

Semantic context effects in the comprehension of reduced pronunciation variants

Marco van de Ven · Benjamin V. Tucker ·
Mirjam Ernestus

© Psychonomic Society, Inc. 2011

Abstract Listeners require context to understand the highly reduced words that occur in casual speech. The present study reports four auditory lexical decision experiments in which the role of semantic context in the comprehension of reduced versus unreduced speech was investigated. [Experiments 1](#) and [2](#) showed semantic priming for combinations of unreduced, but not reduced, primes and low-frequency targets. In [Experiment 3](#), we crossed the reduction of the prime with the reduction of the target. Results showed no semantic priming from reduced primes, regardless of the reduction of the targets. Finally, [Experiment 4](#) showed that reduced and unreduced primes facilitate upcoming low-frequency related words equally if the interstimulus interval is extended. These results suggest that semantically related words need more time to be recognized after reduced primes, but once reduced primes have been fully (semantically) processed, these primes can facilitate the recognition of upcoming words as well as do unreduced primes.

Keywords Acoustic reduction · Word recognition · Speech perception · Semantic priming · Latent semantic analysis · English

Introduction

One important characteristic of spontaneous speech is that many word tokens are much shorter than their corresponding citation forms (e.g., Ernestus, 2000; Johnson, 2004). For example, the English words *yesterday* and *ordinary* can be pronounced like [ˈjeʃeɪ] and [ˈɔɹnəri]. Segments may be shorter, completely missing, or realized differently. This may lead to ambiguity, since, for example, the distinction between long and short vowels and between voiced and voiceless stops may be smaller in spontaneous speech than in careful speech. Previous research has shown that listeners need contextual information (e.g., semantic or acoustic information) to understand highly reduced pronunciation variants (e.g., Ernestus, Baayen & Schreuder, 2002; Kems, Ernestus, Schreuder & Baayen, 2004; van de Ven, Ernestus & Schreuder, 2011), but it is so far unknown what types of contextual information listeners rely on and to what extent. The present study investigated how semantic context contributes to the recognition of mildly reduced and unreduced variants.

Ernestus et al. (2002) extracted pronunciation variants from the Ernestus Corpus of Spontaneous Dutch (Ernestus, 2000). Tokens were classified as having a low degree of reduction if hardly any or no segments were missing. Tokens were classified as having a medium degree of reduction if they were reduced but consisted of more than the initial, final, and stressed segments. The remaining words were classified as having a high degree of reduction. Participants listened to these variants in isolation, within

M. van de Ven (✉)
Max Planck Institute for Psycholinguistics,
Nijmegen, the Netherlands
e-mail: Marco.vandeVen@mpi.nl

B. V. Tucker
University of Alberta,
Edmonton, Canada

M. Ernestus
Radboud University Nijmegen,
Nijmegen, the Netherlands

M. Ernestus
Max Planck Institute for Psycholinguistics,
Nijmegen, the Netherlands

their phonological context (i.e., together with adjacent vowels and any intervening consonants) or within their sentential context. The participants' task was to orthographically transcribe the speech fragments. The results showed that words of a low degree of reduction were well recognized in all three context conditions. Words of a medium or high degree of reduction, however, were recognized only if presented within at least the phonological context (for words with a medium degree of reduction) or the sentential context (for words with a high degree of reduction). Thus, listeners require contextual support to recognize highly reduced pronunciation variants.

A study by Kemps et al. (2004) suggests that listeners reconstruct missing speech sounds when they hear reduced pronunciation variants in context. In a phoneme-monitoring experiment, Dutch participants were presented with target words ending in the derivational suffix *-lijk* [lɔk] (e.g., *eigenlijk* 'actually', *koninklijk* 'royal') extracted from spontaneous speech. They heard canonical realizations of these words and reduced realizations, in which the suffix [lɔk] was produced as [k], and were asked to press a button whenever they heard an [l]. If the suffix was presented in isolation, participants correctly pressed a button only for those variants that contained [l] (i.e., only the unreduced variants). However, when the suffix was presented in sentence context, participants also pressed a button for the reduced variants (i.e., without [l]).

It is unclear what type of contextual information is used by listeners to understand reduced pronunciation variants. Reduced speech may be considered as an adverse listening condition. Some studies predict that, in adverse listening conditions, listeners rely more heavily on any available information, including contextual information (e.g., Hawkins & Smith, 2001). However, more recent work indicates that this is not necessarily the case. For example, the role of semantic contextual information appears marginal in listening to low-pass filtered speech, another type of adverse listening condition (Aydelott & Bates, 2004; Aydelott, Dick & Mills, 2006). Furthermore, research by Andruski, Blumstein and Burton (1994) suggests that semantic context helps English listeners less in their comprehension of obstruents with reduced, as compared with unreduced, voice onset time (VOT) distinctions.

In the present study, we investigated the role of semantic context in the comprehension of reduced pronunciation variants in a series of simple auditory lexical decision experiments with implicit semantic priming. Listeners heard English nouns and pseudonouns, and for each word, they had to make a lexical decision. We examined the effects of a target word's semantic relatedness to the preceding word on the recognition of the target word.

All words were thus produced and presented in isolation, rather than in sentence context, in order to isolate semantic

effects from other higher level information (e.g., syntax and pragmatics). Follow-up research should indicate whether semantic information in the sentence or discourse context influences the recognition of reduced pronunciation variants in the same way as in our experiments.

Previous research has shown that semantic priming effects are strong for words with a low word frequency but are only marginal (or even absent) for words with a high word frequency (e.g., Becker, 1979; Rayner, Ashby & Pollatsek, 2004; Van Petten & Kutas, 1990). Therefore, we predicted that semantic priming effects for reduced speech would be largest for words with a low word frequency as well.

Many studies investigating semantic priming use words that, on the basis of the preceding context, are either highly predictable (e.g., *The opposite of hot is cold*) or unpredictable (e.g., *She read about the flower*; see, e.g., Bradlow & Alexander, 2007; Donnenwerth-Nolan, Tanenhaus & Seidenberg, 1981; Meyer & Schvaneveldt, 1971). In everyday listening situations, however, words are seldom highly predictable. Hence, listeners often need to resort to using subtle semantic information in the context. In the present study, a continuous measure of semantic relatedness was used, and the target words in this study varied from being mildly related to highly related to their preceding words, rather than either semantically highly related or completely unrelated.

We used latent semantic analysis (LSA) to estimate the semantic relatedness of the words (Deerwester, Dumais, Furnas, Landauer & Harshman, 1990). LSA provides a score, ranging from -1 to 1, that indicates to what extent words are semantically related, where a higher LSA score denotes a stronger semantic relatedness. LSA rests on the assumption that semantically related words tend to occur in similar texts. This computational technique uses deep statistical analysis to infer words' semantic relationships beyond their first-order co-occurrences. On the basis of the distributions of words, these words are placed in a multidimensional vector space. LSA scores are obtained by computing the cosine distance between the words' vectors. Previous research has shown that LSA scores can predict human behavior in psycholinguistic experiments—for example, semantic priming in a visual lexical decision task (Landauer & Dumais, 1997).

In the present study, we investigated the role of semantic contextual information in the processing of unreduced and reduced pronunciation variants. We report four auditory lexical decision experiments, in which participants were presented with unreduced and/or reduced isolated words. In **Experiment 1**, participants were presented only with unreduced words, in order to establish the baseline effects of semantic context for listeners presented with clear speech in our experiments.

Experiment 1

Method

Participants Twenty native speakers of English from the participants pool of the Department of Linguistics, University of Alberta, took part in the experiment and received course credit for their participation.

Materials We extracted 154 nouns, with varying word frequencies (range: 40–58,322), from the spoken portion of the Corpus of Contemporary American English (85 million word tokens; Davies, 2008). These nouns were used to construct 77 word pairs (see the [Appendix](#)) that differed in their semantic relatedness. The semantic relatedness of the members of the word pairs ranged from semantically highly related (LSA score: .93; e.g., *saddle–horse*) to mildly related (LSA score: .36; e.g., *snake–beak*). We obtained LSA scores for the word pairs by using the Pairwise Comparison interface at the LSA Website (Landauer, 1998), where we selected the term-to-term comparison type, 300 factors, and, as the topic space, General Reading up to the 1st year of college. Both the LSA scores and the log word frequencies of the second members of the word pairs (i.e., the target words) were normally distributed, and therefore, they could be used as numeric variables in regression analyses. The members of a word pair were presented on consecutive trials, and we investigated the effect of the semantic relatedness of the word pairs on the recognition of the second members of these pairs.

Furthermore, the experiment contained 174 filler words that were semantically unrelated to their preceding and following words and 128 pseudowords. The pseudowords were all phonotactically possible words of English, and two Mann–Whitney tests showed that they had the same number of syllables and segments as the existing English words, on average ($p > .1$ in both cases; mean number of syllables, 1.5 for the pseudowords vs. 1.5 for the existing words; mean number of segments, 5.4 for the pseudowords vs. 5.4 for the existing words). Since we included only a limited number of pseudoword fillers we induced a *yes* response bias, making it particularly difficult to find any priming effects in our data. Furthermore, the many unrelated fillers were included to minimize strategic priming effects. Consequently, any semantic priming effects that showed up were robust effects.

The materials were spoken by a male native speaker of Canadian English, who pronounced the words carefully (mean speech rate: 6.79 segments per second). We presented him with the words in a fully randomized order, such that no words were preceded by semantically related words and the speaker's realizations of the words could not be affected by semantic priming. A different native speaker of Canadian English verified that all the words were

pronounced naturally and clearly. The recordings were made in a sound-attenuated booth at the Alberta Phonetics Laboratory, with an Alesis ML-9600 hard disc recorder and a Countryman E6 directional microphone. The sampling rate was 16 bit/44.1 kHz.

After having extracted the individual words from the recordings, we created three lists, in which the 77 semantically related word pairs, the filler words, and the pseudowords were pseudo-randomized so that no more than six existing words or three pseudowords occurred in succession. Furthermore, we avoided rhyme and/or alliteration between words on consecutive trials. There were, minimally, 6 participants per list.

Procedure The participants were tested individually in a sound-attenuated booth, using E-Prime 2.0 (Schneider, Eschman, & Zuccolotto, 2007) and MBQUARTQP805 Demo headphones. They listened to the stimuli over closed headphones and decided as quickly as possible, for each stimulus, whether it was an existing English word. The next stimulus was presented 1,000 ms after each button press or after a time-out of 4,500 ms from stimulus onset. We selected these timing parameters on the basis of a pilot experiment, which showed the time participants needed to recognize the unreduced materials and the reduced materials used in [Experiment 2](#) and to get ready for the next stimulus. The materials were presented at a comfortable listening level. The experiment lasted approximately 15 min.

Results and discussion

Participants produced 1,523 correct responses, 15 incorrect responses, and two time-outs for the target words (mean response time [RT] from the words' uniqueness points [UPs], excluding the time-outs: 288.04 ms). We analyzed participants' RTs for the correct responses by means of linear mixed-effects models with contrast coding for factors (e.g., Jaeger, 2008). We measured RTs from the words' UPs, which are the segments in the words at which these words diverge from all other words in the language (Marslen-Wilson, 1987). We measured RTs from the words' UPs, rather than from word offsets, because listeners may recognize words before word offset. We preferred the UP over word onset because the unreduced words in this experiment and the reduced words used in [Experiment 2](#) differed in how quickly they become unique and could be recognized. We determined the words' UPs on the basis of CELEX (Baayen, Piepenbrock & Gulikers, 1995).

We restricted all the analyses in this study to the target words, rather than also including the primes and the existing filler words, because the LSA scores with the preceding words were distributed normally only for the target words. Moreover, the primes and existing filler words

were all preceded by semantically unrelated words, and the present study focused on listeners' sensitivities to differences in degrees of semantic relatedness, rather than differences between semantically related and unrelated word pairs.

Furthermore, we restricted all analyses to correct trials directly preceded by other correct trials, excluding those trials for which the data show that the listeners did not recognize the target or the prime (or both). One of the semantically related word pairs (*dance–ballet*) was discarded in our analyses (for all the experiments in this study) because the target word (*ballet*) was not recognized by more than 50% of the participants. We also removed those trials for which participants pressed the button prior to the words' UPs, because, in these cases, participants were probably guessing. In addition, we removed trials for which the RT or the RT for the preceding trial (henceforth, previous RT) was extremely long (>1,500 ms after stimulus onset), thereby removing trials for which the interstimulus interval was very long (>2,500 ms after stimulus onset). We removed these trials because we wished to investigate the effects of the interstimulus interval and compare our results with those from [Experiment 4](#), in which we used a fixed interstimulus interval of 2,500 ms.

The final data set consisted of 1,412 trials. We applied a log transformation to the RTs and to the previous RTs in order to obtain normal distributions, and we analyzed our data by means of a backward stepwise selection procedure, in which predictors and interactions were removed if they did not attain significance at the 5% level.

We included the fixed effect variables of LSA (LSA score) and target word frequency (log word frequency). Furthermore, we included five additional variables mainly to reduce variance in the data set. We included the fixed variables trial number and target word duration (log of the stimulus duration; we took the log of the durations so that the RTs and durations were on the same scale). We also included previous RT (log of the RT on the preceding trial), as an indication of the participants' local response speed. Furthermore, we included the random variables of participant and word.

For all the regression models reported in this study, we excluded data points for which the standardized residuals were smaller than -2.5 or larger than 2.5. We then reran the regression models. A summary of the results is provided in [Table 1](#).

First of all, there were significant main effects for the control variables trial number, previous RT, and target word duration. Participants responded faster toward the end of the experiment, when the preceding RT was short, and to longer words. Since we found similar effects of target word duration if we measured the RTs from word onset or word offset, we interpret this target word duration effect as reflecting how much time listeners had to narrow in on the target word and limit the number of competitors prior to the word's UP.

Table 1 Results for the statistical analysis of the logged response times (RTs) in [Experiment 1](#).

Predictor: Fixed Effects	β	F	p
Intercept	13.294	–	<.0001
Trial number	-0.0004	14.53	<.001
Previous RT	0.257	79.47	<.0001
Target word duration	-1.072	62.05	<.0001
Target word frequency	-0.322	2.37	n.s.
LSA	-3.321	2.06	n.s.
Target word frequency : LSA	0.432	10.11	<.01
Predictor: Random Effects Variance Explained		χ^2	p
Participant	0.053	134.33	<.0001
Word	0.039	51.53	<.0001

Most important for our research question, we found an interaction between LSA and target word frequency. This interaction is shown in [Fig. 1](#). We found semantic priming effects for words with a relatively low word frequency, in line with the literature cited in the Introduction of this article. In addition, we found a semantic interference effect for words in the highest frequency range. Since this interference effect has not been reported in the literature before, we further investigated this effect. We refitted our regression model to a subset of the data consisting of words in the highest range for LSA and/or word frequency and found that this interference effects held for ten word pairs—namely, *borders–country*, *chair–room*, *driver–car*, *game–player*, *gold–silver*, *peace–war*, *peaks–mountain*, *saddle–horse*, *team–coach*, and *tooth–dentist*. The negative main effect of LSA disappeared if we included additional word pairs. Hence, this semantic interference effect was not restricted to two or three tokens obviously sharing some characteristic in our experiment. We will formulate an explanation for this interference effect in the General discussion section of this article.

The question now arose as to whether we could find similar semantic context effects for reduced speech, which is characterized by shorter word durations and missing segments. We addressed this issue in [Experiment 2](#), in which we tested reduced pronunciations of the same words as those we used in [Experiment 1](#).

Experiment 2

Method

Participants Twenty native speakers of English from the same participant pool as that used in [Experiment 1](#) received course credit to take part in the experiment. Participants in [Experiment 2](#) had not participated in [Experiment 1](#).

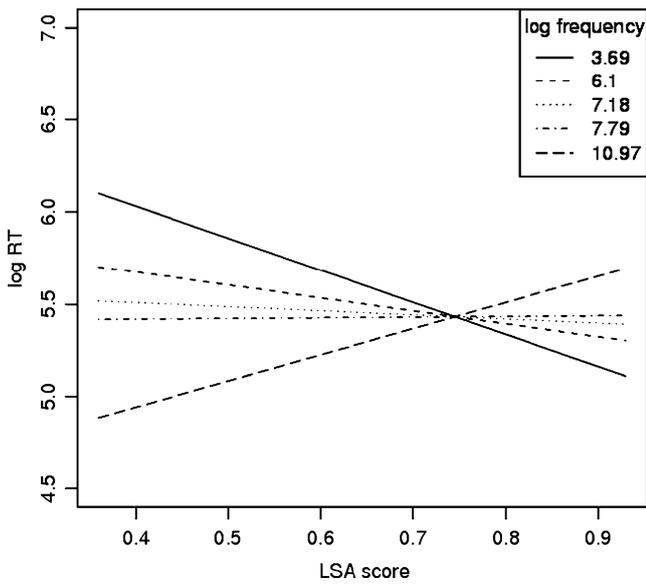


Fig. 1 Combined effects of LSA and target word frequency on log RT for the target words in Experiment 1

Materials We created a new set of recordings of the materials used in Experiment 1. For these new recordings, we asked the same speaker of Canadian English to pronounce the same list of words, but now at a faster speaking rate, in order to elicit more reduced speech. Again, another native speaker of Canadian English verified that the words were produced in a natural manner.

The durations of these reduced realizations were significantly shorter than those of the unreduced realizations tested in Experiment 1, $t(574, 20) = -25.47, p < .0001$ (mean duration, 377.01 ms for the reduced realizations, as compared with 568.10 ms for the unreduced realizations; mean speech rate, 11.4 and 6.8 segments per second for the reduced and unreduced realizations, respectively; see Fig. 2). We also compared the durations of the unreduced and reduced variants by subtracting the duration of the reduced variant from the duration of the unreduced variant and dividing its outcome by the duration of the unreduced variant. The descriptive statistics are provided in Table 2.

The reduced realizations not only were durationally shorter than the unreduced ones, but also contained fewer segments. The reduced and unreduced realizations of each word were phonetically transcribed, and we subtracted the number of segments in the reduced realization (Experiment 2) from the number of segments in the unreduced realization (Experiment 1), dividing its outcome by the number of segments in the unreduced realization. For example, *story* [ˈstɔːri] (five segments) was reduced to [ˈstɔːi] (four segments), and *player* [ˈpleɪər] (six segments) was reduced to [ˈpleɪr] (five segments), resulting in scores of 0.2

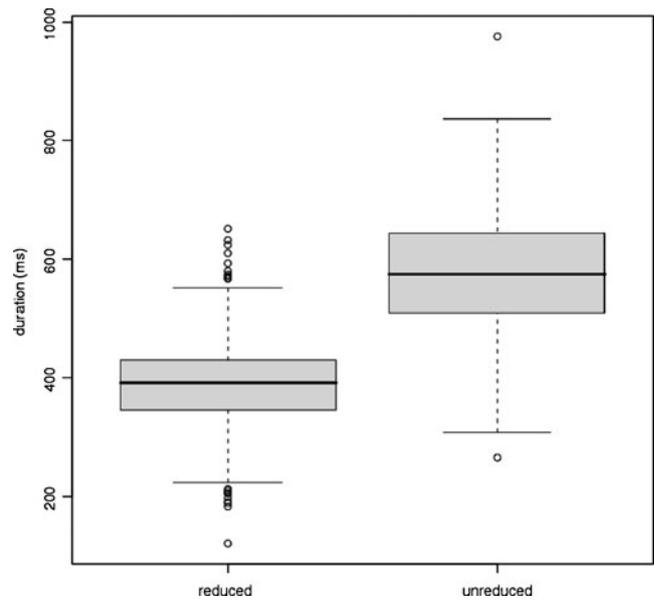


Fig. 2 Word durations for the reduced and unreduced realizations of the stimuli

and 0.17, respectively. The descriptive statistics are reported in Table 2. In most reduced realizations, no segments were completely missing.

We performed Mann–Whitney tests investigating whether there were differences between the degrees of segmental and durational reduction for the primes, targets, existing filler words, and pseudowords. We did not find any differences ($p > .1$ in all cases).

We also tested which segments were typically reduced in our materials. We found that consonants were more frequently missing than vowels (in 18% vs. 2.5% of the words). Most of these missing consonants were plosives (84.51%), followed by approximants (12.68%) and fricatives (2.82%), and consonants were missing, especially in syllable-final position (85.92%). For example, the word *curtains* was realized like [ˈkɜːrtənz] in the unreduced condition and like [ˈkɜːrtns] in the reduced condition. Spectrograms and transcriptions of the unreduced and reduced realizations of this word are provided in Fig. 3. Please note that our materials were only mildly reduced, as compared with the most extremely reduced pronunciations

Table 2 Descriptive statistics for the degree of durational and segmental reduction of the stimuli

Stimulus Type	Average Durational Reduction	Range	Average Segmental Reduction	Range
Target	0.09	0.02–0.18	0.05	0–0.33
Prime	0.09	0.04–0.16	0.03	0–0.40
Filler	0.09	0.02–0.27	0.04	0–0.40
Pseudoword	0.08	0.03–0.16	0.02	0–0.25

that can be found in spontaneous speech and that can be recognized only within their linguistic context. Since the reduced words in our study were produced in isolation, they were only mildly reduced, and they could be recognized in isolation.

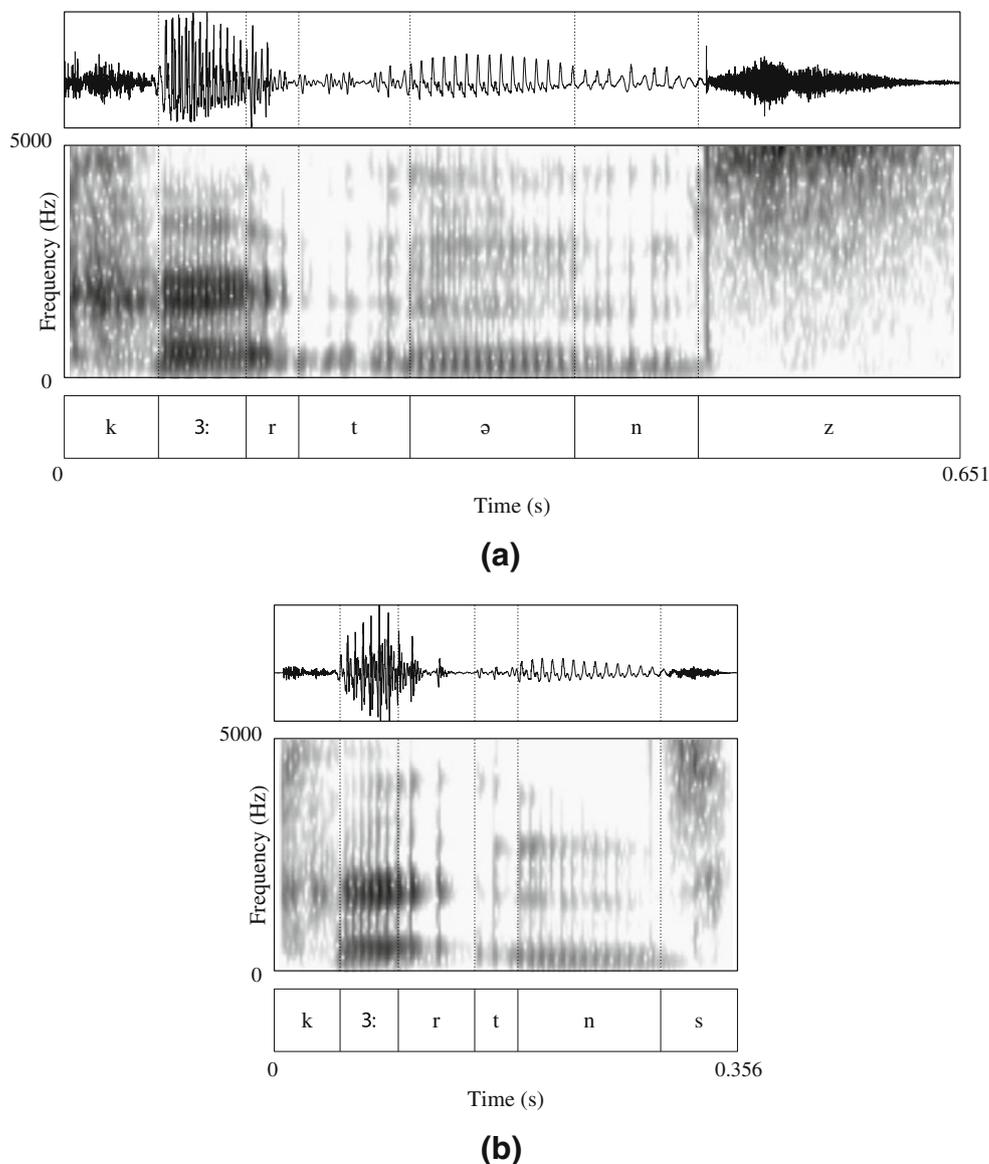
Procedure The procedure was identical to the one in Experiment 1.

Results and discussion

Participants produced 1,477 correct responses, 63 incorrect responses, and no time-outs for the target words (mean RT from the words' UPs: 356.59 ms). We analyzed

and compared the number of errors in Experiments 1 and 2 by means of a linear mixed-effects model with the binomial link function (Jaeger, 2008), with the dependent variable correctness (correct/incorrect), including the same random and fixed variables as in the analyses for the RTs in Experiment 1, in addition to the fixed variable register (indicating whether the primes and targets were reduced or unreduced). We found an effect only for register: Participants in Experiment 2 produced more incorrect responses than did the participants in Experiment 1 (4% of the reduced target words vs. 1% of the unreduced target words) $\beta = -1.255$, $F(1, 2934) = 18.44$, $p < .0001$. This finding suggests that although the reduced words in our experiments were produced in isolation and were mildly reduced, they were nevertheless more difficult to understand in isolation than the unreduced variants.

Fig. 3 Spectrograms and transcriptions for the unreduced (a) and reduced (b) realizations of the word *curtains*



We analyzed the RTs for the correct responses, using the same predictors as those in [Experiment 1](#), in addition to the fixed variable register. After filtering the data, using the procedure described in the Results and discussion section for [Experiment 1](#), the data set consisted of 1,359 trials. A summary of the results is provided in [Table 3](#).

The control variables showed the same effects as in [Experiment 1](#). More important, we found that, in contrast to [Experiment 1](#), [Experiment 2](#) showed no effects of LSA and only a main effect of target word frequency.

The results of these two experiments show interesting differences between the roles of semantic context in the recognition of reduced and unreduced pronunciation variants. In [Experiment 1](#), we found that a higher semantic relatedness with the previous word facilitated the recognition of words with a low frequency, while it hindered the recognition of words with a very high frequency. In [Experiment 2](#), we did not find any effects of semantic relatedness, although the effect of word frequency was omnipresent. In order to test whether the differences between the two experiments attained statistical significance, we fitted a regression model to the combined data sets of [Experiments 1](#) and [2](#), including the same predictors as those for the separate analyses, in addition to the new predictor register.

In this analysis, we could not include target word duration, because this variable now showed a bimodal distribution (i.e., our reduced stimuli were much shorter than our unreduced stimuli). The previous RT and register variables were highly correlated, and we therefore orthogonalized these two variables. We fitted a simple linear regression model with the dependent variable previous RT and the predictor register, and we used the residuals of this model in our regression analysis, instead of the raw previous RTs. We will report only main effects of and interactions with register. First, we found an interaction between word and register ($\chi^2 = 122.08$, $p < .0001$), indicating that certain words showed larger effects of reduction than did others. Second, we found a three-way interaction between register, LSA, and target word frequency,

Table 3 Results for the statistical analysis of the logged response times (RTs) in [Experiment 2](#)

Predictor: Fixed Effects	β	F	p
Intercept	9.44	–	<.0001
Trial number	-0.0005	40.85	<.0001
Previous RT	0.224	70.28	<.0001
Target word duration	-0.793	39.01	<.0001
Target word frequency	-0.041	4.93	<.05
Predictor: Random Effects Variance Explained		χ^2	p
Participant	0.060	224.81	<.0001
Word	0.045	143.69	<.0001

$\beta = 2.006$, $F(1, 2588) = 4.03$, $p < .05$, which confirms the difference in the effects of semantic context in interaction with word frequency in the two experiments.

[Experiments 1](#) and [2](#) contained either unreduced ([Experiment 1](#)) or reduced ([Experiment 2](#)) pronunciation variants, which meant that the speech register of the prime and target were always identical. As a consequence, we cannot determine whether the differences in the effects of semantic context between the two experiments were due to the speech register of the prime, the target, or both. We conducted a third auditory lexical decision experiment, in which we crossed the register of the prime (reduced or unreduced) with the register of the target word (reduced or unreduced).

Experiment 3

Method

Participants Forty-eight native speakers of English from the same participant pool as that used in [Experiments 1](#) and [2](#) were paid to take part in the experiment. They had not participated in the previous experiments.

Materials We used the materials of both [Experiment 1](#) (i.e., unreduced speech) and [Experiment 2](#) (i.e., reduced speech). We crossed the reduction of the prime with the reduction of the target. We created four versions of each of the three pseudo-randomized lists, such that we had all possible combinations of the speech register of the prime and target for all prime–target pairs (i.e., a reduced prime followed by a reduced target, a reduced prime followed by an unreduced target, an unreduced prime followed by a reduced target, and an unreduced prime followed by an unreduced target). We made sure that all combinations (reduced–reduced, reduced–unreduced, unreduced–reduced, and unreduced–unreduced) occurred equally often in each randomization list (19 pairs for each combination). Moreover, the filler words in each list were equally often unreduced as reduced. There were 4 participants per list.

Procedure The procedure was identical to those in the previous experiments.

Results and discussion

Participants produced 3,600 correct responses, 85 incorrect responses, and 11 time-outs for the target words. In order to analyze the errors produced for the target words, we fitted a linear mixed-effects model with the binomial link function

(Jaeger, 2008), using the same random and fixed variables as for the combined analysis of Experiments 1 and 2, in addition to the variables that are described more extensively below for the analysis of the RTs. Participants produced more incorrect responses for reduced than for unreduced targets (3% of the reduced target words vs. 0.3% for the unreduced target words), $\beta = -2.743, F(1, 3635) = 47.05, p < .0001$.

The mean RTs for the correct responses are provided in Table 4. We again analyzed these RTs by means of linear mixed-effects modeling, including the same predictors as those for the combined analysis of Experiments 1 and 2, except that the register variable was now replaced by register of the target and register of the prime, which indicated the speech registers of the target and of the preceding word (reduced or unreduced), respectively. Instead of the raw previous RTs, we used the residuals of a model predicting previous RT as a function of register of the prime, because these variables were highly correlated. After filtering the data, using the procedure described in the Results and discussion section for Experiment 1, the data set consisted of 3,060 trials. A summary of the statistical results is provided in Table 5.

The two control variables (trial number and previous RT) showed the same effects as in the previous experiments. More interesting, participants responded more quickly to unreduced target words, although the exact effects of register of the target differed across words. This interaction was also found in subsequent analyses, but we will not mention it again and will list it only in the tables. Furthermore, this effect of the register of the target was stronger after unreduced primes.

More important for our research question, we found a three-way interaction between register of the prime, LSA, and target word frequency. In order to interpret this three-way interaction, we split up the data by register of the prime.

We first analyzed the target words with unreduced primes. We included only predictors and interactions that were significant in the regression model shown in Table 5. A summary of the statistical results is provided in Table 6. We found the same main effects and interactions as in Table 5, including an interaction between LSA and target

word frequency. This interaction is shown in Fig. 4. Semantically related primes appear beneficial only for target words with a low word frequency. For target words with a higher frequency, semantic priming effects are marginal, while for the words in the highest frequency range, there is a reverse effect of semantic relatedness. These findings are very similar to our findings in Experiment 1, in which all words, including the primes, were unreduced.

Subsequently, we analyzed the target words with reduced primes. A summary of the statistical results is provided in Table 7. We did not find an interaction between register of the target and previous RT. More important, we did not find any effects of LSA and target word frequency, which suggests that, in contrast to unreduced primes, reduced primes hardly influenced the recognition of upcoming semantically related words.

Experiment 3 thus suggests that the absence of semantic context effects in Experiment 2 was due to the reduction of the primes, rather than the targets. This finding indicates that the semantically related words are hardly activated by reduced primes. Unreduced primes, on the other hand, can influence the recognition of both unreduced and reduced targets.

Since reduced primes influenced the recognition of upcoming words to a lesser extent than did unreduced primes, reduced pronunciation variants appear to be less deeply processed at the point at which participants made their lexical decisions for the following words. Experiment 4 tested whether reduced pronunciation variants are permanently processed less deeply, or whether they are processed as deeply as unreduced variants, but later in time.

Experiment 4

Method

Participants Forty-eight native speakers of English, from the same participant pool as those in the previous experiments, received course credit to take part in the experiment. They had not participated in the previous experiments.

Materials The materials were identical to those in Experiment 3, except that we now used the versions of only two randomization lists, instead of three. As a consequence, there were 6 participants per list.

Procedure The procedure was identical to those in the previous experiments, except that we now used a fixed interstimulus interval of 2,500 ms, which is, on average, 500–600 ms longer than the interstimulus interval for the

Table 4 Mean response times (RTs) measured from the words' uniqueness points (excluding time-outs) for unreduced/reduced target words preceded by unreduced/reduced primes

Reduction of the Target	Reduction of the Prime	Mean RT (ms)
Unreduced	unreduced	270.94
Unreduced	reduced	299.52
Reduced	unreduced	473.41
Reduced	reduced	450.42

Table 5 Results for the statistical analysis of the logged response times (RTs) in [Experiment 3](#)

Predictor: Fixed Effects	β	F	p
Intercept	6.872	–	<.0001
Trial number	-0.001	101.83	<.0001
Previous RT	0.268	230.79	<.0001
Target word frequency	-0.092	1.97	n.s.
LSA	-1.235	2.23	n.s.
Register of the prime (unreduced)	1.400	0.99	n.s.
Register of the target (unreduced)	-0.381	188.71	<.0001
Target word frequency : LSA	0.125	7.15	n.s. ¹
Target word frequency :			
register of the prime (unreduced)	-0.197	0.87	n.s.
LSA: register of the prime (unreduced)	-1.940	2.11	n.s.
Register of the target (unreduced) :			
register of the prime (unreduced)	-0.229	51.08	<.0001
Target word frequency : LSA :			
register of the prime (unreduced)	0.293	13.22	<.001
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.046	280.51	<.0001
Word	0.048	204.29	<.0001
Word : register of the target (unreduced)	0.080	131.64	<.0001

trials analyzed in the previous experiments, which was 1,000 ms (the time between a button press and the onset of the next stimulus) + 1,000 ms (the average RT from stimulus onset for unreduced words) or + 900 ms (the average RT from stimulus onset for reduced words).

Results and discussion

Participants produced 3,543 correct responses, 102 incorrect responses, and 51 time-outs for the target words. We fitted a linear mixed-effects model with the binomial link function (Jaeger, 2008) to analyze the targets, using the same random

and fixed variables as for the analysis of [Experiment 3](#). Again, we observed only a main effect for register of the target: Participants produced more incorrect responses for reduced than for unreduced targets (4% of the reduced target words vs. 1% of the unreduced target words), $\beta = -1.947$, $F(1, 3597) = 69.56$, $p < .0001$.

The mean RTs for the correct responses are provided in [Table 8](#). We again analyzed these RTs by means of linear mixed-effects models, including the same predictors as for the analysis of [Experiment 3](#). After filtering the data, using the procedure described in the Results and discussion section for [Experiment 1](#), the data set consisted of 3,129 trials. A summary of the statistical results is provided in [Table 9](#).

Table 6 Results for the statistical analysis of the logged response times (RTs) for target words preceded by unreduced primes in [Experiment 3](#)

Predictor: Fixed Effects	β	F	p
Intercept	6.905	–	<.0001
Trial number	-0.001	47.53	<.0001
Previous RT	0.275	149.18	<.0001
Target word frequency	-0.318	2.66	n.s.
LSA	-3.474	0.96	n.s.
Register of the target (unreduced)	-0.597	218.09	<.0001
Target word frequency : LSA	0.460	13.58	<.001
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.041	86.00	<.0001
Word	0.056	77.03	<.0001
Word: register of the target (unreduced)	0.083	42.58	<.0001

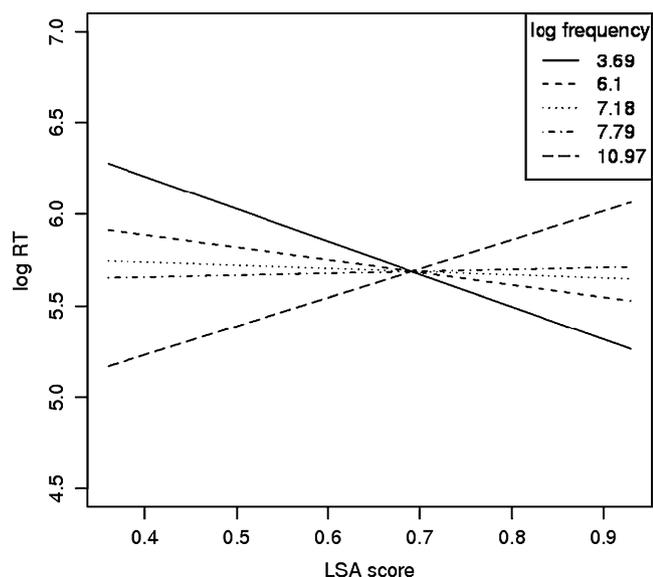


Fig. 4 The combined effects of LSA and target word frequency on log RT for the target words preceded by unreduced primes in Experiment 3

Importantly for our research question, we obtained two interactions with LSA. We found a two-way interaction between LSA and target word frequency: While a higher LSA score elicited faster responses to low-frequency words, we found inhibition for words in the highest word frequency range (see Fig. 5). In contrast to Experiment 3, this interaction between semantic relatedness and lexical frequency was significant for both targets preceded by reduced and unreduced primes (rather than only for targets preceded by unreduced primes).

In addition, we found a three-way interaction between LSA, register of the prime, and register of the target. In order to interpret this three-way interaction, we split the data by register of the prime.

Table 7 Results for the statistical analysis of the logged response times (RTs) for target words preceded by reduced primes in Experiment 3

Predictor: Fixed Effects	β	F	p
Intercept	4.088	–	<.0001
Trial number	-0.001	57.55	<.0001
Previous RT	0.313	95.63	<.0001
Register of the target (unreduced)	-0.372	94.53	<.0001
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.048	125.82	<.0001
Word	0.046	92.01	<.0001
Word : register of the target (unreduced)	0.072	44.65	<.0001

Table 8 Mean response times (RTs) measured from the words' uniqueness points (excluding time-outs) for unreduced/reduced target words preceded by unreduced/reduced primes

Reduction of the Target	Reduction of the Prime	Mean RT (ms)
Unreduced	unreduced	201.49
Unreduced	reduced	223.13
Reduced	unreduced	388.23
Reduced	reduced	363.83

First, we analyzed the targets preceded by unreduced primes. We included only predictors and interactions that were significant in the regression model shown in Table 9. A summary of the statistical results is provided in Table 10. All the main effects and interactions in Table 9 remained significant, including the interaction between register of the target and LSA, depicted in Fig. 6. We show in these figures the RTs for words of a low frequency (the minimum word frequency of our target words), an intermediate frequency (the mean word frequency of our target words), and a high frequency (the maximum word frequency of our target words) separately, since we know that the effect of LSA is modulated by target word frequency. This figure shows that there is stronger semantic priming for unreduced than for reduced targets with low or intermediate frequencies, while there is weaker semantic interference for unreduced than for reduced targets in the highest word frequency range. We will come back to this interaction in the General discussion section.

Subsequently, we analyzed the target words preceded by reduced primes. A summary of the statistical results is provided in Table 11. The results are very similar to those of the analysis of the targets preceded by unreduced primes, except that we did not find an interaction between LSA and register of the target. This suggests that after hearing a reduced prime, there were similar semantic priming effects for unreduced and reduced targets.

Experiments 3 and 4 show some important effects of the interstimulus interval on the semantic context effects induced by unreduced and reduced primes: When processing time was limited (Experiment 3), participants showed semantic context effects, modulated by target word frequency, only for unreduced primes, whereas an extended processing time (Experiment 4) led to such semantic effects for both unreduced and reduced primes. Furthermore, we found that with extended processing time, semantic context effects induced by unreduced primes was smaller for reduced than for unreduced targets. In order to test whether these differences attained statistical significance, we fitted two final regression models, one for targets preceded by unreduced primes and one for targets preceded by reduced primes, to the combined data sets in Experiments 3 and 4. We analyzed targets preceded by unreduced and reduced

Table 9 Results for the statistical analysis of the logged response times (RTs) in Experiment 4

Predictor: Fixed Effects	β	F	p
Intercept	7.309	–	<.0001
Previous RT	0.199	133.17	<.0001
Target word frequency	-0.184	0.50	n.s.
LSA	-2.281	1.23	n.s.
Register of the prime (unreduced)	-0.089	1.09	n.s.
Register of the target (unreduced)	-0.583	144.43	<.0001
LSA : register of the prime (unreduced)	0.243	1.02	n.s.
LSA : register of the target (unreduced)	0.113	0.84	n.s.
Register of the prime (unreduced) : register of the target (unreduced)	0.360	26.59	<.0001
Target word frequency : LSA	0.275	4.68	<.05
LSA : register of the prime (unreduced) : register of the target (unreduced)	-0.811	9.95	<.01
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.030	227.30	<.0001
Participant : register of the target (unreduced)	0.021	62.75	<.0001
Word	0.065	324.26	<.0001
Word : register of the target (unreduced)	0.133	195.10	<.0001

primes separately to simplify the analyses (i.e., to avoid having to test for four-way interactions). We included the same predictors as for the separate analyses, except for register of the prime. Furthermore, we added the fixed variable of experiment (Experiment 3 or 4). We will report only interactions with experiment.

First of all, we analyzed the targets preceded by unreduced primes. We found a main effect of experiment, $\beta = 0.341$, $F(1, 3066) = 5.59$, $p < .05$, and an interaction between experiment and previous RT, $\beta = -0.077$, $F(1, 3066) = 9.19$,

$p < .01$, indicating that listeners responded more slowly in Experiment 4, except when the response to the preceding trial was fast. Furthermore, we found an interaction between experiment, LSA, and register of the target, which indicates that the interaction between LSA and register of the target was present only in Experiment 4. Subsequently, we analyzed the targets preceded by reduced primes. We found a three-way interaction between experiment, LSA, and target word frequency, $\beta = 0.199$, $F(1, 2961) = 5.73$, $p < .05$, indicating that the interaction between LSA and target word frequency was present only after longer processing time.

In summary, our results indicate that acoustically reduced pronunciation variants can induce semantic priming effects (in interaction with lexical frequency), but only if there is more time available to process these reduced variants.

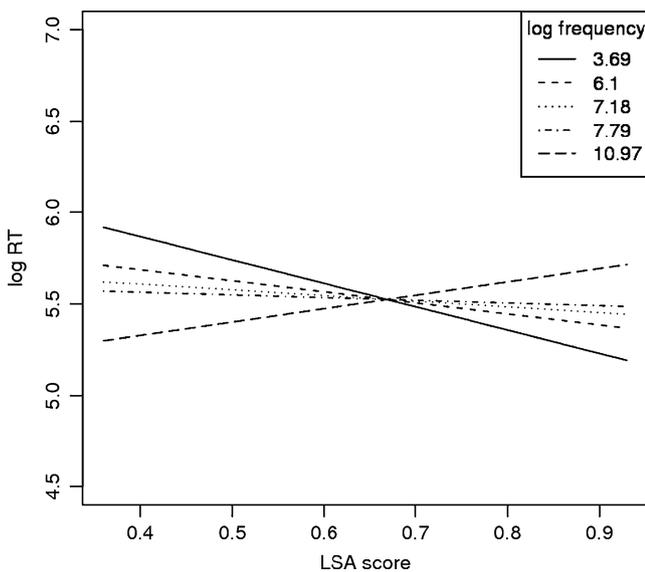


Fig. 5 The combined effects of LSA and target word frequency on log RT in Experiment 4

General discussion

Previous research has shown that listeners need contextual resources to understand reduced pronunciation variants (Ernestus et al., 2002). It is unclear, however, which resources play a role in the understanding of reduced speech. If we assume that reduced speech is an adverse listening condition, previous research has suggested that listeners pay more attention to any type of information, including semantic contextual information, than in the comprehension of clear speech (Hawkins & Smith, 2001). However, research by Aydelott and Bates (2004) and Aydelott et al. (2006) suggested that for one particular adverse listening condition, listening to low-pass filtered

Table 10 Results for the statistical analysis of the logged response times (RTs) for target words preceded by unreduced primes in [Experiment 4](#)

Predictor: Fixed Effects	β	F	p
Intercept	6.689	–	<.0001
Previous RT	0.231	85.67	<.0001
Target word frequency	-0.293	0.63	n.s.
LSA	-3.080	0.85	n.s.
Register of the target (unreduced)	-0.327	438.11	<.0001
Target word frequency : LSA	0.419	6.59	<.05
Register of the target (unreduced) : LSA	-0.477	5.72	<.05
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.023	63.97	<.0001
Participant : register of the target (unreduced)	0.010	11.20	<.001
Word	0.070	156.43	<.0001

speech, the role of semantic contextual information is marginal, although these studies could not ascertain whether listeners actually understood the preceding semantic context. Furthermore, Andruski et al. (1994) demonstrated that smaller VOT distinctions may reduce semantic priming at short interstimulus intervals (50 ms), but not at longer intervals (250 ms), for English listeners. The role of semantic context in adverse listening conditions is thus unknown, including its role in the processing of reduced speech. The present study directly compared the role of semantic (separately from syntactic) information in the processing of unreduced and reduced speech. Since previous research indicated that semantic priming is stronger for words with a lower word frequency (e.g., Becker, 1979; Rayner et al., 2004; Van Petten & Kutas, 1990), we investigated the

interaction between the semantic relatedness of a word with its preceding word and word frequency.

We reported four simple auditory lexical decision experiments with implicit semantic priming in English, in which listeners had to make a lexical decision for each word they heard (which were all nouns). The semantic relatedness of the target words with their primes (presented in immediate succession) ranged from mildly related to highly related (only filler pairs were unrelated), since the present study focused on listeners' sensitivity to differences in the extent to which words are semantically related, instead of differences between semantically highly related and unrelated words. The semantic relatedness of words was estimated by means of LSA (Deerwester et al., 1990). Listeners heard only unreduced pronunciation variants in

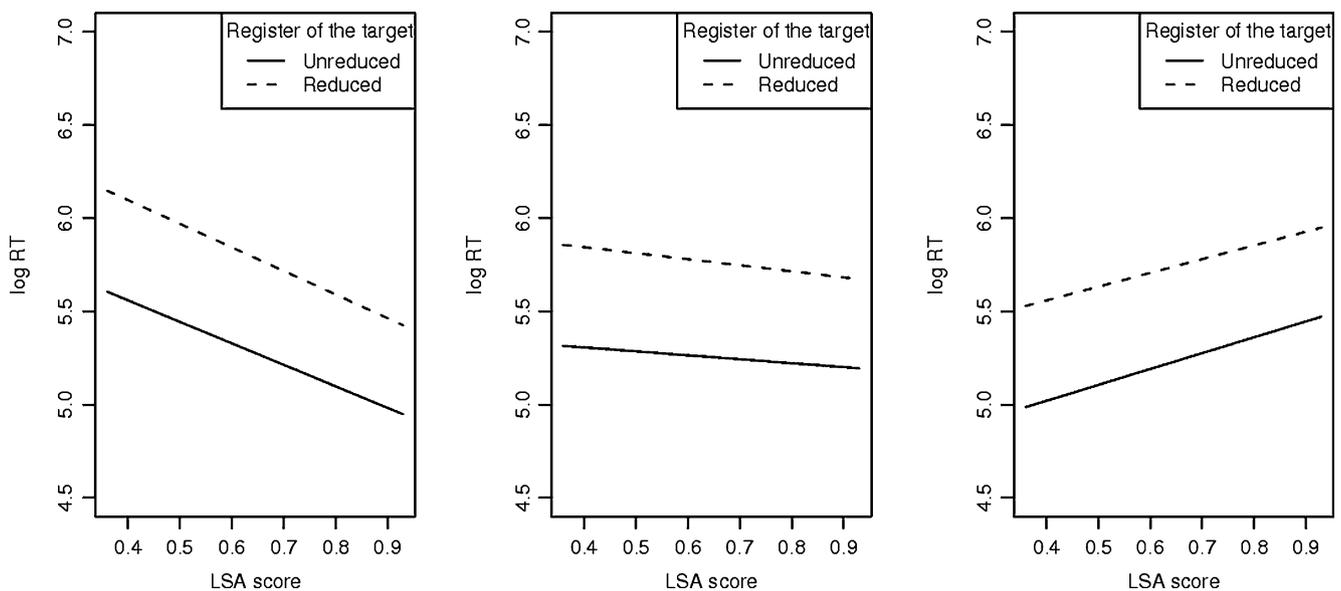


Fig. 6 The effects of LSA on log RT for unreduced and reduced targets with low (left), intermediate (middle), and high (right) word frequencies, preceded by unreduced primes in [Experiment 4](#)

Table 11 Results for the statistical analysis of the response times (RTs) for target words preceded by reduced primes in Experiment 4

Predictor: Fixed Effects	β	F	p
Intercept	5.856	–	<.0001
Previous RT	0.249	59.83	<.0001
Target word frequency	-0.205	0.74	n.s.
LSA	-2.419	0.41	n.s.
Register of the target (unreduced)	-0.460	141.32	<.0001
Target word frequency : LSA	0.314	4.82	<.05
Predictor: Random Effects	Variance Explained	χ^2	p
Participant	0.024	101.50	<.0001
Participant : register of the target (unreduced)	0.041	33.80	<.0001
Word	0.056	109.84	<.0001

Experiment 1, only reduced variants in Experiment 2, and both in Experiments 3 and 4. We recorded and presented all words in isolation (instead of in sentences) in order to carefully control the semantic relatedness of each word with its preceding word and to exclude influences from other higher level information (e.g., syntax and pragmatics). As a natural consequence, the stimuli in our experiments were only mildly reduced as compared with the highly reduced forms that can occur in sentence-medial positions in actual spontaneous speech, and they could be recognized in isolation.

First of all, we investigated, as a baseline, the role of semantic information in the processing of unreduced speech. For words with frequencies ranging from 0.5 to 85 per million (based on Davies, 2008), covering the typical range from low- to high-frequency words according to Carroll (1967), we found the same semantic effects that have been documented before (e.g., Becker, 1979; Rayner et al., 2004; Van Petten & Kutas, 1990)—that is, semantic priming for target words with a low word frequency and no priming for target words with a high word frequency.

Several accounts have been proposed for the facilitatory effects of semantic information in human speech processing. Some researchers have proposed that semantic information facilitates word recognition by allowing listeners to reduce lexical search space (e.g., McClelland & Elman, 1986). In these models, the activation of a word spreads to related words. Other researchers have suggested that the semantic congruency between a word and its preceding context facilitates semantic integration (e.g., Van Petten & Kutas, 1990).

Surprisingly, we obtained very different results for words with frequencies in our highest word frequency range (from 94 to 686.15 per million). These words show semantic interference, which has not been reported in the literature to our knowledge. This interference effect was found each time we investigated the role of semantic

priming after unreduced primes (i.e., in Experiments 1, 3, and 4). Additional analyses showed that these interference effects were not restricted to a couple of atypical word pairs, but they were based on ten different word pairs in our experiments.

The absence of this interference effect in the literature may be explained by the fact that most studies (e.g., Becker, 1979; Rayner et al., 2004; Van Petten & Kutas, 1990) investigated few very high-frequency words that were highly related to their preceding words, since they collapsed over various degrees of semantic relatedness. For example, Becker performed a visual lexical decision experiment with semantically highly related, moderately related, and unrelated word pairs. Importantly, the example provided for the highly related word pairs in this study (*freezing–cold*) is only mildly related (LSA score: .43), as compared with the highly related target word pairs in our study. Hence, that study collapsed over words with mildly related to highly related preceding words, which explains why this study, and most other studies reported in the literature, could not detect the semantic interference effect found in the present study. Studies are necessary that explicitly investigate the comprehension of high-frequency words preceded by semantically highly related words to see whether our results can be replicated using different word pairs—preferably, also in different languages.

The question arises how these interference effects can be integrated into existing psycholinguistic models of speech comprehension. The interference can be explained in connectionist models of speech comprehension such as TRACE (McClelland & Elman, 1986). In these models, high-frequency words have relatively high resting activations, and (as was mentioned above) semantic priming is explained as the spreading of activation of words to semantically related words, which consequently can be recognized more quickly. This implies that, before the recognition of the prime words,

semantically related words with very high resting activation levels may get activation levels that are higher than the activation levels of the prime words, due to this spreading of activation, which would inhibit the recognition of these prime words. These highly frequent words therefore need to be suppressed before listeners can recognize the prime words. Apparently, this suppression can lead to activation levels that are lower than the words' resting activation levels.

Subsequently, we investigated semantic context effects in the processing of reduced speech. In [Experiment 2](#), where listeners were presented only with reduced words, we did not find any semantic context effects. In [Experiment 3](#), we investigated whether these marginalized semantic context effects for reduced speech were due the reduction of the primes, the reduction of the targets, or both. In this experiment, listeners heard both unreduced and reduced words. We found semantic effects after unreduced, but not after reduced, primes, regardless of whether the target was reduced. These results suggest that semantic information in reduced words plays a smaller role, as compared with unreduced words, in the recognition of upcoming related words. Hence, it seems that semantic effects are attenuated if the preceding semantic context is more difficult to process.

These findings raise the question of why semantic priming effects are smaller for reduced primes. On the one hand, reduced pronunciation variants may not activate semantically related words as well as do unreduced pronunciation variants. Alternatively, it may take longer before reduced pronunciation variants activate semantically related words, and semantic effects emerge only if there is ample time to process the reduced variants.

In [Experiment 4](#), listeners had more time (500 or 600 ms longer, on average, depending on the speech register of the prime) to process the prime before we presented the target. We found semantic effects (in interaction with word frequency) from unreduced as well as reduced primes. This suggests that the activation of semantically related words takes more time for reduced speech but that, eventually, semantic effects from reduced speech have a similar magnitude as those from unreduced speech.

Interestingly, the results of [Experiment 4](#) suggest that, after a longer interstimulus interval, unreduced primes show stronger semantic context effects for unreduced targets than for reduced targets. A possible explanation may be that a reduced pronunciation is somewhat unexpected after an unreduced word, since, in natural situations, completely unreduced words, like the ones in our experiment, tend to be surrounded by other clearly pronounced words (e.g., in formal conversations). Reduced words, in contrast, tend to be surrounded by words of all types of reduction degrees. This difference may appear only after a longer processing

time, when listeners have had more time to fully process the primes.

What do these observations reveal about the role of semantic information in the processing of spontaneous speech? They suggest that semantic information plays a role in the processing of both unreduced and reduced speech if there is sufficient time to fully process the words that contain this semantic information, including access to the semantic entries of these words in the lexicon. Since semantically related words (in particular, nouns) are often separated by other words (e.g., function words) in many languages, there is probably sufficient time to fully process reduced variants in most cases. These findings predict that whenever semantically related words occur in immediate succession, these words are less reduced, or listeners resort to different contextual resources. For example, listeners may rely more heavily on acoustic information in the context of reduced words.

The present study serves only as a starting point for establishing the role of semantic context in the processing of spontaneous speech. We tested the recognition of isolated words in the context of other isolated words. Follow-up research is required to establish whether semantic information also influences the processing of reduced words embedded in natural sentences, since it has been suggested that, at least in cross-modal priming, associative priming effects may be smaller under such conditions (e.g., Norris, Cutler, McQueen & Butterfield, 2006).

Furthermore, the present study includes various types of pronunciation variation, including durational reduction and segmental reduction. Previous research by Aydelott-Utman, Blumstein and Burton (2000) suggests that, at least in identity priming, durational reduction inhibits priming, while syllable structure variation (e.g., [pə'lis] versus ['plis] for *police*) may actually enhance priming. More research is required to investigate whether all types of variation influence semantic priming similarly.

In summary, the present study illustrates that the effects of a word's semantic relatedness with its preceding word occur later in the processing of reduced speech than in the processing of unreduced speech. This finding suggests that semantic information in reduced words facilitates the recognition of upcoming words only if there is sufficient processing time. Since reduced words are often separated by other words, listeners probably benefit as much from semantic information in their processing of reduced speech as they do in their processing of unreduced speech.

Acknowledgements A partial analysis of the data from Experiments 1 and 2 has been presented at Interspeech (van de Ven, Tucker & Ernestus, 2009). This research was funded by a European Young Investigator Award to the third author. We wish to thank Robert Schreuder, Iris Hanique, and the three anonymous reviewers for their comments on earlier versions of the manuscript.

Appendix

Table 12 Semantically related word pairs tested in this study and their LSA scores

1. axe–knife .52	40. lock–door .7
2. balloon–airplane .53	41. milk–cheese .82
3. basement–stairs .59	42. month–receipts .63
4. beard–eyebrows .65	43. morning–breakfast .79
5. bees–honey .81	44. music–concert .8
6. blanket–pillow .54	45. nurses–patients .76
7. boat–dock .77	46. ocean–shelf .55
8. book–pages .62	47. paper–ink .72
9. borders–country .51	48. parade–crowds .58
10. bottle–glass .55	49. party–birthday .43
11. candy–chocolate .62	50. peace–war .76
12. cat–mouse .72	51. peaks–mountain .9
13. chair–room .75	52. pie–butter .66
14. church–pope .75	53. pollen–flower .88
15. cloud–drops .65	54. presents–christmas .45
16. coasts–sea .66	55. rabbit–fox .64
17. collar–dress .53	56. river–flood .67
18. cousin–aunt .68	57. saddle–horse .93
19. curtains–window .76	58. school–junior .59
20. dance–ballet .81	59. sister–baby .37
21. dishes–kitchen .75	60. sky–mist .57
22. driver–car .79	61. snake–beak .36
23. farm–crop .72	62. spoon–bowl .6
24. feather–eagle .66	63. storm–lightning .69
25. fire–chimney .63	64. story–plot .78
26. game–player .88	65. sunrise–noon .51
27. gold–silver .88	66. symbols–map .58
28. hook–fish .66	67. teacher–classroom .9
29. horn–bull .4	68. team–coach .89
30. household–wives .57	69. tooth–dentist .89
31. house–porch .65	70. tower–castle .5
32. huts–village .72	71. town–streets .58
33. ice–glaciers .84	72. voice–tone .61
34. idea–topic .64	73. water–bucket .48
35. island–mainland .77	74. wax–plaster .52
36. jacket–shirt .72	75. web–spider .77
37. king–princess .64	76. wheel–gears .75
38. knee–elbow .74	77. wine–beer .85
39. lid–box .57	

References

- Andruski, J., Blumstein, S., & Burton, M. (1994). The effect of subphonetic differences on lexical access. *Cognition*, *52*, 163–187.

- Aydelott, J., & Bates, E. (2004). Effects of acoustic distortion and semantic context on lexical access. *Language and Cognitive Processes*, *19*, 29–56.
- Aydelott, J., Dick, F., & Mills, D. L. (2006). Effects of acoustic distortion and semantic context on event-related potentials to spoken words. *Psychophysiology*, *43*, 454–464.
- Aydelott-Utman, J., Blumstein, S., & Burton, M. (2000). Effects of subphonetic and syllable structure variation on word recognition. *Perception & Psychophysics*, *62*, 1297–1311.
- Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX lexical database (CD-ROM)*. Philadelphia: University of Pennsylvania, Linguistic Data Consortium.
- Becker, C. (1979). Semantic context and word frequency effects in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 252–259.
- Bradlow, A. R., & Alexander, J. A. (2007). Semantic and phonetic enhancements for speech-in-noise recognition by native and non-native listeners. *The Journal of the Acoustical Society of America*, *121*, 2339–2349.
- Carroll, J. B. (1967). On sampling from a lognormal model of word frequency distribution. In H. Kučera & W. N. Francis (Eds.), *Computational analysis of present-day American English* (pp. 406–424). Providence: Brown University Press.
- Davies, M. (2008). *The corpus of contemporary American English (coca): 385 million words, 1990–present* [On-line], Retrieved from <http://www.americancorpus.org>
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, *41*, 391–407.
- 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, R., & Schreuder, R. (2002). The recognition of reduced word forms. *Brain and Language*, *81*, 162–173.
- Hawkins, S., & Smith, R. (2001). Polysp: A polysystemic, phonetically-rich approach to speech understanding. *Italian Journal of Linguistics*, *13*, 99–188.
- Jaeger, F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, *59*, 434–446.
- Johnson, K. (2004). Massive reduction in conversational American English. In *Spontaneous speech: Data and analysis. Proceedings of the 1st session of the 10th international symposium* (pp. 29–54). Tokyo: National International Institute for Japanese Language.
- Kemps, R., Ernestus, M., Schreuder, R., & Baayen, R. (2004). Processing reduced word forms: The suffix restoration effect. *Brain and Language*, *19*, 117–127.
- Landauer, T. (1998). *Lsa@cuboulder* [Online], Retrieved from <http://lsa.colorado.edu>
- Landauer, T., & Dumais, S. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.
- Marslen-Wilson, W. D. (1987). Functional parallelism in spoken word recognition. *Cognition*, *25*, 71–102.
- McClelland, J. L., & Elman, J. (1986). The TRACE model of speech perception. *Cognitive Psychology*, *18*, 1–86.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.

- Norris, D., Cutler, A., McQueen, J., & Butterfield, S. (2006). Phonological and conceptual activation in speech comprehension. *Cognitive Psychology*, *53*, 146–193.
- Rayner, K., Ashby, J., & Pollatsek, A. (2004). The effects of frequency and predictability on eye fixations in reading: Implications for the E-Z reader model. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 720–732.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2007). E-Prime getting started guide. Psychology software tools.
- van de Ven, M., Ernestus, M., & Schreuder, R. (2011). *Contextual influences in the understanding of spontaneous speech*. Manuscript submitted for publication.
- van de Ven, M., Tucker, B., & Ernestus, M. (2009). Semantic context effects in the recognition of acoustically unreduced and reduced words. In *Proceedings of the 10th Annual Conference of the International Speech Communication Association* (pp. 1867–1870). Brighton.
- Van Petten, C., & Kutas, M. (1990). Interactions between sentence context and word frequency in event-related brain potentials. *Memory & Cognition*, *18*, 380–393.