

## PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/16998>

Please be advised that this information was generated on 2020-11-29 and may be subject to change.

# CAN LEXICAL KNOWLEDGE MODULATE PRELEXICAL REPRESENTATIONS OVER TIME?

**James M. McQueen**

Max-Planck-Institute for Psycholinguistics, Nijmegen, The Netherlands

**Dennis Norris**

MRC Cognition and Brain Sciences Unit, Cambridge, U.K.

and **Anne Cutler**

Max-Planck-Institute for Psycholinguistics, Nijmegen, The Netherlands

## ABSTRACT

The results of a study on perceptual learning are reported. Dutch subjects made lexical decisions on a list of words and nonwords. Embedded in the list were either [f]- or [s]-final words in which the final fricative had been replaced by an ambiguous sound, midway between [f] and [s]. One group of listeners heard ambiguous [f]-final Dutch words like [kara?] (based on *karaf*, *carafe*) and unambiguous [s]-final words (e.g., *karkas*, *carcase*). A second group heard the reverse (e.g., ambiguous [karka?] and unambiguous *karaf*). After this training phase, listeners labelled ambiguous fricatives on an [f]-[s] continuum. The subjects who had heard [?] in [f]-final words categorised these fricatives as [f] reliably more often than those who had heard [?] in [s]-final words. These results suggest that speech recognition is dynamic: the system adjusts to the constraints of each particular listening situation. The lexicon can provide this adjustment process with a training signal.

## 1. INTRODUCTION

A battle has been raging in psycholinguistics for many years around the question of feedback. Does lexical knowledge influence earlier levels of speech processing? Models have been proposed in which lexical representations modulate prelexical representations via top-down connections (most notably TRACE [1,2]); other models have been proposed which explicitly rule out this kind of feedback (e.g., the FLMP [3,4] and the Race model [5,6]). Experiments have been carried out which have been taken as demonstrations of the existence of feedback (e.g. [7,8,9,10,11]); others have challenged the notion of feedback (e.g. [6,12,13,14,15]).

In the midst of this battle, however, little attention has been paid to another question: is feedback beneficial? As we have recently argued [16], feedback can do nothing to improve spoken word recognition, at least not if the speech perception system is operating optimally. If earlier stages of processing pass information up to the lexicon in an optimal fashion, and a word is then selected for recognition, all that feedback can do is update the prelexical level so that it agrees with the lexical decision that has already been made. Although feedback could improve phoneme recognition (by filling in missing or ambiguous sounds), it does so at a cost. It acts to overwrite any information at the prelexical level which mismatches with lexical

knowledge. That is, it acts to replace the information which was actually present in the physical speech signal with lexical information. In the worst case, this could lead to a kind of hallucination [3,16].

Arguments such as these led us to develop the Merge model [16], which makes the strong assumption that there is no feedback of information from the lexicon to prelexical processing during spoken word recognition. Although there are many demonstrations that listeners use lexical knowledge in tasks requiring phonemic decisions, Merge explains these effects as arising from the feedforward influence of the lexicon on postlexical decision units. These units continuously receive and integrate information from the prelexical and lexical levels. Merge can explain a wide range of data on lexical involvement in phonemic decisions to both words and nonwords [16]. These effects arise in words because activation of word representations at the lexical level increases the activation of the decision units for the phonemes in those words. Lexical effects in nonwords come about because the representations of words which sound similar to those nonwords will be activated when those nonwords are heard, so again can influence the activation of decision units.

There is, however, one piece of data which may challenge Merge's assumption that there is no feedback. Samuel [11] paired phonemic restoration with selective adaptation. He asked listeners to categorise a [bɪ]-[dɪ] continuum during an adaptation phase in which they heard multiple repetitions of five different words interspersed with categorisation trials. In one condition, the words each contained a [b] and no [d]'s (e.g., *alphabet*), in another, the words each had a [d] but no [b]'s (e.g., *armadillo*). Selective adaptation effects were observed: there were more [b] decisions when the adaptors contained [d]'s than when they contained [b]'s. Crucially, the same effect was observed when the critical phonemes in the adaptor words were replaced with signal-correlated noise. Adaptation occurred as if the listeners had heard the phonemes that had been replaced with noise. Samuel argued that this was the result of feedback of lexical knowledge to the prelexical level, the presumed locus of the adaptation effect.

If adaptation does indeed take place at the prelexical level, these results could be particularly problematic for the Merge model. While Merge has a ready explanation for direct lexical effects on phonemic decision-making, it cannot explain influences of the lexicon on prelexical processes. Given Merge's success in accounting for

many other lexical effects, it is worth asking whether there is anything special about the paradigm used by Samuel [11]. One critical difference between Samuel's experiment and previous studies which have attempted to demonstrate lexical feedback is that Samuel's experiment did not use an on-line measure of the immediate effects of lexical information on phoneme categorisation. In his experiment, the lexical effects are assumed to take place during the adaptation phase. The categorisation data is then used to infer that lexical feedback had altered prelexical processing.

It is quite possible that the adaptation data do not actually reflect an immediate facilitatory effect of lexical feedback at all, but rather a longer-term effect of perceptual learning. Other apparent demonstrations of lexical feedback can be explained by learning processes which do not entail feedback [17]. As we have already argued, top-down feedback cannot actually help prelexical processing. Lexical information, however, could facilitate the process of learning perceptual categories. Consider what might happen if you were to encounter a talker whose productions reflect a different placement of a particular phoneme boundary from your own. It would be easier to understand her if you could retune your perceptual categories. Even though she may produce unusual, and possibly ambiguous, phonemes, it might be possible to use lexical and contextual information to determine which phonemes she intended to produce. Lexically-derived knowledge of the target phoneme could be used to train low-level perceptual processes, that is, it could be used to adjust category boundaries, so as to increase the probability of correctly classifying the spoken input. This kind of feedback would thus be beneficial to the recognition system.

Samuel's [11] data might therefore reflect a perceptual learning process rather than immediate on-line feedback. Listeners in his experiment could be using lexical information to adjust their interpretation of acoustic-phonetic information in response to exposure to the adaptors. Specifically, they could learn that a signal-correlated noise version of a [b], for example, in the context of a word like *alpha\*et*, is an acceptable token of the [b] category. This could produce an adaptation effect (i.e., as if they had heard a real [b]).

This study was an attempt to address this possibility. We exposed listeners to the speech of a talker who appeared to produce exemplars of a phonetic category in an idiosyncratic way, and then tested whether categorisation of a continuum of speech sounds including this category was affected. The idiosyncratic sound was an ambiguous fricative, created by editing natural speech to be midway between [f] and [s]. Exposure to this sound occurred in a list of words and nonwords, to which Dutch listeners made lexical decisions. Embedded in the list were either [f]- or [s]-final words in which the final fricative had been replaced by the ambiguous sound [?]. One group of listeners heard ambiguous [f]-final Dutch words like [kara?] (based on *karaf*, *carafe*) and unambiguous [s]-final words (e.g., *karkas*, *carcase*), as if, for this talker, the

boundary between [f] and [s] had been shifted towards [s]. A second group heard the reverse (e.g., ambiguous [karka?] and unambiguous *karaf*), as if the [f]-[s] boundary had been shifted towards [f]. A control group heard the ambiguous sounds at the end of strings which would be nonwords whether they ended with [f] or [s].

After this training phase, listeners in all three groups were asked to label ambiguous fricatives on the [f]-[s] continuum. We predicted that if listeners use lexical knowledge to adjust their interpretation of ambiguous fricative information, the listeners who had heard [?] in [f]-final words would tend to categorise sounds on the continuum as [f] more often than the control group, and that the listeners who had heard [?] in [s]-final words would tend to label the sounds as [s] more often than the control group. This training experiment is described as Experiment 2 below. In Experiment 1, we selected the ambiguous fricative.

## 2. EXPERIMENT 1

### 2.1. Method

#### 2.1.1. Subjects

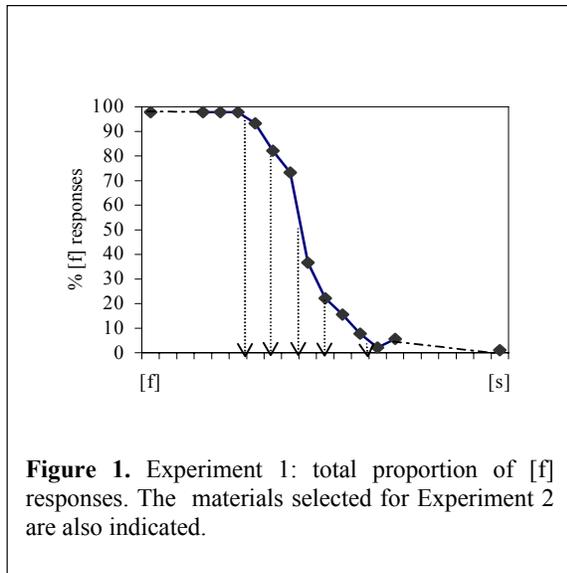
Nine members of the MPI for Psycholinguistics subject panel were paid to take part. They were all native speakers of Dutch, with no known hearing disorders.

#### 2.1.2. Materials and Stimulus Construction

A female native speaker of Dutch recorded a number of tokens of the syllables [ɛf] and [ɛs] in a sound-treated booth onto Digital Audio Tape. The syllables are the names for the letters F and S. This recording was then re-digitized (at 16 kHz) and examined with the Xwaves speech editor. Two friction noises, 272 ms in duration, one [f] and one [s], were excised from the recording, cutting at a zero-crossing at the onset of friction energy. A 21-step continuum was then made by adding the amplitudes of the two waveforms sample by sample in different proportions [14]. The proportions were equally spaced in 21 steps from 0.0 to 1.0 and were added pairwise so as to sum to 1.0. The amplitude of each of these steps was then scaled down to 35% of each original, to increase ambiguity. They were then spliced onto an [ɛ] context (96 ms); this was the vocalic portion of the utterance from which the [f] endpoint was taken, ending at the same splice point. Fourteen steps from this [ɛf]-[ɛs] continuum were then selected for use in the experiment. In addition to the endpoints, the steps were selected from the area of the continuum which informal listening indicated was the most ambiguous: steps 1 ([f]), 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, and 21 ([s]).

#### 2.1.3. Procedure

Each stimulus was presented ten times, in pseudo-random order, such that all stimuli were distributed across the entire running order. Listeners were tested in sound-damped booths. They were presented with a stimulus once every 2.6 seconds at a comfortable listening level over headphones. They were asked to



indicate whether the final sound of each token was [f] or [s] by pressing one of two labelled buttons, "F" or "S". Half the subjects made [s] responses with their dominant hand; the other half made [f] responses with their dominant hand. They were asked to respond on every trial, as fast and as accurately as possible. There was a short pause half way through the main sequence of 140 trials, which was preceded by a short practice block containing one token of each of the 14 stimuli.

## 2.2. Results and Discussion

The total proportion of [f] responses to each member of the 14-step continuum is plotted in Figure 1. The continuum was labelled quite categorically: listeners only found tokens in the range from step 7 to step 13 to be ambiguous. The continuum was most ambiguous between steps 9 and 10. We used these results to select an ambiguous fricative sound for use in the lexical decision phase of Experiment 2, as well as five sounds for use in the categorisation phase.

## 3. EXPERIMENT 2

### 3.1. Method

#### 3.1.1. Subjects

Forty nine new subjects from the MPI subject panel were paid for their participation. They were all native speakers of Dutch, with no known hearing disorders.

#### 3.1.2. Materials and Stimulus Construction

In order to select the most ambiguous fricative from the [ɛf]-[ɛs] continuum, we generated a more fine-grained range of sounds between the same two endpoints as were used in Experiment 1. A new, 42-step continuum was made by adding the same natural [f] and [s] endpoints sample by sample, in 42 equally-spaced steps. Step 18 on this continuum corresponds to half way between steps 9

and 10 on the 21-step continuum. This fricative noise was used as the ambiguous sound [ʔ] in the lexical decision phase. It was also used in the categorisation phase, along with steps 12, 15 (8 in the 21-step continuum), 21 (11 in the 21-step continuum) and 26. As shown in Figure 1, these steps spanned the range from a sound which was likely to be labelled almost always as [f] to one which was likely to be labelled mostly as [s].

Forty Dutch words were selected as experimental items. Twenty contained no [v], [z] or [s] and ended with an [f]. There were five with one syllable (e.g., *druif*, grape), five with two syllables (e.g., *karaf*, carafe), five with three syllables (e.g., *biograaf*, biographer) and five with four syllables (e.g., *locomotief*, locomotive). The other 20 contained no [v], [z] or [f] and ended with an [s]. There were again five of each length (e.g., *muis*, mouse; *karkas*, carcass; *geitkaas*, goat's cheese; and *problematisch*, problematic). The words were paired across sets on their final vowels. Word frequency was matched between the sets as a whole (13 per million for the [f]-words, 14 per million for the [s]-words), and stress patterns were matched between sets. A further 20 nonwords were then constructed which contained no [v], [z], [f] or [s] (five of each of the four lengths). These items were chosen such that they would be nonwords whether they ended with [f] or [s]. This set was again matched in stress patterns to the two sets of words.

One hundred filler words were selected (25 of each length), and 100 filler nonwords were constructed (again, 25 of each length). The sounds [v], [z], [f] and [s] did not occur in any of these items. The nonwords tended to become nonwords (i.e., were no longer consistent with any real Dutch words) before their final phonemes.

All 260 items were recorded by the same talker who recorded the [ɛf]-[ɛs] continuum (in the same recording session). The [f]-final words and fillers were recorded as such, the [s]-final words with a final [f] and a final [s] (e.g., *karkaf* and *karkas*) and the experimental nonwords with a final [f]. The experimental items were examined with Xwaves. [ʔ]-final versions of the [f]-final words were made by replacing the [f] with [ʔ], which was spliced onto the final vowel of each word (at a zero crossing at the onset of frication). [ʔ]-final versions of the [s]-final words were made by replacing the [f] in the [f]-final versions of these words with [ʔ]. [ʔ]-final nonwords were made by replacing the [f] in the original versions with [ʔ]. Thus, for all [ʔ]-final items, the ambiguous fricative was spliced onto vowels taken from natural [f]-final tokens (as indeed were the fricatives in the [ɛf]-[ɛs] continuum). This meant that any cues to place of articulation in the vowels in the [ʔ]-final items consistently cued labiodental place. Individual speech files for 100 experimental items were then made: natural and [ʔ]-final versions of each of the 40 words, and [ʔ]-final versions of the 20 nonwords.

#### 3.1.3. Procedure

Three lists of items were constructed, each comprising 100 words and 100 nonwords. One list contained the [ʔ]-final versions of the 20 [f]-final words, the natural

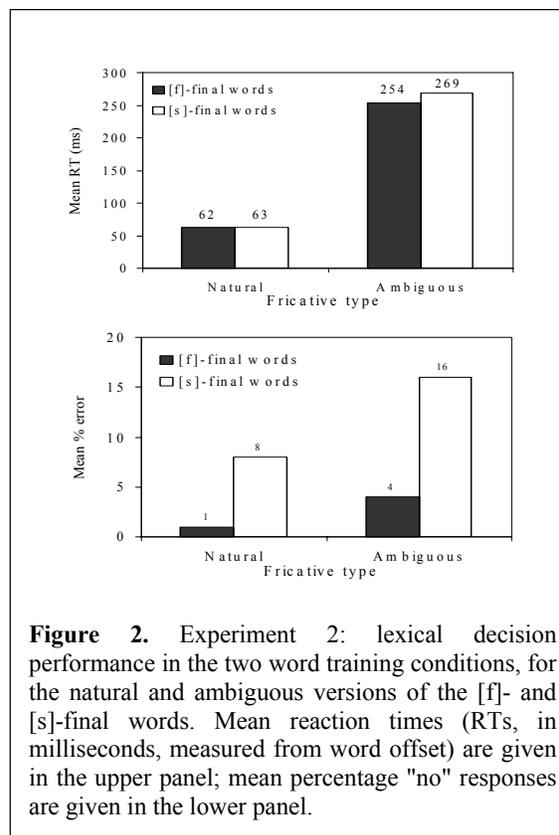
versions of the 20 [s]-final words, 60 filler words (15 of each length), and all the filler nonwords. The second list was identical, except that it contained the [ʔ]-final versions of the [s]-final words (instead of the natural versions) and the natural versions of the [f]-final words (instead of the ambiguous versions). The third list contained the 20 [ʔ]-final experimental nonwords, 80 filler nonwords (20 of each length), and all the filler words. A pseudo-random running order of the two experimental-word lists was then made, such that the order of presentation of all items was identical on the two lists, except for which version of each experimental word was to appear on a given trial. The natural and [ʔ]-final experimental items were spread equally through the list, except that they did not occur in the first 12 trials. There were no more than four words or four nonwords in a row. A second running order was then made for each of these two lists: the same first 12 items in the same order, followed by the remaining items in reverse order. Two versions of the experimental-nonword list were based on these two running orders, with the fillers which were the same as in the other two lists appearing in the same positions and the [ʔ]-final nonwords appearing in the same slots as the experimental words.

Listeners were tested individually in a sound-damped booth. They were given written instructions that they would hear a list of words and nonwords over headphones, and that they were to indicate whether each item was a real Dutch word or a nonsense word by pressing one of two labelled buttons, "JA" (yes) or "NEE" (no). They were asked to respond as fast and as accurately as possible. Reaction times (RTs) were measured from item onset and adjusted so as to measure from item offset prior to analysis. Listeners were not informed about the presence of ambiguous sounds. They were told that there would be a second part to the experiment, but not told what that would be. Each listener then heard the items from one running order of one of the three lists, presented at a comfortable listening level with an inter-onset interval of 2.6 seconds. They were then given written instructions for categorisation (as in Experiment 1), while the experimenter placed "F" and "S" labels on the response buttons. Half the subjects made "S" (and "JA") responses with their dominant hand; the other half had the reverse. The five sounds on the continuum were then each presented six times, in pseudo-random order, as in Experiment 1, except that there was no practice block, and no pause during this test.

### 3.2. Results and Discussion

#### 3.2.1. Lexical Decision

Three of the 17 subjects who heard ambiguous [s]-final and natural [f]-final words (e.g., [karkaʔ] and *karaf*) labelled more than 50% of the [ʔ]-final items as nonwords and were therefore excluded from all further analyses (since they were unwilling to treat the [ʔ]-final items as tokens of [s]-final words, they are unlikely to show the same kind of training effect as the other

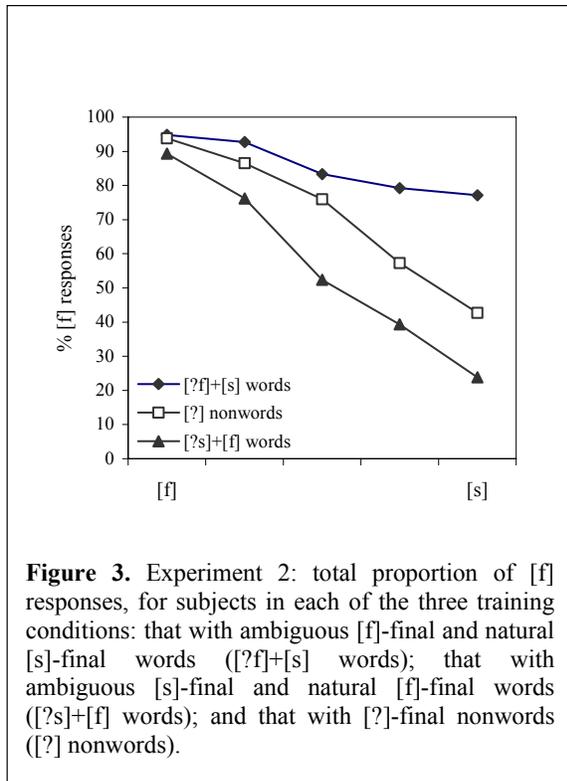


**Figure 2.** Experiment 2: lexical decision performance in the two word training conditions, for the natural and ambiguous versions of the [f]- and [s]-final words. Mean reaction times (RTs, in milliseconds, measured from word offset) are given in the upper panel; mean percentage "no" responses are given in the lower panel.

subjects). All 16 subjects who heard ambiguous [f]-final and natural [s]-final words (e.g., [karaʔ] and *karkas*) passed this criterion. The [ʔ]-final items in the third condition were labelled as nonwords on 97% of trials.

Lexical decision performance in the two word training conditions is summarised in Figure 2. Analyses of Variance (ANOVAs) on the RT data revealed a clear ambiguity effect, in the form of an interaction between training condition and type of final fricative in the original words:  $F_1(1,28) = 484.52, p < 0.001$ ;  $F_2(1,38) = 266.31, p < 0.001$ . Neither of the main effects was significant. The listeners in the [ʔ]-final nonword condition were as slow to decide that the [ʔ]-final items were nonwords (266 ms, on average) as the listeners in the word condition were to decide that the [ʔ]-final items were words (261 ms, on average).

Although listeners were thus slower to make decisions about the [ʔ]-final items, those in the word training conditions nevertheless tended to identify them as tokens of the original words. The overall proportion of "no" responses to [ʔ]-final words was only 10%. The proportion of "no" responses to the original words was 4%. This ambiguity effect was significant:  $F_1(1,28) = 10.96, p < 0.005$ ;  $F_2(1,38) = 6.48, p < 0.05$ . While there was again no main effect of training condition, there was a main effect of type of fricative: there were more "no" responses to the [s]-final words ( $F_1(1,28) = 27.26, p < 0.001$ ;  $F_2(1,38) = 7.40, p < 0.01$ ). This appears to be due to two factors: first, three of the natural [s]-final words had an error rate over 30%; and second, four of



the ambiguous [s]-final words (all of them monosyllabic) were judged to be nonwords more than 30% of the time. In spite of these biases, most subjects accepted most of the [?]-final versions of words as tokens of those words.

### 3.2.2. Categorisation

A strong effect of training condition on categorisation performance was observed. As shown in Figure 3, listeners who had heard ambiguous [f]-final and natural [s]-final words (e.g., [kara?] and *karkas*) were very strongly biased to label all the sounds on the continuum as [f], while those who had heard ambiguous [s]-final and natural [f]-final words (e.g., [karka?] and *karaf*) were the most likely to label the sounds as [s]. The performance of the listeners in the third ([?]-final nonword) group was intermediate. In an ANOVA on the proportion of [f] responses for each fricative there was a significant effect of training condition ( $F(2,43) = 6.56, p < 0.005$ ). There was also an effect of fricative sound ( $F(4,172) = 48.06, p < 0.001$ ) and an interaction of these two factors, reflecting the fact that the training effect varied across the continuum ( $F(8,172) = 4.73, p < 0.001$ ). There were reliably more [f] responses after exposure to ambiguous [f]-final and natural [s]-final words than after exposure to ambiguous [s]-final and natural [f]-final words ( $F(1,28) = 17.40, p < 0.001$ ). Neither word-training condition, however, was significantly different from the nonword-training condition.

It is obvious that there is an [f]-bias in Experiment 2, relative to Experiment 1. Subjects were more likely to label the ambiguous fricatives as [f] after lexical decision than in the original categorisation task. The same bias

was seen in the fact that three subjects had to be excluded from the [karka?] plus *karaf* condition, and in the lexical decision data (where subjects were unwilling to accept [?]-final tokens as instances of monosyllabic words, that is, in the words where the fricative information carried the most proportional weight).

This bias may reflect a shift in cue weighting. Listeners in Experiment 1 heard only one vowel, taken from an [ɛf] context. Since any formant-transition information in this vowel signalling an [f] was therefore not informative in the [f]-[s] decision, listeners may have tended to ignore it. In contrast, the listeners in the lexical decision phase in Experiment 2 heard a relatively large number of different natural utterances with only 10% of the items ending in [?]. This may have increased their reliance on vocalic cues, relative to the Experiment 1 listeners, making the ambiguous sound more [f]-like. This increased reliance on information in the vowel could then have continued in the short categorisation phase in Experiment 2. While this bias is orthogonal to the lexically-mediated training effect, it does suggest, like the lexical effect, that listeners continually adjust how they interpret speech information in the light of the constraints of a given listening situation.

## 4. GENERAL DISCUSSION

Listeners who had heard [?] in [f]-final words during a short lexical decision session, and had tended to identify them as real words, then categorised those sounds as [f] more often than listeners who had heard [?] in [s]-final words in the training session. This suggests that listeners can use their lexical knowledge to adjust the way they evaluate the acoustic-phonetic information in the speech of a particular talker. The training session provided two sources of lexical information that listeners could use to adjust the boundaries of their fricative categories. They heard ambiguous sounds at the ends of words (like [?] in [kara?], consistent with *karaf*); here the lexicon provides support for the interpretation of [?] as an [f]. In addition, they heard natural tokens of words ending with the other fricative endpoint (like *karkas*); here the lexicon provides indirect support that [?] is not an [s].

These results shed new light on the value of feedback in the speech recognition system. As we argued earlier, the direct influence of the lexicon on prelexical representations can do nothing more than cause the lower level to agree with a decision already made at the lexical level, and is therefore of no value to the recognition process. This holds, however, for the recognition of any specific word at any given state of the speech recognition system, assuming all else to remain constant. The present results show how feedback over a longer period of time, during which the system does not remain constant, can indeed offer benefit. If longer-term feedback from the lexicon provides the prelexical level with information on how to interpret particular speech sounds in a given listening situation, this could act to improve word recognition. The learning would act to modulate prelexical representations and hence influence

the processing of other words, improving their recognition. It may thus be the case that lexical influences on prelexical processing will only be observed in situations where long-term adjustments to the idiosyncrasies of a given speech input can be made. This kind of perceptual learning may often take place in real-world situations, as, for example, when we are exposed to the speech of a non-native speaker of our language, or to the speech of someone with a marked dialect.

It will be interesting to establish the parameters under which this kind of feedback may work. How much training exposure is required? How quickly are any adjustments undone when the next new talker is heard? What kind of biasing information is needed to induce noticeable adjustments to phonetic categories (i.e., how consistent does the training have to be)?

In conclusion, we emphasise again that the kind of feedback at issue here differs from that instantiated in TRACE [1], where there is an automatic and mandatory influence of all words on all of their constituent phonemes. Such feedback would not benefit word recognition, and could even hinder it, and we propose that for this reason it does not operate. But the argument that the TRACE kind of feedback does not operate [16], and the demonstration that effects attributed to TRACE-style feedback can be accounted for by learning processes [17], do not imply that there is no feedback anywhere in the system. Instead, we suggest, learning itself may involve a certain type of feedback, a type which only acts when it is beneficial to spoken word recognition. When the lexicon can provide the prelexical level with useful information for further speech processing, it will do so. The speech recognition system is more dynamic than is often assumed. The system adjusts to the constraints of each particular listening situation, and can apparently use feedback from the lexicon as a training signal for these adjustments, exactly in the situation where this leads to more efficient lexical access.

## 5. ACKNOWLEDGEMENTS

We thank Marloes Weijers, Anne Rutgers and Tau van Dijk for their help with this research. Correspondence may be sent to J.M. McQueen, MPI for Psycholinguistics, Postbus 310, NL 6500 AH Nijmegen, or to james.mcqueen@mpi.nl, dennis.norris@mrc-cbu.cam.ac.uk, or anne.cutler@mpi.nl.

## 6. REFERENCES

[1] McClelland, J.L. & Elman, J.L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, **18**, 1–86.  
 [2] McClelland, J.L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology*, **23**, 1–44.  
 [3] Massaro, D.W. (1989). Testing between the TRACE model and the Fuzzy Logical Model of Speech Perception. *Cognitive Psychology*, **21**, 398–421.

[4] Massaro, D.W. & Oden, G.C. (1995). Independence of lexical context and phonological information in speech perception. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1053–1064.  
 [5] Cutler, A. & Norris, D. (1979). Monitoring sentence comprehension. In W.E. Cooper & E.C.T. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 113–134). Hillsdale, NJ: Erlbaum.  
 [6] Cutler, A., Mehler, J., Norris, D., & Segui, J. (1987). Phoneme identification and the lexicon. *Cognitive Psychology*, **19**, 141–177.  
 [7] Ganong, W.F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance*, **6**, 110–125.  
 [8] Connine, C.M. & Clifton, C. (1987). Interactive use of lexical information in speech perception. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 291–299.  
 [9] Elman, J.L. & McClelland, J.L. (1988). Cognitive penetration of the mechanisms of perception: Compensation for coarticulation of lexically restored phonemes. *Journal of Memory and Language*, **27**, 143–165.  
 [10] Samuel, A.G. (1996). Does lexical information influence the perceptual restoration of phonemes? *Journal of Experimental Psychology: General*, **125**, 28–51.  
 [11] Samuel, A.G. (1997). Lexical activation produces potent phonemic percepts. *Cognitive Psychology*, **32**, 97–127.  
 [12] Eimas, P.D., Marcovitz Hornstein, S.B., & Payton, P. (1990). Attention and the role of dual codes in phoneme monitoring. *Journal of Memory and Language*, **29**, 160–180.  
 [13] Frauenfelder, U.H., Segui, J., & Dijkstra, T. (1990). Lexical effects in phonemic processing: Facilitatory or inhibitory? *Journal of Experimental Psychology: Human Perception and Performance*, **16**, 77–91.  
 [14] McQueen, J.M. (1991). The influence of the lexicon on phonetic categorization: Stimulus quality in word-final ambiguity. *Journal of Experimental Psychology: Human Perception and Performance*, **17**, 433–443.  
 [15] Pitt, M.A. & McQueen, J.M. (1998). Is compensation for coarticulation mediated by the lexicon? *Journal of Memory and Language*, **39**, 347–370.  
 [16] Norris, D., McQueen, J.M. & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary. *Behavioral and Brain Sciences*, **23**, 299–325.  
 [17] Norris, D. (1993). Bottom-up connectionist models of 'interaction'. In G. Altmann & R. Shillcock (Eds.), *Cognitive models of speech processing: The second Sperlonga meeting* (pp. 211–234). Hillsdale, NJ: Erlbaum.