

Processing reduced word forms: The suffix restoration effect

Rachèl Kemps,^{a,*} Mirjam Ernestus,^a Robert Schreuder,^b and Harald Baayen^{a,b}

^a Max Planck Institute for Psycholinguistics, P.O. Box 310, 6500 AH Nijmegen, The Netherlands

^b University of Nijmegen, The Netherlands

Accepted 2 December 2003

Available online 31 January 2004

Abstract

Listeners cannot recognize highly reduced word forms in isolation, but they can do so when these forms are presented in context (Ernestus, Baayen, & Schreuder, 2002). This suggests that not all possible surface forms of words have equal status in the mental lexicon. The present study shows that the reduced forms are linked to the canonical representations in the mental lexicon, and that these latter representations induce reconstruction processes. Listeners restore suffixes that are partly or completely missing in reduced word forms. A series of phoneme-monitoring experiments reveals the nature of this restoration: the basis for suffix restoration is mainly phonological in nature, but orthography has an influence as well.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Casual reduced speech; Phonemic restoration; Suffix restoration; Phoneme monitoring; Orthography versus phonology; Dutch

1. Introduction

In spontaneous speech, words are often typically produced with fewer segments, or even syllables, than when they are carefully pronounced in isolation. Especially highly frequent words may occur in highly reduced form. For instance, the Dutch word *eigenlijk* ('actually'), with the canonical pronunciation [eɪχələk], may in casual speech be realised as [eɪχək], and the Dutch word *natuurlijk* ('of course'), with the canonical pronunciation [natyrlək], may be realised as [tyk] (Ernestus, 2000). In both examples, the suffix *-lijk* [lək] is reduced to [k]. Ernestus, Baayen, and Schreuder, 2002 have shown that Dutch listeners recognise highly reduced word forms (taken from a corpus of spontaneous speech) only when such forms are presented in their full context (i.e., in a context of several words). When they are presented in isolation or in a very limited context, listeners do not recognise them. Apparently, not all possible phonetic variants of words are represented in the mental lexicon or, if they are, they are not equally accessible. Otherwise, recognition of reduced word forms in isolation would not be problematic. Furthermore, their findings suggest

that hearing reduced word forms in context induces a process of reconstruction, which makes the reduced forms difficult to distinguish from their non-reduced counterparts. The present study investigates whether reconstruction does indeed take place and what the precise nature of this reconstruction process, which links highly reduced forms such as [eɪχək] to their canonical form ([eɪχələk]), might be.

Phonemic restoration is a powerful auditory illusion that was first reported by Warren (1970). Phonemic restoration typically occurs when part of an utterance is deleted and replaced with an extraneous sound (e.g., a cough or white noise). Listeners report that such an utterance sounds intact. They appear to 'hear' parts of words that are not really present in the acoustic signal (Samuel, 1996a). For phonemic restoration to be effective, the extraneous sound must have some spectral resemblance to the missing speech sound(s) (Samuel, 1981a, 1981b). Replacing speech sounds with silence does not lead to phonemic restoration (Warren, 1970).

Taft and Hambly (1985) showed that listeners might also restore reduced vowels to their full counterparts. In a syllable-monitoring task, listeners tended to accept a match between a target syllable and a word when the vowel of the target syllable was full while the vowel of the word was reduced, but only when the full vowel in

* Corresponding author. Fax: +31-24-352-12-13.

E-mail address: rachel.kemps@mpi.nl (R. Kemps).

question was consistent with the spelling of the reduced vowel. For example, while listeners heard [læ ɟ] in *la-goon* [ləgu:n], they did not hear [lɔɟ] in the same word.

Casual reduced speech is fundamentally of a very different nature than the stimuli used in the traditional phonemic restoration studies, which tend to be carefully realised stimuli. In spontaneous reduced speech, phonemes are missing, but they are neither replaced by an extraneous sound nor by reduced phonemes. They are simply not realised. Their absence is inherent to the type of speech, and it is systematic and highly frequent. Furthermore, the reductions in spontaneous speech are more dramatic than, for instance, the vowel reductions studied by Taft and Hambly (1985): Complete morphemes may be reduced to a single phoneme, or they may not be realised at all. Experiment 1 was designed to answer the question whether listeners might nevertheless restore missing speech sounds in highly reduced speech.

2. Experiment 1

We presented listeners with words ending in the highly frequent derivational suffix *-(e)lijk* [(ə)lək] ('-ly'). In spontaneous speech, this suffix is often severely reduced. For instance, *vreselijk* [frɛsələk] ('terrible') may be realised as [frɛsk], and *eigenlijk* [ɛiχələk] ('actually') may be realised as [ɛiχək]. In these examples, the phoneme [l] is present in the non-reduced realisations, but absent in the reduced realisations. We presented our participants with both non-reduced realisations (containing the phoneme [l]) and reduced realisations (not containing the phoneme [l]), and asked them to perform a phoneme-monitoring task, with the target phoneme [l]. We presented the target words in two conditions. In the first part of the experiment, we presented the target words in a context of several words (Full Context). In the second part of the experiment, we only presented the suffix or whatever was left of it (Minimal Context). In Full Context, participants were expected to recognise the reduced word forms (cf. Ernestus et al., 2002), and thus, restoration of missing phonemes on the basis of the activated non-reduced lexical representations of the target words should be possible. In this condition, our participants may report the presence of the phoneme [l], not only for the non-reduced forms, but also for the reduced forms. In other words, we expected many 'false positive' responses (i.e., reporting the presence of the phoneme [l] when this phoneme is not actually present in the acoustic signal). In Minimal Context, however, only small fragments of the target words were presented. As a result, the representations of the target words should not be activated and thus restoration of missing phonemes should not be possible. This condition enabled us to establish whether listeners are capable of accurately detecting the presence of the phoneme [l] when their

perception cannot be influenced by the activated representations of the target words. Accurate detection performance in Minimal Context paired with many false positives in Full Context would constitute solid evidence for restoration of missing phonemes in reduced word forms.

2.1. Method

2.1.1. Participants

Fifty-one participants, mostly students at the University of Nijmegen, were paid to participate in the experiment. All were native speakers of Dutch.

2.1.2. Materials

We selected 11 Dutch words ending in the derivational suffix *-(e)lijk* [(ə)lək] ('-ly') as target stimuli. From a corpus of spontaneous speech (Ernestus, 2000), we selected both a non-reduced and a reduced realisation of every target word (see Appendix A). In the non-reduced variants, the suffix was always fully realised. In the reduced variants, the suffix was either completely or partly reduced, but in no case the phoneme [l] was present in the realisation (as established by two trained phoneticians).

The filler words in our experiment were 60 words, half of which contained a derivational affix other than *-(e)lijk* (e.g., *be-*, *ge-*, *ont-*, *-baar*, *-isch*), and half of which did not contain a derivational affix. Half of all filler words contained the phoneme [l]. For every filler word, two non-reduced tokens were selected from the corpus of spontaneous speech mentioned above. This resulted in 60 (non-reduced) realisations of the filler words containing the phoneme [l], and 60 (non-reduced) realisations of the filler words without the phoneme [l].

The target items and the filler items were put into a list and the order of presentation was pseudo-randomised three times (target items never occurred consecutively), resulting in three lists. Ten practice items were selected from the speech corpus: five non-reduced realisations of words containing the phoneme [l] and five non-reduced realisations of words without [l]. The order of presentation of the practice items was also randomised three times.

All items were presented in the two conditions. In Full Context (first part of the experiment), the items were presented in a context of several words. The context never contained the phoneme [l]. In Minimal Context (second part of the experiment), only the suffix *-(e)lijk* was presented for the non-reduced versions of the target words. For the reduced versions of the target words, we presented whatever was left of the suffix. For example, for the non-reduced realisation of *eigenlijk* ([ɛiχələk]), we presented [ələk], and for the reduced realisation of *eigenlijk* ([ɛiχək]), we presented [ək]. For the filler and the practice items containing the phoneme [l], we presented

only the phoneme [l] plus approximately two to three surrounding phonemes. Finally, for the filler and the practice items without the phoneme [l], we presented a randomly selected portion of signal consisting of approximately two to three phonemes.

2.1.3. Procedure

Participants performed a phoneme-monitoring task, with the target phoneme [l]. They were instructed (on paper) to listen carefully to the stretches of speech that were presented to them, and to decide whether these stretches contained the sound [l]. The participants responded by pressing one of two buttons on a button box. Immediately after each button press, the participants indicated how confident they were about their decision, by circling one of the numbers on a rating scale ranging from 1 to 5 (1 = least confident, 5 = most confident). Each trial consisted of the presentation of a warning tone (377 Hz) during 500 ms, followed after an interval of 200 ms by the auditory stimulus. Stimuli were presented through Sennheiser headphones. After circling one of the numbers on the rating scale, the participants moved on to the next trial by pressing either button on the button box. The total experiment consisted of 284 trials: The first 142 trials constituted the Full Context Condition, the second 142 trials constituted the Minimal Context Condition. Prior to both conditions, the corresponding set of practice trials was presented, followed by a short pause. Within both conditions, there was a pause half way through, that is, after 71 trials. The two conditions were separated by another pause. The total duration of the experimental session was approximately 30 min.

2.2. Results and discussion

In total, 2244 target trials were presented (11 target stimuli \times 2 types \times 2 conditions \times 51 participants). Due to technical failure (the software not registering some button presses), the responses to only 2013 trials were recorded (i.e., the proportion of missing trials was 10.3%). Table 1 summarises the numbers of 'yes'-responses ('Yes, I heard the sound [l]') to the non-reduced and the reduced target words, in both Full Context and Minimal Context.

The shaded cells contain the numbers of false positive responses (i.e., reporting the presence of the phoneme [l] when no [l] was actually present in the acoustic signal).

The number of false positive responses was considerably higher in Full Context than in Minimal Context. A logistic regression analysis with the ratio of the numbers of 'yes'- and 'no'-responses for a given item as the dependent variable, and Type of word form (non-reduced versus reduced word form) and Context (Full Context versus Minimal Context) as factors yielded significant main effects of Type of word form ($\chi^2(1) = 620.7$, $p < .0001$) and Context ($\chi^2(1) = 299.8$, $p < .0001$), and a significant interaction of Type of word form and Context ($\chi^2(1) = 162.75$, $p < .0001$). Logistic regression analyses on the data for the two Context Conditions separately revealed significant simple effects of Type of word form for both Context Conditions (Full Context: $\chi^2(1) = 81.8$, $p < .0001$; Minimal Context: $\chi^2(1) = 811.0$, $p < .0001$). The pattern in the interaction between Type of word form and Context shows that, when the critical stretches of speech (i.e., the non-reduced and the reduced suffixes) are presented outside their linguistic context, listeners are capable of accurately detecting the phoneme [l]. In contrast, when these stretches of speech are presented within their linguistic context, listeners tend to incorrectly report the presence of the phoneme [l] in the reduced word forms. In other words, listeners restore the missing phoneme [l].

For the analysis of the confidence scores (and for all following analyses of confidence scores), we pooled the scores for all 'yes'-responses over all items and all subjects, since not all items and not all participants produced false positives and hits. As the scores were not distributed normally, we analysed them by means of Wilcoxon's Test. The average confidence scores of the 'yes'-responses in Full Context were not significantly different for the non-reduced and the reduced word forms (4.9 and 4.8, respectively, Wilcoxon's $W = 33379.5$, $p = .16$, two-tailed). This shows that the participants were equally confident about having heard the phoneme [l], whether or not this phoneme [l] was actually present in the acoustic signal. Apparently, the restoration escaped participants' awareness (as opposed to being the result of a conscious decision strategy). This was supported by the participants' remarks after the experiment, which gave no indication of the use of a conscious decision strategy.

Reaction times were measured from the onset of the first segment following the [l], or in the case of the reduced word forms, from the onset of the first segment that followed the position at which the [l] would have been realised. For the analysis of the reaction times (and for all following analyses of reaction times), we pooled the reaction time data for the 'yes'-responses over all items and all subjects, and we applied Wilcoxon's Test, for the same reasons described above for the confidence scores. The analysis showed that, although restoration of missing phonemes may occur at an unconscious level, it does take time. In Full Context, the reaction times

Table 1
Numbers of 'yes'-responses to non-reduced and reduced word forms in Full Context and in Minimal Context in Experiment 1

Type of word form	Full Context	Minimal Context
Non-reduced ([l] present)	418	439
Reduced (no [l] present)	308	16

corresponding to the false positives were significantly longer (242 ms on average) than the reaction times corresponding to the hits (i.e., correctly reporting the presence of [ɪ]; Wilcoxon's $W = 84063$, $p < .0001$). Our explanation for this finding is that the false positive responses are mediated by a lexical level of representation that becomes available at or after lexical access, whereas the hits reflect a phonetic level of processing (cf. Hallé, Chéreau, & Segui, 2000).

To summarise, listeners restore phonemes that are missing in reduced word forms. Confidence scores suggest that listeners are not aware of the fact that they do. Nevertheless, the restoration takes time, probably because it is dependent on the lexical information that becomes available only after the intended word has been recognised.

3. Experiment 2

In the Minimal Context Condition in Experiment 1, we presented very short stretches of speech, only containing the suffix itself, such that the participants would not recognise the words. It is possible that a trace of a 'missing' [ɪ] was present in the Full Context, for instance, in the form of a change in the quality of the preceding vowel, that was absent in the short stretches of speech presented in the Minimal Context. If so, then it is no wonder that participants more often responded 'no' in Minimal Context than in Full Context (when the whole form, including the possible [ɪ]-trace, was presented to them). In order to establish whether traces of the [ɪ] were indeed present in our reduced word forms, we ran a control experiment, employing the same phoneme-monitoring task as in Experiment 1. We presented the same target stimuli as in the Minimal Context Condition of Experiment 1, but now a larger portion of them was presented, so that any possible acoustic trace of [ɪ] would be included.

3.1. Method

3.1.1. Participants

Ten 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.

3.1.2. Materials

The target stimuli in this experiment were the same 11 non-reduced and 11 reduced realisations as in Experiment 1. We now presented the suffix *-(e)lijk* plus the preceding vowel and any intervening consonant. For example, for the non-reduced realisation of *vreselijk* [frɛslək], we presented [ɛslək]. For the reduced realisation of *vreselijk* [frɛsk], we presented [ɛsk]. Any trace of

the [ɪ] left in the preceding vowel ([ɛ]) is now included. Twenty-two non-reduced filler items were added, pseudo-randomly selected from the set of filler items of Experiment 1 (one realisation per word). Half of these filler items contained the phoneme [ɪ]. In addition, 10 non-reduced practice items were presented, half of which contained an [ɪ]. For the filler items and the practice items, we presented stretches of 4–5 phonemes.

The target items and the filler items were put into a list and the order of presentation was pseudo-randomised three times (the first three items were filler items, no more than two target items occurred consecutively, and the non-reduced and the reduced variants of one word never occurred consecutively). The order of presentation of the practice items was also randomised three times.

3.1.3. Procedure

The participants again performed a phoneme-monitoring task (target phoneme [ɪ]). For the details of the procedure, see Experiment 1. The total experiment consisted of 10 practice trials and 44 experimental trials. The total duration of the experimental session was approximately 5 min.

3.2. Results and discussion

In total, 220 target trials were presented (11 target stimuli \times 2 types \times 10 participants). Due to technical failure, the response to one trial was not recorded (i.e., the proportion of missing trials was 0.5%).

For the non-reduced forms, 9 out of 11 forms received 100% 'yes'-responses, indicating that the phoneme [ɪ] was clearly perceivable in these forms. One form (*waarschijnlijk*) received 87% 'yes'-responses and one form (*eerlijk*) received only 56% 'yes'-responses. For the reduced forms, 8 out of 11 forms received 100% 'no'-responses, indicating that these forms did not contain the phoneme [ɪ] or any perceivable trace of it. One form (*mogelijk*) received 56% 'no'-responses, one form (*duidelijk*) received 33% 'no'-responses, and one form (*onmiddellijk*) received 11% 'no'-responses. Clearly, these three forms contained some acoustic trace leading the participants to perceive the phoneme [ɪ].

Given these results, we re-analysed the data of Experiment 1, now excluding the non-reduced form of *eerlijk* and the reduced forms of *mogelijk*, *duidelijk*, and *onmiddellijk* from both the Full Context Condition and the Minimal Context Condition. Table 2 summarises the resulting numbers of 'yes'-responses to the non-reduced and the reduced target words, in both Full Context and Minimal Context.

The number of false positive responses was still considerably higher in Full Context than in Minimal Context (shaded cells). A logistic regression analysis with the ratio of the numbers of 'yes'- and 'no'-responses

Table 2

Numbers of 'yes'-responses to non-reduced and reduced word forms in Full Context and in Minimal Context in Experiment 1, after exclusion of the items *eerlijk* (non-reduced form), and *mogelijk*, *duidelijk*, and *onmiddellijk* (reduced form)

Type of word form	Full Context	Minimal Context
Non-reduced ([l] present)	378	435
Reduced (no [l] present)	184	7

for a given item as the dependent variable, and Type of word form (non-reduced versus reduced word form) and Context (Full Context versus Minimal Context) as factors again yielded significant main effects of Type of word form ($\chi^2(2) = 721.3$, $p < .0001$) and Context ($\chi^2(1) = 126.3$, $p < .0001$), and a significant interaction of Type of word form and Context ($\chi^2(1) = 156.3$, $p < .0001$). Logistic regression analyses on the data for the two Context Conditions separately yielded significant simple effects of Type of word form for both conditions (Full Context: $\chi^2(1) = 118.4$, $p < .0001$; Minimal Context: $\chi^2(1) = 816.5$, $p < .0001$). In other words, the analyses of the response data without the problematic items lead to similar results as the original analyses.

The same holds for the analysis of the confidence scores. In Full Context, the average confidence scores for the 'yes'-responses did not differ for the non-reduced word forms and the reduced word forms (4.9 and 4.8, respectively; Wilcoxon's $W = 33379.5$, $p = .16$, two-tailed). In addition, the reaction times were still longer (448 ms on average) for the false positives than for the hits (Wilcoxon's $W = 54190$, $p < .0001$, two-tailed).

In summary, listeners are able to accurately detect the presence of the phoneme [l] in phoneme strings that are presented without any linguistic context. When the same phoneme strings are presented within their natural context, listeners frequently incorrectly report the presence of the phoneme [l] in reduced word forms, just as confidently as when they correctly report the presence of the phoneme [l] in non-reduced word forms. In other words, listeners restore the missing phoneme [l].

4. Experiment 3

Restoration is based on the information that becomes available once a lexical representation is accessed (Samuel, 1981a, 1987, 1996b; Warren, 1970). This information may be phonological or orthographic in nature. Previous studies have shown that orthographic information is influential in phoneme-monitoring experiments (Dijkstra, Roelofs, & Fieuws, 1995; Donnenwerth-Nolan, Tanenhaus, & Seidenberg, 1981; Hallé et al., 2000; Seidenberg & Tanenhaus, 1979; Taft & Hambly, 1985; Treiman & Cassar, 1997), especially when the experiment contains many words with deviant spellings (Cutler, Treiman, & Van Ooijen, 1998).

All these experiments concerned careful laboratory speech, which is relatively slow compared to natural, spontaneous speech. It is not clear what the nature of the lexically provided information may be that results in the restoration of highly reduced word forms in spontaneous speech, as observed in Experiment 1. Is the restoration the result of the activation of the phonological code of the word, is it the result of the activation of the orthographic code, or is it the result of the activation of both?

An answer to this question may be found in a comparison of the results of our first two experiments with the results of a similar experiment in which restoration, if at all, must necessarily take place on the basis of orthographic information. In Experiment 3, we presented our participants with items that either did or did not contain a mismatch between the orthographic and the phonological code. The matching items were words for which the orthographic code contained the digraph *ei/ij* (the two possible spellings of [ɛi] in Dutch) and the phonological code contained the phoneme [ɛi] (e.g., *bi-jna* [bɛina] 'almost'). The mismatching items were words ending in the suffix *-(e)lijk*. This suffix forms one of the few exceptions to the generalisation that the digraph *ij* in Dutch is realised as [ɛi], as, in this suffix, the digraph *ij* is always pronounced as [ə]. Thus, in words ending in *-(e)lijk*, there is a mismatch between the orthographic code and the phonological code. Since [ɛi] is pronounced as [ə] in only very few morphemes, this mismatch is more severe than the vowel reductions studied by Taft and Hambly (1985).

Again, the participants performed a phoneme-monitoring task, but now the target phoneme was [ɛi] instead of [l]. As in Experiment 1, all items were presented in two conditions. In the Full Context Condition, the items were presented in a context of several words. If the restoration phenomenon that we have observed in the corresponding condition in Experiment 1 were (at least in part) the result of the activation of orthographic codes, we would expect our participants to (incorrectly) report the presence of the phoneme [ɛi] in the words with the suffix *-(e)lijk* in this condition. In contrast, if the restoration phenomenon were solely the result of the activation of phonological codes, we would expect our participants to (correctly) report the absence of the phoneme [ɛi] in the words with the suffix *-(e)lijk*.

In the Minimal Context Condition, we presented only the phoneme corresponding to the digraph *ij/ei* (i.e., [ə] for the words with *-(e)lijk* and [ɛi] for all other words) plus two to three surrounding phonemes. In this condition, we expect participants to accurately detect the presence of the phoneme [ɛi]. If restoration of phonemes occurs on the basis of activated orthographic representations, we may expect accurate detection performance in Minimal Context paired with many false positive responses in Full Context.

4.1. Method

4.1.1. Participants

Forty-seven participants, 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.

4.1.2. Materials

From the corpus of spontaneous speech (Ernestus, 2000), we selected 22 non-reduced realisations of words with the suffix *-(e)lijk* [ə]lək], that is, words with a mismatch between the orthographic code and the phonological code (see Appendix B). Additionally, we selected 22 non-reduced realisations of words in which there is no such mismatch: words of which the orthographic code contains the digraph *ei* or *ij* and the phonological code contains the diphone [ɛi] (e.g., *bijna* [bɛina], see Appendix B).

Sixty-four filler items were added: 20 non-reduced realisations of words containing both *ijlei* and [ɛi], with derivational affixes other than *-(e)lijk* (e.g., *be-*, *ge-*, *ont-*, *-baar*, *-isch*); 22 non-reduced realisations of words not containing *ijlei* or [ɛi], with derivational affixes other than *-(e)lijk*; and 22 non-reduced realisations of words, without derivational affixes, not containing *ijlei* or [ɛi]. Twenty practice items were selected, half of which contained both *eilij* and [ɛi] and half of which did not contain *eilij* or [ɛi]. The target items and the filler items were put into a list and the order of presentation was pseudo-randomised three times (no more than two target items occurred consecutively), resulting in three lists. The order of presentation of the practice items was also randomised three times.

In the Full Context Condition (first part of the experiment), the items were presented in a context of several words. In the Minimal Context Condition (second part of the experiment), we presented only the phoneme corresponding to the digraph *ijlei* (i.e., [ə] for the words with *-(e)lijk* and [ɛi] for all other words) plus two to three surrounding phonemes. For example, for the realisation *eigenlijk* [ɛiχələk], we presented [ələk], and for the realisation *bijna* [bɛina] ('almost'), we presented [bɛin]. For the items without the digraph *ijlei*, a randomly selected portion of signal consisting of two to three phonemes was presented.

4.1.3. Procedure

Participants performed a phoneme-monitoring task, with the target phoneme [ɛi]. For the details of the procedure, see Experiment 1. In contrast to Experiment 1, the participants were instructed orally. Moreover, as in Experiment 1, the confidence rating scale ranged from 1 to 5 (1 = least confident, 5 = most confident), but now participants were also given the option to indicate that they had given the wrong response, by circling "F."

The total experiment consisted of 216 trials: the first 108 trials constituted the Full Context Condition, the second 108 trials constituted the Minimal Context Condition. Both conditions were presented in three blocks of 36 trials, separated by short pauses. Prior to each condition, the corresponding set of practice trials was presented, followed by a short pause. The total duration of the experimental session was approximately 30 min.

4.2. Results and discussion

In total, 4136 target trials were presented (22 target stimuli × 2 types × 2 conditions × 47 participants). We excluded those trials from the analyses for which participants circled "F" on the rating scale form (143 trials, 3.5%). Due to technical failure, another 599 responses (14.5%) were missing. Table 3 summarises the remaining numbers of 'yes'-responses to the matching word forms (both *ijlei* and [ɛi] present) and to the mismatching word forms (*ijlei* but not [ɛi] present), in both Full Context and Minimal Context.

The shaded cells contain the numbers of false positives (i.e., reporting the presence of the phoneme [ɛi] when no [ɛi] was actually present in the acoustic signal). Although we observed surprisingly many false positive responses in Minimal Context, the number of false positive responses was considerably higher in Full Context. A logistic regression analysis with the ratio of the numbers of 'yes'- and 'no'-responses for a given item as the dependent variable, and Type of word form (matching versus mismatching word form) and Context (Full Context versus Minimal Context) as factors yielded significant main effects of Type of word form ($\chi^2(2) = 1120.6$, $p < .0001$) and Context ($\chi^2(1) = 48.0$, $p < .0001$), and, importantly, a significant interaction of Type of word form and Context ($\chi^2(1) = 15.8$, $p < .0001$). Logistic regression analyses on the data for the two Context Conditions separately revealed significant simple effects of Type of word form for both Context Conditions (Full Context: $\chi^2(1) = 396.0$, $p < .0001$; Minimal Context: $\chi^2(1) = 760.0$, $p < .0001$). The pattern in the interaction between Type of word form and Context shows that, when the critical stretches of speech are presented outside their linguistic context, listeners are more or less capable of accurately detecting the phoneme [ɛi]. In contrast, they produce many false positives when the stretches of speech are presented in context. This shows that listeners base their responses at

Table 3

Numbers of 'yes'-responses to matching word forms (containing *ijlei* and [ɛi]) and mismatching word forms (containing *ij* but not [ɛi]) in Full Context and in Minimal Context in Experiment 3

Type of word form	Full Context	Minimal Context
Matching ([ɛi] present)	745	800
Mismatching (no [ɛi] present)	400	241

least partly on the orthographic representations of the words.

In fact, after the experiment, many participants reported that they had been aware of the presence of items that contained a mismatch between the phonological and the orthographic code, and that they had chosen to base their decisions on orthography. They had adopted the conscious decision strategy of reporting the phoneme [ɛi] whenever the orthographic code contained the digraph *ij* or *ei*.

This decision strategy is also reflected in the reaction times, which were not different for the false positive responses and the correct 'yes'-responses in this experiment (Wilcoxon's $W = 157094.5$, $p = .13$, two-tailed). This contrasts with Experiment 1, in which the false positive responses were significantly slower than the correct 'yes'-responses. Apparently, the listeners in Experiment 3 always based their decision on lexical orthographic information, irrespective of stimulus type.

The average confidence scores for the 'yes'-responses in Full Context were significantly lower for the mismatching items (i.e., for the items with *-(e)lijk*) than for the matching items (4.7 and 5.0, respectively; Wilcoxon's $W = 114727.5$, $p < .0001$, two-tailed). This reflects the participants' awareness of the mismatches between orthography and the presented acoustic realisations. In other words, although many participants used an orthography-based strategy, the actual realisations still affected the confidence scores.

The considerable number of false positive responses in Minimal Context suggests that, also in this condition, our listeners recognised the phoneme string [(ə)lək] as the suffix *-(e)lijk*. In other words, also in Minimal Context, the orthographic code (containing *ij*) appears to have occasionally been activated, and to have lead the participants to report the presence of the phoneme [ɛi].

The participants probably applied the orthography-based strategy because the instructions had not mentioned possible mismatches between orthography and acoustic realisations, and the participants were therefore uncertain how to deal with such mismatches. The reason that they consciously relied more on orthography than on the acoustic signal may be that spontaneous speech is relatively fast and poorly intelligible, introducing uncertainty about whether certain segments are present.

5. Experiment 4

In Experiment 1, as opposed to in Experiment 3, there was no indication of the use of a strategy, neither in reaction times nor in participants' remarks. The results of Experiment 3 are therefore not informative about whether the restoration in Experiment 1 may have occurred on the basis of orthography. We carried out a fourth experiment, investigating whether orthography

has an influence even when the instructions are explicit about the orthography–phonology mismatches and how to deal with such mismatches, that is, when it is emphasised that participants' decisions should be based on phonological information. In this experiment, participants were instructed to listen carefully to the acoustic signal, and if they did not hear [ɛi], as would be the case for all mismatching forms, they were to respond 'no,' even if the orthography contained the digraph *ij*. In order to prevent participants from now adopting the strategy of responding 'no' whenever they heard [(ə)lək], without carefully listening to the acoustic signal, we included as fillers a number of catch trials: trials including the suffix *-(e)lijk* (suffix without [ɛi]), but with the phoneme [ɛi] at another position in the word (e.g., *eigenlijk* [ɛiχələk]) or in the context, so that participants would occasionally have to respond 'yes' to trials containing [(ə)lək]. Finally, it was stressed that the participants had to respond as quickly as possible. If orthography is difficult to ignore, we may expect relatively many false positive responses in the Full Context Condition of this experiment, despite the emphasis on listening to the acoustic signal and responding as fast as possible. Restoration of reduced word forms must then indeed be a consequence, at least in part, of the automatic activation of orthographic information.

5.1. Method

5.1.1. Participants

Twenty-one participants, 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, 2, or 3.

5.1.2. Materials

Exactly the same materials were used as in Experiment 3, plus 53 additional items (all taken from the corpus of spontaneous speech). Of these 53 additional items, 28 functioned as catch trials (in the Full Context Condition) in order to prevent participants from adopting the decision strategy "if *-(e)lijk* [(ə)lək], then 'no'-button." These items contained the suffix *-(e)lijk* and the phoneme [ɛi] at some other position in the word or in the context. The other 25 additional items consisted of 21 tokens of *bijvoorbeeld* ('for example'), in which *ij* was realized as [ə], and 4 tokens of *bijzonder* ('special'), with [i] as the first vowel. These items reminded the participants to press the 'yes'-button only if they heard [ɛi]. Two tokens of *bijvoorbeeld* and one token of *bijzonder* were included in the practice set. The additional items were evenly distributed over the trial lists that were used in Experiment 3, with the restriction that if the same word ending in *-(e)lijk* occurred both in a target trial and in a catch trial, the target trial was always presented first (e.g., the word *natuurlijk* occurred

first in a context without [ɛi], and only later in the list in a context with [ɛi]).

As in Experiments 1 and 3, all items were presented in two conditions. In the first part of the experiment (Full Context Condition), the items were presented in a context of several words. In the second part of the experiment (Minimal Context Condition), we presented only the phoneme corresponding to the digraph *ij/lei* plus two or three surrounding phonemes. If neither the word nor the context contained *ij/lei*, we presented a randomly chosen phoneme sequence.

5.1.3. Procedure

Participants performed exactly the same task as in Experiment 3: a phoneme-monitoring task (target phoneme [ɛi]), with every trial followed by a confidence rating. For the details of the procedure, see Experiment 3. Again, the participants were instructed orally, but now special emphasis was laid on listening carefully to the acoustic signal and absolutely ignoring the orthography, and on responding as quickly as possible. Participants were provided with examples of words containing the digraph *ij/lei* without containing the phoneme [ɛi], such as *bijzonder* ([bizɔndəʀ]) and *bijvoorbeeld* ([bəvo:ʀ be:lt]). Moreover, one of the examples was a word containing the suffix *-(e)lijk*. Participants were specifically instructed to respond 'no' to such words.

The total experiment consisted of 316 trials: The first 158 trials constituted the Full Context Condition, the second 158 trials constituted the Minimal Context Condition. Both conditions were presented in three blocks of approximately equal length, separated by short pauses. Prior to both the Full Context Condition and the Minimal Context Condition, the set of corresponding practice trials was presented, followed by a short pause. The total duration of the experimental session was approximately 45 min.

5.2. Results and discussion

In total, 1848 target trials were presented (22 target stimuli \times 2 types \times 2 conditions \times 21 participants). Due to technical failure, for one participant the data of the Minimal Context Condition were missing, and for another participant the data of the Full Context Condition were missing. Distributed over the other participants, another 26 responses were missing. This amounted to 114 missing trials (6.2%). Furthermore, we excluded those trials from the analyses for which participants circled "F" on the rating scale form (100 trials, 5.4%). Table 4 summarises the remaining numbers of 'yes'-responses to the matching word forms (both *ij/lei* and [ɛi] present) and to the mismatching word forms (*ij/lei* but not [ɛi] present), in both Full Context and Minimal Context.

The shaded cells contain the numbers of false positive responses. Although the overall proportion of false

Table 4

Numbers of 'yes'-responses to matching word forms (containing *ij/lei* and [ɛi]) and mismatching word forms (containing *ij* but not [ɛi]) in Full Context and in Minimal Context in Experiment 4

Type of word form	Full Context	Minimal Context
Matching ([ɛi] present)	366	389
Mismatching (no [ɛi] present)	53	19

positive responses was considerably lower than in Experiments 1 and 3, the number of false positive responses in Full Context was again higher than in Minimal Context. A logistic regression analysis with the ratio of the numbers of 'yes'- and 'no'-responses for a given item as the dependent variable, and Type of word form (matching versus mismatching word form) and Context (Full Context versus Minimal Context) as factors yielded a significant main effect of Type of word form ($\chi^2(2) = 1268.8$, $p < .0001$), a significant main effect of Context ($\chi^2(1) = 17.8$, $p < .0001$), and a significant interaction of Type of word form and Context ($\chi^2(1) = 4.4$, $p < .05$). Logistic regression analyses on the data for the two Context Conditions separately revealed significant simple effects of Type of word form for both Context Conditions (Full Context: $\chi^2(1) = 569.6$, $p < .0001$; Minimal Context: $\chi^2(1) = 717.3$, $p < .0001$). The interaction pattern suggests that, in Full Context, orthographic codes are activated automatically, and that, despite specific instructions to ignore the spelling of the words, participants' responses are still partly based on these orthographic codes.

A closer look at the response data in Experiment 4 shows that the processes underlying the overt response behaviour may not have been the same for all participants in this experiment. Only three participants were responsible for 65% of all false positive responses. The remaining 35% of the false positive responses were more or less equally distributed over the other participants. We ran separate analyses of the confidence scores and the reaction times for these two groups of participants. These analyses showed that, for the three participants with relatively many false positive responses, the pattern in the reaction times was similar to that in Experiment 3 (no difference in reaction times between hits and false positives in Full Context; Wilcoxon's $W = 708$, $p = .25$, two-tailed), while the pattern in the confidence ratings was similar to that in Experiment 1 (no difference in confidence ratings between the false positives and the hits in Full Context; Wilcoxon's $W = 865$, $p = .75$, two-tailed). For the other participants, we obtained the opposite result. Their pattern in the reaction times was similar to that in Experiment 1 (slower reaction times—537 ms on average—for false positives than for hits in Full Context; Wilcoxon's $W = 1253$, $p < .01$, two-tailed), while their pattern in the confidence ratings was similar to that in Experiment 3 (lower confidence scores for the false positives than for the hits—4.1 and 5.0,

respectively—in Full Context; Wilcoxon's $W = 1473.5$, $p < .0001$, two-tailed).

We think that the three exceptional participants applied the same strategy as the participants in Experiment 3, that is, scanning the orthographic code for *ij/ei*, although they were explicitly instructed not to do so. They were more confident in their false positive responses than the participants in Experiment 3, basically because they ignored the instructions altogether.

The pattern in the reaction time data for the other participants in Experiment 4 suggests that the false positive responses by these participants were not the result of a conscious decision strategy, but, instead, that they were the result of unconscious restoration. This restoration must necessarily have taken place on the basis of orthographic information. Apparently, orthography plays a role in the restoration of reduced word forms: when listening to spontaneous speech in which large differences exist between orthography and phonology, orthography cannot be ignored.

While the on-line responses of these participants were indicative of the occurrence of restoration, the off-line confidence ratings suggest that, in retrospect, participants have been aware of the mismatching nature of the items that triggered the false positives (contrary to in Experiment 1). This is easily explained by the fact that the instructions for Experiment 4 were explicitly aimed at making the participants aware of the orthography–phonology mismatch.

To conclude, in the case of a mismatch between orthography and phonology, receiving only very general instructions (Experiment 3) leads to great uncertainty (as shown by the confidence scores and the participants' remarks), and, as a result, to the use of a conscious decision strategy (as suggested by the reaction times and reported by the participants). Receiving adequate, specific instructions (Experiment 4) reduces uncertainty about the task, but the mismatching nature of the items eliciting the false positives still affects the (off-line) confidence scores. The relatively small number of false positives in Experiment 4 as compared to Experiment 1 suggests that there is a role for orthography, but that the phonemic restoration of reduced word forms, as it took place in Experiment 1, occurs mainly on the basis of phonological information.

6. General discussion

In the first two experiments of this study, we investigated the processing of highly reduced word forms, as they occur in casual spoken Dutch. In these reduced forms, suffixes may be either completely or partly missing. Highly reduced word forms are not recognised when they are presented in isolation, whereas they are recognised when they occur in their natural context (Ernestus

et al., 2002). A possible explanation for this finding may be that when highly reduced forms are recognised in their natural context, restoration of partly or completely missing suffixes takes place. The results of our first two experiments confirm this hypothesis. Listeners report the presence of a missing phoneme [l] in reduced forms, but only when these forms are presented in a context of several words, that is, only when the reduced forms can be recognised. When the critical stretches of phonemes are presented in isolation, listeners accurately discriminate items containing [l] and items not containing [l].

In the classic phonemic restoration studies, restoration has been shown to occur for carefully realised stimuli in which one or more phonemes have been replaced by a (spectrally resembling) sound. Our experiments show that partly or completely missing suffixes are restored in naturally reduced speech, when the missing phonemes are not replaced by another sound and a large portion of the word is missing. This shows that restoration is a natural, highly frequent process.

Experiments 3 and 4 investigated the nature of the lexically provided information that the restoration of reduced word forms is based on. Does restoration occur on the basis of phonology, orthography, or both? The results of Experiments 3 and 4 suggest that orthographic information plays a role: When presented with stimuli with mismatching orthography and phonology, and when explicitly asked to ignore orthography, listeners still occasionally base their responses on the orthographic information. We conclude that orthography is activated automatically when a word is recognised in a phoneme-monitoring task, and that it is difficult to ignore. This conclusion is in line with other evidence that phonological processing may be influenced by orthographic representations (Dijkstra et al., 1995; Donnenwerth-Nolan et al., 1981; Hallé et al., 2000; Seidenberg & Tanenhaus, 1979; Taft & Hambly, 1985; Treiman & Cassar, 1997). The relatively small number of false positive responses in Experiment 4 as compared to Experiment 1 suggests that, although orthography may play a role, phonology is the main source of information that the restoration in Experiment 1 was based on.

To conclude, previous research has shown that listeners cannot recognise highly reduced word forms when they are presented in isolation (Ernestus et al., 2002). The present study sheds more light on this issue, by investigating the processing of casually realised word forms, in which the suffix *-(e)lijk* is either partly or completely reduced. By means of a traditional phoneme-monitoring task, we showed that, when these reduced word forms are presented in context, the suffixes that are missing in these forms are restored: listeners 'hear' the suffixes that are missing. The conscious percept is based on the activated canonical representation, not so much on (a pre-lexical representation of) the acoustic signal itself: the activated representations in the lexicon determine what we think

we hear. In isolation, restoration does not occur. As a consequence, listeners may not recognise reduced word forms when presented in isolation.

The restoration phenomenon shows that the reduced forms are linked to the canonical representations in the mental lexicon, and that the canonical, non-reduced representations, as well as the corresponding orthographic representations, are highly activated upon hearing reduced forms. This finding has implications for episodic models assuming that all surface forms of words are stored in the mental lexicon (e.g., Goldinger, 1998). Even though, in such an architecture, lexical access is mediated by a representation of a reduced form, the non-reduced word form reaches the highest level of activation. It overrules the activation of the representation of the actual reduced form, which results in restoration: The phoneme-monitoring response cannot be executed on the basis of the reduced representation. Activation of the canonical representation is consistent with models that assume that all words are represented by just one representation in the lexicon. These models, however, face the challenge of how to map highly reduced word forms such as [tyk] onto the non-reduced canonical representation ([natyɾlək]—*natuurlijk*).

Restoration enables listeners to understand spontaneous speech, without comprehension being hampered at a conscious level by the drastic reductions inherent to this type of speech. The restoration of highly reduced word forms is partly based on orthographic information, but phonological information appears to be most important.

Appendix A. Non-reduced and reduced variants of the target words in Experiment 1

Word	Non-reduced form	Reduced form
<i>duidelijk</i> ‘clear(ly)’	[dœydlək]	[dœyɾə]
<i>eerlijk</i> ‘honest(ly)’	[ɪ:lək [⌈]]	[ɪ:]
<i>eigenlijk</i> ‘actual(ly)’	[eɪχələk [⌈]]	[eɪχək]
<i>mogelijk</i> ‘possible/ possibly’	[moχələk]	[moχə g [⌈]]
<i>namelijk</i> ‘namely’	[namələk]	[nam g]
<i>natuurlijk</i> ‘of course’	[nətylək]	[tyk]
<i>onmiddellijk</i> ‘immediate(ly)’	[ɔ̃mɪdələk]	[ɔ̃mɪdək [⌈]]
<i>uiteindelijk</i> ‘eventual(ly)’	[œyɛilə g [⌈]]	[œyɛɪŋ]
<i>verschrikkelijk</i> ‘terrible/terribly’	[fəsɾ ɪklək]	[fsɾk]
<i>vreselijk</i> ‘terrible/terribly’	[frɛslək]	[frɛsk]
<i>waarschijnlijk</i> ‘probable/probably’	[v ^o sχɛlək [⌈]]	[fsχɛŋk [⌈]]

Appendix B. Target items in Experiment 3

Twenty-two items with the suffix *-(e)lijk* [(ə)lək] (ij in the orthographic code, but no [ɛi] in the phonological code): *afhankelijk* ‘dependent,’ *afzonderlijk* ‘separate,’ *behoorlijk* ‘decent,’ *duidelijk* ‘clear,’ *eerlijk* ‘honest,’ *heerlijk* ‘delightful,’ *makkelijk* ‘easy,’ *moeilijk* ‘difficult,’ *mogelijk* ‘possible,’ *nadrukkelijk* ‘emphatic,’ *namelijk* ‘namely,’ *natuurlijk* ‘of course,’ *onmiddellijk* ‘immediate,’ *onredelijk* ‘unreasonable,’ *ontzaglijk* ‘immense,’ *oostelijk* ‘eastern,’ *persoonlijk* ‘personal,’ *redelijk* ‘reasonable,’ *verantwoordelijk* ‘responsible,’ *verschrikkelijk* ‘awful,’ *vreselijk* ‘awful,’ *vrolijk* ‘cheerful.’

Twenty-two items with *eilij* in the orthographic code and [ɛi] in the phonological code: *beide* ‘both,’ *bijna* ‘almost,’ *blij* ‘happy,’ *blijkt* ‘appears,’ *blijven* ‘to stay,’ *eigen* ‘own,’ *einde* ‘end,’ *kijken* ‘to look,’ *klein* ‘small,’ *kwijt* ‘lost,’ *lijkt* ‘seems,’ *partij* ‘party,’ *pijn* ‘pain,’ *prijs* ‘price,’ *slijten* ‘to wear out,’ *termijn* ‘term,’ *trein* ‘train,’ *vrij* ‘free,’ *weinig* ‘few,’ *wijk* ‘neighbourhood,’ *zei* ‘said,’ *zijn* ‘to be.’

References

- Cutler, A., Treiman, R., & Van Ooijen, B. (1998). Orthografiekonkistentie ephexkts in foneme detektion? *Proceedings of the Fifth International Conference on Spoken Language Processing*, 6, 2783–2786.
- Dijkstra, T., Roelofs, A., & Fieuws, S. (1995). Orthographic effects on phoneme monitoring. *Canadian Journal of Experimental Psychology*, 49, 264–271.
- Donnenwerth-Nolan, S., Tanenhaus, M., & Seidenberg, M. (1981). Multiple code activation in word recognition: Evidence from rhyme monitoring. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 170–180.
- Ernestus, M. (2000). *Voice assimilation and segment reduction in casual Dutch: A corpus-based study of the phonology–phonetics interface*. Utrecht: LOT.
- Ernestus, M., Baayen, H., & Schreuder, R. (2002). The recognition of reduced word forms. *Brain and Language*, 81, 162–173.
- Goldinger, S. D. (1998). Echoes of echoes. An episodic theory of lexical access. *Psychological Review*, 105, 251–279.
- Hallé, P. A., Chéreau, C., & Segui, J. (2000). Where is the /b/ in absurde [apsyɾd]? It is in French listeners’ mind. *Journal of Memory and Language*, 43, 618–639.
- Samuel, A. G. (1981a). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, 110, 474–494.
- Samuel, A. G. (1981b). The role of bottom-up confirmation in the phonemic restoration illusion. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 1124–1131.
- Samuel, A. G. (1987). Lexical uniqueness effects on phonemic restoration. *Journal of Memory and Language*, 26, 36–56.
- Samuel, A. G. (1996a). Phoneme restoration. *Language and Cognitive Processes*, 11, 647–653.
- Samuel, A. G. (1996b). Does lexical information influence the perceptual restoration of phonemes? *Journal of Experimental Psychology: General*, 125, 28–51.
- Seidenberg, M., & Tanenhaus, M. (1979). Orthographic effects on rhyme monitoring. *Journal of Experimental*

- Psychology: Human Learning and Memory*, 5, 546–554.
- Taft, M., & Hambly, G. (1985). The influence of orthography on phonological representations in the lexicon. *Journal of Memory and Language*, 24, 320–335.
- Treiman, R., & Cassar, M. (1997). Can children and adults focus on sound as opposed to spelling in a phoneme counting task? *Developmental Psychology*, 33, 771–780.
- Warren, R. M. (1970). Perceptual restoration of missing speech sounds. *Science*, 167, 392–393.