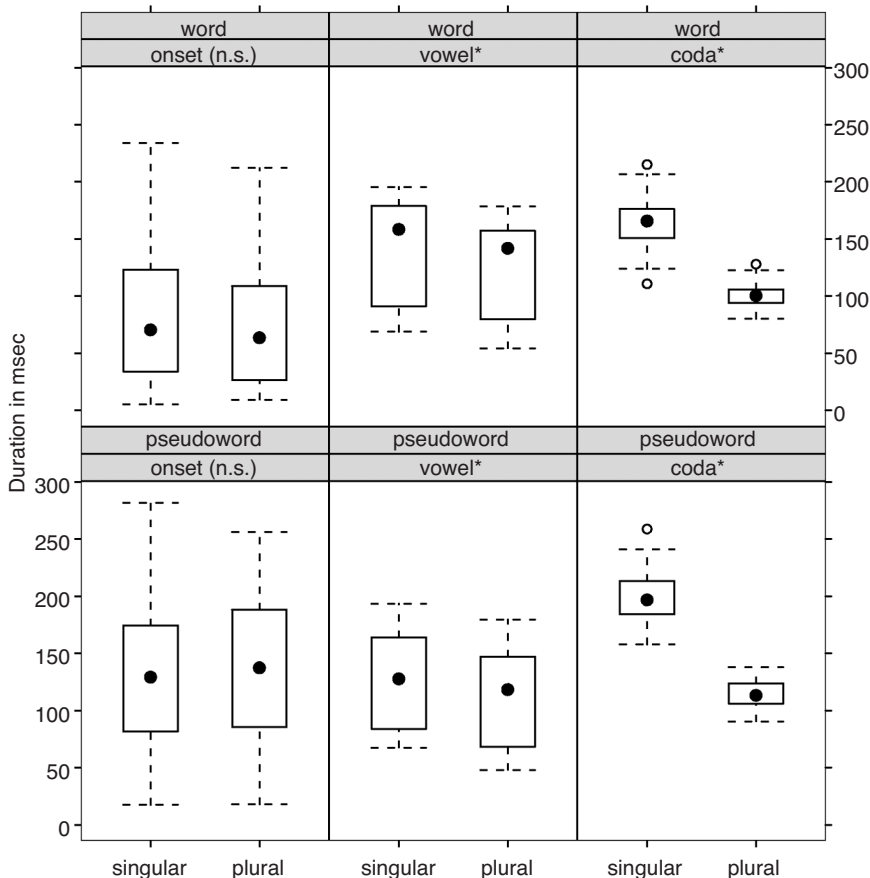


Figure 5. Duration (in milliseconds) of the onset, the vowel, and the coda in singulars and in the stems of plurals for words (upper panel) and pseudowords (lower panel). *Difference in duration significant in two-tailed pairwise *t* test and in two-tailed paired Wilcoxon test ($p < .0001$).

onset of the stem of the corresponding plural form. For the pairs of onsets of existing words, there is a 7-msec difference that fails to reach significance [$t(47) = 1.7, p = .10$]. For the onset pairs in pseudowords, there is an 8-msec difference in the opposite direction (the onsets of stems in plurals tend to be longer than those of singulars) that also does not reach significance [$t(46) = -1.5, p = .14$].

These small and nonsignificant differences in duration of the onset contrast with the longer and very significant difference in duration for the vowels (17 msec for the words and 15 msec for the pseudowords). For the codas, the difference in duration is even greater (63 msec for the words and 87 msec for pseudowords, most of which is due to the release noise duration of the final plosive). Considered jointly, this pattern of results suggests that the duration of the onset is a stable anchor point against which the duration of the vowel, as well as the duration of the coda, can be calibrated. If the durations of vowel and coda, as compared with that of the onset, are relatively long, the incoming speech signal is likely to be a singular. If these durations are relatively short, the likelihood increases that it will be part of a morphological continuation form. In other words, we think that the relative durations of the vowel and the coda with respect to the onset provide the acoustic information that, in our experiments, gives rise to the prosodic mismatch effect.

Relative durations differ from word to word. For instance, the relative duration of the vowel with respect to the onset will depend on whether the vowel is phonemically long or short, as well as on the number of segments in the onset. Similarly, the relative length of the coda varies with the number of segments in the coda and in the onset. In addition, specific combinations of segments in the syllable may affect their duration (Waals, 1999). We therefore hypothesize that the relevant information is provided lexically, with a given lexical form—in our experiments, a given singular or its plural—having a prototypical distribution of relative segmental durations. In other words, we propose that a lexical entry specifies not only the segments and their order, but also the relative durations of the vowel and the coda with respect to the onset. (In the subsymbolic approach of Johnson, 1997, the item-specific distributions would presumably be coded in the weights of the connections in the network mapping vector-quantized speech input onto lexical representations.) This view is consistent with the finding that the correlation between prosodic (durational) mismatch and RTs was stronger for words than for pseudowords, suggesting item-specific support for the prosody-based expectation regarding the number of syllables to follow for existing words.

The prosodic mismatch effect for pseudowords (Experiment 4) points to the existence of a general rule or of an analogical mechanism for building up an expectation of whether an unstressed syllable will follow, since no lexical entries are available for pseudowords. Given an

analogical mechanism that generalizes over stored exemplars, the prosodic mismatch effect in pseudowords can be viewed as resulting from implicit knowledge of prosodic structure that emerges from the patterns that are present in the lexicon. In a subsymbolic framework, the prosodic mismatch effect for pseudowords would reflect the implicit generalizations of the network with respect to the co-occurrences of segmental durations and syllable structure. In more general terms, the prosodic mismatch effect for pseudowords probably reflects the unconditional probabilities for the co-occurrences of segmental durations and syllable structure. In the case of words, these unconditional probabilities might be supplemented by conditional probabilities on the basis of the co-occurrences of the sequence of segments constituting a word's form representation, the durations of these segments, and their syllable structure. The hypothesis that durational structure is part of the lexical representations of words is compatible with Goldinger's (1998) episodic (or exemplar-based) theory, according to which experience with spoken word tokens leaves detailed traces of these tokens in memory. It is also compatible with the linguistic distributional evidence brought together by Bybee (2001), evidence that shows that phonologically redundant information is stored in the (mental) lexicon. Furthermore, it is compatible with Pierrehumbert's (2002) exemplar-based framework, in which each individual word has an associated probability distribution (exemplar cloud) for each of its segments.

The importance of durational information is also supported by the pattern of frequency effects in our experiments, a pattern that strongly suggests that the durational information in the stem codetermines which of two representations (singular or plural) becomes most active. For all the experiments, we conducted multilevel covariance analyses (Pinheiro & Bates, 2000) in which RTs were predicted by duration, singular surface frequency, plural surface frequency, and where applicable, durational difference (between normal and constructed forms). We will discuss only the effects of singular surface frequency and plural surface frequency here.

In the number decision experiments, we observed effects of singular surface frequency in all cases *except* those in which both the segmental and the durational information pointed to a plural form (i.e., in the case of the normal plural forms in Experiment 2). In other words, if either source of information (segmental or durational) in the acoustic signal points to a singular form, the singular representation is activated, even when there is a mismatch between the different sources of information in the signal.⁴ Plural surface frequency, on the other hand, has an effect whenever the durational information points to the plural form, irrespective of what form the segmental information points to (i.e., in the case of the constructed singular forms in Experiment 1 and the normal plural forms in Experiment 2).⁵ In other words, in a number decision task, the durational information in the stem appears to codetermine whether the singular or the plural representation is activated: Durational cues to the plural

form lead to activation of the plural representation, and durational cues to the singular form lead to activation of the singular representation.

When in a number decision experiment, segmental information points to a singular form whereas durational information points to a plural form (i.e., in the case of the constructed singular forms in Experiment 1), we observe competition between the singular and the plural forms: Both the singular and the plural representations are activated. In the normal case (i.e., in the case of the normal singular forms in Experiment 1 and in the case of the normal plural forms in Experiment 2), no competition is observed: Only the correct representations are activated. The ambiguity between the singular and the plural forms appears to be resolved through the durational differences in the stem. This finding reduces the competition problem that is the result of having stored lexical representations for inflected forms in lexical memory. Given the prosodic differences documented in this study, the inflected form might well be a less strong cohort competitor for the uninflected form, and vice versa.

In the lexical decision experiment (Experiment 4), we observed a different pattern of frequency effects. There were facilitatory effects of both singular surface frequency and plural surface frequency, for both normal and constructed singular forms.⁶ We observed no competition, contrary to the pattern in the number decision experiments. Interestingly, for lexical decision, the relevant information is whether the perceived segments form an existing word. Since the distinction between the singular and the plural is irrelevant in lexical decision, the support for the singular and the plural is pooled: Both the singular and the plural representations support a positive lexical decision.

The prosodic mismatch effect documented in this study has important consequences for our understanding of the morphological structure of complex words. The way words are written in such languages as Dutch and English suggests that they consist of stems and affixes that are strung together like beads on a string. Phonemic transcriptions convey the same impression. Our experiments show that this impression is wrong. Plurals are not just singulars with an additional suffix. The precise acoustic realization of the stem provides crucial information to the listener about the morphological context in which the stem appears.

REFERENCES

- ALEGRE, M., & GORDON, P. (1999). Frequency effects and the representational status of regular inflections. *Journal of Memory & Language*, **40**, 41-61.
- ARONOFF, M. (1976). *Word formation in generative grammar*. Cambridge, MA: MIT Press.
- BAAYEN, R. H., BURANI, C., & SCHREUDER, R. (1997). Effects of semantic markedness in the processing of regular nominal singulars and plurals in Italian. In G. E. Booij & J. v. Marle (Eds.), *Yearbook of morphology* (pp. 13-34). Dordrecht: Kluwer.
- BAAYEN, R. H., DIJKSTRA, T., & SCHREUDER, R. (1997). Singulars and plurals in Dutch: Evidence for a parallel dual route model. *Journal of Memory & Language*, **37**, 94-117.
- BAAYEN, R. H., MCQUEEN, J., DIJKSTRA, T., & SCHREUDER, R. (2003). Frequency effects in regular inflectional morphology: Revisiting Dutch plurals. In R. Baayen & R. Schreuder (Eds.), *Morphological structure in language processing* (pp. 355-390). Berlin: Mouton de Gruyter.
- BAAYEN, R. H., PIEPENBROCK, R., & VAN RIJN, H. (1993). *The CELEX lexical database* [CD-ROM]. Philadelphia: University of Pennsylvania, Linguistic Data Consortium.
- BAAYEN, R. H., SCHREUDER, R., DE JONG, N. H., & KROTT, A. (2002). Dutch inflection: The rules that prove the exception. In S. Nootboom, F. Weerman, & F. Wijnen (Eds.), *Storage and computation in the language faculty* (pp. 61-92). Dordrecht: Kluwer.
- BERTRAM, R., LAINE, M., BAAYEN, R.-H., SCHREUDER, R., & HYÖNÄ, J. (1999). Affixal homonymy triggers full-form storage even with inflected words, even in a morphologically rich language. *Cognition*, **74**, B13-B25.
- BOERSMA, P., & WEENINK, D. J. M. (1996). *PRAAT: A system for doing phonetics by computer, Version 3.4* (Rep. 132). Amsterdam: Institute of Phonetic Sciences of the University of Amsterdam. Available at <http://www.fon.hum.uva.nl/praat/>.
- BOOIJ, G. E. (1995). *The phonology of Dutch*. Oxford: University of Oxford Press, Clarendon Press.
- BYBEE, J. L. (2001). *Phonology and language use*. Cambridge: Cambridge University Press.
- CAMBIER-LANGEVELD, T. (2000). *Temporal marking of accents and boundaries*. Amsterdam: LOT.
- DAVIS, M. H., MARSLÉN-WILSON, W. D., & GASKELL, M. G. (2002). Leading up the lexical garden-path: Segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **28**, 218-244.
- DRESSLER, W. U., MAYERHALER, W., PANAGL, O., & WURZEL, W. U. (1987). *Leitmotifs in natural morphology*. Amsterdam: Benjamins.
- FOWLER, C. A. (1977). *Timing control in speech production*. Bloomington: Indiana University Linguistics Club.
- GOLDINGER, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, **105**, 251-279.
- JOHNSON, K. (1997). The auditory/perceptual basis for speech segmentation. *Ohio State University Working Papers in Linguistics*, **50**, 101-113.
- LAHIRI, A., & MARSLÉN-WILSON, W. D. (1991). The mental representation of lexical form: A phonological approach to the recognition lexicon. *Cognition*, **38**, 245-294.
- LEHISTE, I. (1972). *Suprasegmentals*. Cambridge, MA: MIT Press.
- LINDBLOM, B., & RAPP, K. (1973). Some temporal regularities of spoken Swedish. *PILUS* (Papers from the Institute of Linguistics, University of Stockholm), **21**, 1-59.
- LORCH, R. F., & MYERS, J. L. (1990). Regression analyses of repeated measures data in cognitive research. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 149-157.
- MARSLÉN-WILSON, W. D. (1990). Activation, competition, and frequency in lexical access. In G. T. M. Altmann (Ed.), *Cognitive models of speech processing: Psycholinguistic and computational perspectives* (pp. 148-172). Cambridge, MA: MIT Press.
- MARSLÉN-WILSON, W. D., MOSS, H. E., & VAN HALEN, S. (1996). Perceptual distance and competition in lexical access. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 1376-1392.
- MCCARTHY, J. J., & PRINCE, A. (1993). Generalized alignment. In G. E. Booij & J. v. Marle (Eds.), *Yearbook of morphology* (pp. 79-154). Dordrecht: Kluwer.
- MCCLELLAND, J. L., & ELMAN, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, **18**, 1-86.
- MUNHALL, K., FOWLER, C., HAWKINS, S., & SALTZMAN, E. (1992). "Compensatory shortening" in monosyllables of spoken English. *Journal of Phonetics*, **29**, 225-239.
- NOOTEBOOM, S. G. (1972). *Production and perception of vowel duration: A study of the durational properties of vowels in Dutch*. Unpublished doctoral dissertation, University of Utrecht.
- NORRIS, D. G. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition*, **52**, 189-234.
- PIERREHUMBERT, J. (2002). Word-specific phonetics. In C. Gussenhoven & N. Warner (Eds.), *Laboratory phonology VII* (pp. 101-140). Berlin: Mouton de Gruyter.

- PINHEIRO, J. C., & BATES, D. M. (2000). *Mixed-effects models in S and S-PLUS: Statistics and computing*. New York: Springer-Verlag.
- PISONI, D. B., & LUCE, P. A. (1987). Acoustic-phonetic representation in word recognition. *Cognition*, **25**, 21-52.
- RAAIJMAKERS, J. G. W., SCHRIJNEMAKERS, J. M. C., & GREMMEN, F. (1999). How to deal with "the language-as-fixed-effect-fallacy": Common misconceptions and alternative solutions. *Journal of Memory & Language*, **41**, 416-426.
- SALVERDA, A., DAHAN, D., & MCQUEEN, J. (2003). The role of prosodic boundaries in the resolution of lexical embedding in speech comprehension. *Cognition*, **90**, 51-89.
- SCHREUDER, R., DE JONG, N. H., KROTT, A., & BAAYEN, R. H. (1999). Rules and rote: Beyond the linguistic either-or fallacy. *Behavioral & Brain Sciences*, **22**, 1038-1039.
- SERENO, J., & JONGMAN, A. (1997). Processing of English inflectional morphology. *Memory & Cognition*, **25**, 425-437.
- WAALS, J. (1999). *An experimental view of the Dutch syllable*. The Hague: Holland Academic Graphics.

NOTES

1. These counts are based on the phonological transcriptions in the CELEX lexical database. These transcriptions explicitly code final devoicing of underlyingly voiced consonants. For example, the Dutch singular noun *hond* ("dog"), underlyingly ending in a voiced [d], is transcribed as [fɔnt], with a voiceless [t]. The plural *honden* ("dogs") is transcribed as [fɔndən], with a voiced [d]. The singular is thus not counted as embedded in the plural. Furthermore, these counts do not take differences in syllabic structure between longer words and their shorter embedded words into account. For example, both *boek-verkoper* ("book seller") and *boe-ken* ("books") were counted as continuation forms of *boek* ("book"), even though *boe-ken* does not share its syllabic structure with *boek*.

2. In the present study, we concentrated on the perceptual effects of durational and intonational information in the speech signal. Conceivably, monosyllables differ from the stems of their bisyllabic inflected forms in other respects as well, such as in the quality of the vowel and in dynamic spectral information. Subsequent research is needed to uncover the perceptual effects of such differences between monosyllables and the stems of their inflectionally related bisyllabic forms.

3. In the analyses reported here, log RTs were predicted as measured from word offset. A model predicting log RTs as measured from word onset yielded the following results: an inhibitory effect of duration [the

longer the duration, the longer the response latencies: $t(45) = 20.0, p < .0001$] and an inhibitory effect of durational difference [the larger the durational mismatch, the longer the response latencies: $t(45) = -2.7, p < .05$]. The effect of durational difference remained significant after partialing out the effects of the other predictors [$F(1,1035) = 4.1, p < .05$].

4. In Experiment 1 (prosodic [durational and intonational] difference between normal and constructed singular forms), singular surface frequency had a facilitatory effect on RTs to both the normal singular forms [with segmental and prosodic cues pointing to the singular; $t(1044) = -2.7, p < .01$] and the constructed singular forms [with segmental cues pointing to the singular but prosodic cues pointing to the plural; $t(1034) = -3.8, p < .001$]. The higher the singular surface frequency, the easier it was for the subjects to give the *singular* response to both the normal and the constructed singular forms. In Experiment 2 (prosodic [durational and intonational] difference between normal and constructed plural forms), we observed a facilitatory effect of singular surface frequency [$t(998) = -2.9, p < .01$] for the constructed plural forms only (i.e., for the forms that carried the prosodic characteristics of the singular). In Experiment 3, in which the two types of singular forms differed in intonation, but not in duration (and in fact, both carried the durational characteristics of the singular), we observed a facilitatory effect of singular surface frequency for both normal and constructed singular forms [$t(2269) = -2.8, p < .01$].

5. In Experiment 1, plural surface frequency had an inhibiting effect on the RTs to the constructed singular forms only (i.e., to the forms that carried the prosodic [durational and intonational] characteristics of the plural). The higher the plural surface frequency, the more *difficult* it was for the subjects to give the *singular* response to the constructed singular forms. In Experiment 2, we observed a facilitatory effect of plural surface frequency for the normal plural forms [$t(999) = -2.7, p < .01$]. In Experiment 3 (intonational but no durational difference), there was no effect of plural surface frequency, either for the normal or for the constructed forms [$t(2269) = 0.7, p = .50$]. The latter finding shows that only the presence of intonational cues to a particular form is not sufficient to activate that form.

6. In Experiment 4 (lexical decision), we observed facilitatory effects of both singular surface frequency and plural surface frequency for normal singular forms [singular surface frequency, $t(874) = -3.7, p < .001$, one-tailed; plural surface frequency, $t(874) = -2.7, p < .01$, one-tailed] and for constructed singular forms [singular surface frequency, $t(864) = -1.6, p = .05$, one-tailed; plural surface frequency, $t(864) = -1.8, p < .05$, one-tailed]. (We applied one-tailed tests because frequency effects are always facilitatory for lexical decision.)

APPENDIX A
The Experimental Nouns Used in Experiment 1

Singular	Plural	Singular	Plural
1. beek	beken	25. lat	latten
2. boot	boten	26. lip	lippen
3. bout	bouten	27. map	mappen
4. breuk	breuken	28. noot	noten
5. brok	brokken	29. pet	petten
6. buik	buiken	30. peuk	peuken
7. dijk	dijken	31. plaat	platen
8. draak	draken	32. pruijk	pruikken
9. duit	duiten	33. rat	ratten
10. feit	feiten	34. reep	repen
11. geit	geiten	35. rok	rokken
12. graat	graten	36. schaap	schapen
13. grap	grappen	37. spreuk	spreuken
14. grot	grotten	38. straat	straten
15. heup	heupen	39. struijk	struijken
16. kaak	kaken	40. taak	taken
17. kip	kippen	41. tak	takken
18. klip	klippen	42. vak	vakken
19. knaap	knapen	43. vlok	vlokken
20. knop	knoppen	44. wet	wetten
21. krat	kratten	45. wrak	wrakken
22. kruik	kruiken	46. wrat	wratten
23. kuit	kuiten	47. zaak	zaken
24. lap	lappen	48. zweep	zwepen

APPENDIX B
The Pseudowords Used in Experiment 4

Singular	Plural	Singular	Plural
1. beep	bepen	25. kaat	katen
2. bijk	bijken	26. knaat	knaten
3. brek	brekken	27. paak	paken
4. breut	breuten	28. peut	peuten
5. draap	drapen	29. plaak	plaken
6. fap	fappen	30. plik	plikken
7. feik	feiken	31. rak	rakken
8. fek	fekken	32. schoet	schoeten
9. fip	fippen	33. soot	soten
10. fnok	fnokken	34. sprek	sprekken
11. foot	foten	35. strat	stratten
12. frac	frappen	36. stroek	stroeken
13. fruijk	fruijken	37. suik	suiken
14. gaak	gaken	38. tek	tekken
15. get	getten	39. trak	trakken
16. geup	geupen	40. trit	tritten
17. glit	glitten	41. veek	veken
18. gop	goppen	42. weip	weipen
19. gouk	gouken	43. wop	woppen
20. graak	graken	44. wot	wotten
21. grat	gratten	45. wuik	wuiken
22. grok	grokken	46. zaap	zapen
23. gruijk	gruijken	47. zweep	zwoepen
24. guik	guiken		

(Manuscript received September 8, 2003;
revision accepted for publication June 24, 2004.)