Polish Studies in English Language and Literature

Edited by Jacek Fisiak

Vol. 31

PETER LANG
Frankfurt am Main · Berlin · Bern · Bruxelles · New York · Oxford · Wien

Magdalena Wrembel
Małgorzata Kul
Katarzyna Dziubalska-Kołaczyk
(eds.)

Achievements and Perspectives in SLA of Speech: New Sounds 2010

Volume I

PETER LANG
Internationaler Verlag der Wissenschaften

Bibliographic Information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data is available in the internet at http://dnb.d-nb.de.

Typesetting by
Pracownia Wydawnicza IFA UAM

ISSN 1436-7513
ISBN 978-3-631-60722-0

www.peterlang.de

Contents

L1 knowledge and the perception of casual speech processes in L2

Annelie Tuinman[1, 2] *and Anne Cutler*[1, 2, 3]
[1] *Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands*
[2] *Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, The Netherlands*
[3] *MARCS Auditory Laboratories, University of Western Sydney, Australia*

ABSTRACT

Casual speech processes appear in every language. We here examine how listeners deal with casual speech processes in a second language, in the light of the processes in their native language. We compared a match case, where a second-language process (/t/-reduction) also appears in native speech, with a mismatch case, where a second-language process (/r/-insertion) is absent from native speech. In each case, native and non-native listeners judged stimuli in which a given phoneme (in sentence context) varied along a continuum from absent to present. Second-language listeners in general mimicked native performance in the match case, but deviated significantly from native performance in the mismatch case. Together these results make it clear that the mapping from first to second language is as important in the interpretation of casual speech processes as in other dimensions of speech perception. Unfamiliar casual speech processes in a second language can delay correct interpretation of speakers' messages. Casual speech processes that are already familiar from native speech, however, are easy to cope with; indeed, our results even suggest that it is possible for subtle difference in their occurrence patterns across the two languages to be detected, and to be accommodated to in second-language listening.

KEYWORDS: L1, L2, word recognition, casual speech, phonemes, reduction, insertion, Dutch, German, English

1. Introduction

The perception of speech in a second language (L2) cannot be understood except in the light of the perceiver's first-language (L1) knowledge; so much has been known to speech and language researchers for at least the better part of a century (see, e.g., Polivanov, 1931: "The phonological representations of our native language are so tightly coupled to our perception that even when we hear words or sentences from a language with quite different phonology, we tend to analyse these words in terms of the phonemic representations of the native language"). Experimental reports now provide abundant documentation of L1 effects on

speech perception in L2, and models of the L1-to-L2 influence (e.g., Best and Tyler, 2007; Flege, 1995) offer detailed accounts of its genesis.

Like models of speech perception in general, however, models of L2 speech perception are understandably based on a somewhat idealised situation. The mapping of a phoneme or sequence of phonemes to stored representations in the lexicon can be predicted very well by the perceptual models; but the modelled situation will only arise if the input actually presents an acoustic form corresponding to each proposed segment. As listeners and speech researchers know only too well, however, real speech abounds with casual speech processes such as assimilation, reduction, deletion and intrusion, all of which lead to phonetic forms which may deviate drastically from the canonical pronunciation of the words intended by the speaker.

In recent years, psycholinguistics has turned increasingly to investigation of real speech, and how listeners deal with the non-canonical forms it presents. A grossly over-simplified summary of the accrued results to date is that while listeners are extremely good at exploiting the fine phonetic detail of utterances and identifying intended words even when casual speech processes have altered them from their canonical form, the alterations can often (temporarily) mislead listeners, and can often result in word recognition being harder than it would have been for the canonically pronounced versions. The fine differences between intended phonemes and phonemes resulting from a casual speech process have been shown to be exploited by listeners, for example in the case of place of articulation assimilation (e.g., to distinguish the /p/ of English *ripe* in *ripe berries* from the assimilated final phoneme of *right berries*; Gow, 2002), in neutralisation (e.g., to distinguish the final /p/ of Dutch *slip* from the devoiced final sound of *slib*; Warner et al., 2004), and in liaison (e.g., to distinguish the word-initial /p/ in French *trop partisan* from the liaison realisation of a word-final /p/ in *trop artisan*; Spinelli et al., 2003). Listeners are successful at identifying word forms despite assimilation of place (Gaskell and Marslen-Wilson, 1996; Gow, 2001) or of voice (Snoeren et al., 2008) and despite reduction (Ernestus et al., 2002) or other non-canonical realisation (e.g., Alphen and McQueen, 2006; Sumner and Samuel, 2005). Despite all this success at dealing with real-speech forms, however, listeners are also often misled. Thus in a phoneme detection task they respond to phonemes which are not actually in the input at all, because they have been deleted in a casual pronunciation (Kemps et al., 2004), and they respond to phonemes which are accidentally there, such as a medial /p/ in a casual pronunciation of *something* (Warner and Weber, 2001). Their word recognition response times are slowed by many different types of casual-speech forms (Andruski et al., 1994; LoCasto and Connine, 2002; Racine and Grosjean, 2000), and they can be seriously misled, at least temporarily, into assuming that a quite different word is being heard (Brouwer et al., 2008).

All of this research has, of course, been carried out with L1 listeners. But if even these experienced listeners are misled, what is going to happen when L2 listeners hear the same sort of input? Hear it they will, because all languages manifest casual speech processes, and L2 listeners cannot permanently confine themselves to speech situations in which the input is as close to canonical perfection as it is in the classroom or on language tapes. In the current investigation we make a start on charting the perceptual effects of casual speech processes in L2 listening.

Interestingly for our purposes, there are some casual speech processes which are found in many languages, and some which are found in only few. Thus assimilation is widespread across languages, with assimilation of place being extremely common (indeed obligatory in some languages, such as Japanese), and assimilation of voice also quite frequent; but patterns such as the Hungarian assimilation of /lr/ sequences to /rr/ are much rarer (Mitterer et al., 2006). Listeners are better at dealing with the processes that are more common in their L1; e.g., for English speakers, assimilation of place is easier to deal with in word recognition tasks than assimilation of voice, while for French speakers the reverse is the case (Darcy et al., 2007). It is therefore interesting to ask whether this advantage can be transferred to L2 input, i.e., whether listeners will also find it easy to deal with an L2 process if they already have experience with the same process in their L1.

One process that is found in many languages is the reduction, even as far as complete deletion, of /t/ in certain phonetic contexts (Guy, 1980). Empirical evidence shows this process, usually called /t/-reduction, to be a gradient phenomenon, with deletion as one end of the continuum (Mitterer and Ernestus, 2006). In the Germanic languages English, German and Dutch, /t/-reduction produces in each case very similar effects. For instance, /t/ is very likely to be effectively deleted either following /s/ or preceding a bilabial, with the result that most utterances of English *postman*, German *Postbeamter* 'postal worker' or Dutch *postbode* 'postman' are equally unlikely to contain much of a detectable trace of /t/. The perception of /t/ in various degrees of reduction has been intensively studied in Dutch, e.g., by Mitterer and Ernestus (2006), Janse, Nooteboom and Quené (2007) and Mitterer and McQueen (2009); from this body of work we know that Dutch listeners are highly sensitive to the patterning of this phonological process in their L1. The fact that the process is more likely to operate before labial than before other coronal segments, for example, informs their perceptual interpretations, including the choice between word forms; thus listeners make use of the initial phoneme of the following word to decide whether an utterance was intended to be *kas* 'greenhouse' or *kast* 'cupboard' (Janse et al., 2007; Mitterer and McQueen, 2009). If native listeners find it useful to draw on knowledge of the production patterning in perception, the same knowledge would surely be of use to L2 listeners as well. For our /t/-reduction experiments, we

therefore used Dutch speech presented to L1 speakers and to speakers of German with proficient L2 Dutch. The latter's L1 knowledge of the production patterning should be similar to that of the Dutch listeners, given that both languages have /t/-reduction and it patterns similarly in each. We refer to this as a case in which the L1 and L2 casual-speech processes match.

We compare this common process with a far less common process, namely the insertion of /r/ between words beginning and ending with vowels, in British English sequences such as *idea of*. This process is unknown in many other languages, for instance in Dutch (Collins and Mees, 1999). For the /r/-insertion case, we used British English speech presented to L1 listeners and to speakers of Dutch with proficient L2 English, including exposure to the British English standard. Because English has the /r/-insertion process but Dutch does not, we refer to this case as one in which L1 and L2 casual-speech processes mismatch.

The experiments which we carried out were broadly similar, in that in each case we constructed a phonetic continuum of stimuli, in which the phoneme under investigation (/t/, /r/) varied from effectively absent through partially realised to indubitably present. In each case, the listeners' task amounted to judging whether the phoneme was present or absent. In each case we compared the L2 listeners' judgements with the judgements given by L1 listeners presented with the same continuum of stimuli. In the match case, native speakers of German with proficient L2 Dutch and native speakers of Dutch decided in two perception experiments whether or not target words ended in /t/; the target words were verbs in the first experiment and nouns and adjectives in the second. In the mismatch case, native speakers of Dutch with proficient L2 English and native speakers of British English decided whether target words were tokens of *ice* or of *rice*.

2. Perception of reduced /t/

2.1. Method

Twenty-one native speakers of Dutch and 16 native speakers of German participated in the first experiment, and 16 native speakers of Dutch and 16 native speakers of German in the second experiment. All were undergraduate students at the Radboud University Nijmegen. The German participants therefore necessarily had a high level of proficiency in Dutch as L2.

Five realizations of /t/, from full production to complete deletion, were presented in two acoustic contexts, after /n/ (where /t/-reduction is unlikely) and after /s/ (where /t/-reduction occurs frequently). The selection of the /s/ and /n/ context and the different forms of reduction were based on patterns observed in a corpus study; the range of /t/ values in the materials had also previously been used in a perceptual study (in both cases, Mitterer and Ernestus, 2006). In each

sentence, listeners judged whether the target word ended in /t/ or not. To ensure optimal control over the acoustic form of the materials, all stimuli were synthesised using a Klatt (1980) formant synthesiser, using formant values and bandwidths determined from segment-by-segment measurements of recordings of the stimulus set by a male native speaker of Dutch.

We conducted two experiments, both using the same range of /t/ values and the same manipulation of preceding phonetic context. In the first experiment (Syntax Experiment), the target words were verbs (e.g., *ren* 'run', *kus* 'kiss'). This made it possible to use grammar, namely a preceding subject pronoun *ik* 'I' versus *zij* 'she', to influence whether or not the ending should properly be /t/; the Dutch present-tense third-person singular inflection is /t/, while the first-person inflection is null, so that the grammatical forms are *ren* and *kus* after *ik*, but *rent* and *kust* after *zij*. In the second experiment (Lexicality Experiment), the target words were nouns and adjectives. In this case, lexical information exercised a similar effect; interpreting /t/ made the target item a correct Dutch word or not, in that *charmant* 'charming' and *kanon* 'gun' are real words, but *charman* and *kanont* are not. We can compare the relative sensitivity of L1 and L2 listeners to the acoustic evidence for /t/, as well as their sensitivity to the preceding phoneme, and to the two types of higher-level information that require versus rule out a /t/.

2.2. Results

The results for the Syntax Experiment with verbs and the Lexicality Experiment with nouns and adjectives were analyzed separately, in each case with a linear mixed-effects model (Baayen, 2008). In linear mixed-effects modelling, a linear regression analysis partials out the effects of each random factor involved in the design, and, starting with a full model, from which insignificant interactions are pruned in stepwise fashion, computes the significance of the effects of each fixed factor, separately and in combination. We chose an alpha level of .05, and every effect referred to below as significant is significant to at least this level.

For both experiments, Participant served as the sole random factor, and fixed factors were Native Language, Coda Signal (full production to complete deletion), Preceding Phonetic Context (/n/, /s/) and the higher-level factor (/t/ predicted or not) relevant for the experiment, i.e., Grammar in the Syntax Experiment, and Lexicality in the Lexicality Experiment. For the Native Language variable, the German speaker group was mapped on the intercept (note that the choice of which group is mapped on the intercept, i.e. serves as base for comparison with the other group, is arbitrary and has no effect on the statistical outcome). Figure 1 shows the mean /t/-response percentages (i.e., how often listeners reported the presence of a /t/).

Figure 1. Syntax and Lexicality Experiments: Percentages of /t/-responses by Dutch (above) and German listeners (below).

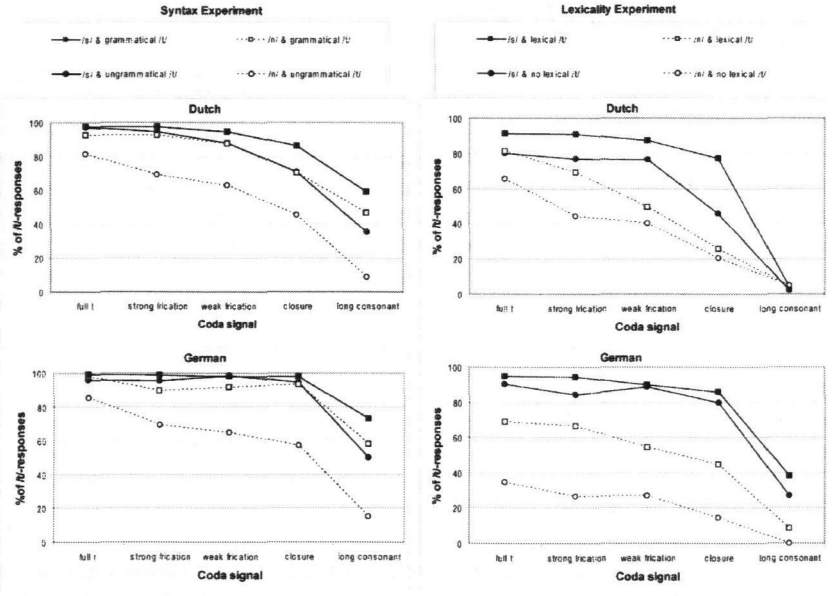

The linear mixed-effects analysis for the Syntax Experiment revealed an overall significant tendency for more /t/-responses (i.e., reports that a /t/ was present) by German (L2) than by Dutch (L1) participants. This tendency, however, was modulated by a number of interactions with the other factors, in particular the combination of Native Language and Grammar. To understand these interactions, we examined the effects of Native Language, Preceding Context and Grammar at each of the five levels of Coda Signal separately. These analyses revealed that for the full /t/ and strong frication coda signal (the two leftmost steps in the figure) there were no overall differences between L1 and L2 listeners, but for the full /t/ Germans had a larger effect of Grammar and a smaller effect of Preceding Context than L1 listeners. For the weak frication coda signal (midpoint in the figure), the Germans gave more /t/-responses than the Dutch and the effect was enlarged where Grammar predicted a /t/. For the closure coda signal (fourth step in the figure), an overall effect of Native Language was affected by Grammar and Preceding Context: Germans gave more /t/-responses in all cases except when Preceding Context was /n/ and Grammar predicted no /t/. Finally, for the long consonant coda signal (rightmost step in the figure)—where there is no hint of /t/, but rather just a lengthened /n/ or /s/—Germans overall gave more /t/-responses than the Dutch.

The analysis of the Lexicality Experiment showed main effects of both Preceding Context and Lexicality; more /t/-responses were given after /s/ than after /n/, and more /t/-responses were given if an existing word resulted. Overall, the effect of Preceding Context was larger for the L2 listeners. However, the reported effects were again modulated by interactions. Therefore, we once again examined the effects of Native Language, Preceding Context and Lexicality on all five levels of Coda Signal separately. The strong Coda Signals (two leftmost steps in the figure) showed a consistent pattern, with main effects for Preceding Context and Lexicality and a significant interaction between Native Language and Preceding Context—L2 listeners had a larger Preceding Context effect than the L1 listeners did. For the weak Coda Signals the effect of Lexicality was larger for L2 listeners, but depended on the Preceding Context, as the figure shows.

In summary, the German L2 listeners' responses are, overall, sensitive to the same factors as the Dutch L1 listeners' responses; on each side of Figure 1 a very similar pattern is visible in the top and bottom graphs, for L1 and L2 listeners respectively. However, in a few cases the L2 responses include significantly more reports of /t/ than the L1 responses. A comparison with production data from the two languages involved suggests an explanation for this pattern. In German, reduction of morphologically functional /t/ (such as a verb inflection) rarely occurs, and is thus presumably avoided by speakers. In Dutch, in contrast, reduction is equally likely for morphological and for non-morphological /t/. Separate production experiments with non-Dutch-speaking Germans and Dutch native speakers confirmed this pattern (Tuinman, 2006). In the light of this comparison, the tendency of the L2 group to produce more /t/ responses than the L1 group, especially when Grammar predicted a /t/, suggests that the German listeners were sensitive to this difference between their L1 and their L2, and tried, perhaps a little too hard, to compensate for it.

3. Perception of intrusive /r/

3.1. Method

Eighteen native speakers of Dutch and 18 native speakers of British English took part in the experiment. The British English participants were students at the University of Sussex. The Dutch participants were students at the Radboud University Nijmegen, and had a high level of proficiency in English as a second language. On average, they had received seven years of English instruction in primary and secondary education. Note that the target variety for Dutch high-school English-language teaching is British English. Also, although other varieties of English are also available via the Dutch media (where movies and television programs are sub-titled, never dubbed), British English television channels (BBC1

and BBC2 at least) form part of the media package available throughout the Netherlands and can thus be received in every household, and long-wave radio broadcasts from the UK can also be received in the Netherlands.

We presented listeners with gradient realizations of /r/ in four sentence contexts varying in contextual and orthographic bias for the presence of /r/. On each trial, listeners judged for the critical word *(r)ice* whether they had heard *ice* or *rice*. There were two pairs of sentences, differing only in whether the word *more* was included. One pair *(the social worker saw (more) ice was given to the poor)* was intended to have a semantic bias towards *rice*; the other pair *(the little girl saw (more) ice was given to her brother)* was assumed to be less biased. Orthographic bias was established by the words *saw* and *more* preceding the target word *(r)ice*. As the phrase *more ice* includes /r/ in the spelling, the perception of /r/ in the speech signal can be attributed to *more*, while in the case of *saw ice*, perception of an /r/ sound cannot be mapped to the spelling. An orthographic bias should therefore manifest itself in terms of more reports of *rice* after *saw* than after *more*.

Again we controlled the form of the stimuli with speech synthesis. For British English, diphone synthesis (giving a more natural-sounding result than formant synthesis) is available, so we used MBROLA (Dutoit et al., 1996), but again modelled the values we chose on the productions of the four sentences by a native British English speaker (female, from London). Prior measures of the same speaker's natural productions had established that she reliably produces /r/ in contexts such as *saw a*, and that such intrusive /r/ productions are significantly shorter (mean 69 ms) than her onset /r/s (mean 89 ms; Tuinman et al., 2007). The /r/ tokens in our stimuli varied from very short (25 ms) to very long (121 ms).

3.2. Results

Figure 2 shows the percentage of *rice* judgments as a function of the three independent variables, separately for the L2 and L1 listeners. The principal result is clear at a glance: in contrast to the case in Experiment 1, here the response patterns for the two listener groups differ. L1 listeners based their responses mainly on durational variation, while L2 listeners barely used durational information, but produced a clear separation between responses in sentences with *saw* (squares) versus *more* (circles).

Figure 2. Percentages of „rice" identifications, for British English listeners (top) and for Dutch listeners (bottom).

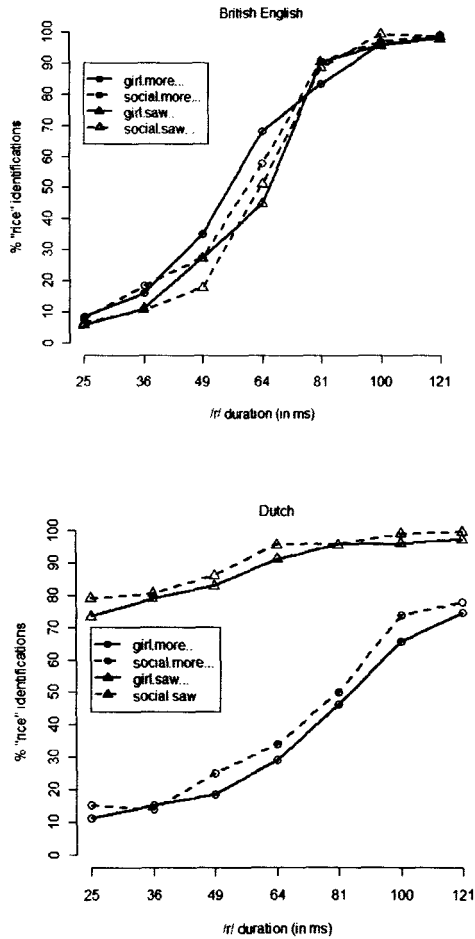

The results were once again analyzed with a linear mixed-effects model, with Participant as random factor. The fixed factors in this case were Native Language, Orthography, Context, and Duration. Duration was entered as a numerical variable, centered on zero, and Orthography and Context, being binary contrasts, were coded as -0.5 for -/r/ bias, +0.5 for +/r/ bias. For Native Language, the L1 group was mapped on the intercept. Analysis again started with a full model, with

insignificant interactions pruned in stepwise fashion. Once again, any main effect or interaction referred to as significant remained, after this pruning process, significant at the .05 level at least.

The L1 listener group showed a significant effect of the duration of /ɪ/, but no effect of the semantic context at all. There was a small effect of orthography in their responses, but it was negative (i.e., the L1 listeners in fact gave somewhat more *rice* responses when the preceding word was *more* than when the preceding word was *saw*). The L2 listeners, in contrast, showed a significant effect of orthography in the predicted direction, i.e., they gave more *rice* responses when the preceding word was *saw* rather than *more*. The Dutch L2 listeners were also significantly less categorical in their responses than the British L1 listeners (see Figure 2: the upper graph shows a typically categorical function, but the functions in the lower graph are almost flat in comparison). In other words, the Dutch listeners were less reliant than the L1 listeners on duration as a signal of whether the /r/ was an onset or not. Additionally, the Dutch group was influenced by semantic context, and made less use of /r/ duration when there was an orthographic bias.

4. Conclusion

The combined results of our two sets of studies motivate the inescapable conclusion that casual speech processes are subject to the same tight coupling with L1 listening experience as every other aspect of L2 speech perception. The L2 listeners' performance in the match case (/t/-reduction experiment) related quite differently to the native performance baseline than the performance of the L2 listeners in the mismatch case (/r/-intrusion experiment). The listeners' experience with the casual speech processes of their L1 crucially determined the way they dealt with the effects of casual speech processes in their L2.

In the match case, the German L2 listeners were broadly sensitive to the same range of factors in the Dutch input as the Dutch L1 listeners: they were sensitive to the acoustic form of the stimuli, and they were in general more likely to report that a /t/ was present (a) when the phoneme preceding it was /s/ rather than /n/, (b) when the /t/ made a pronoun+verb into a grammatical string, and (c) when the presence of a /t/ at the end of a string made a real word. Their similar sensitivity to the acoustic realisation of /t/, and to the probabilities as a function of preceding phonetic context, reflect the parallelism between Dutch and German in the patterning of /t/ deletion.

In the mismatch case, however, the Dutch listeners showed quite a different pattern of sensitivity than the English L1 listeners: while the L1 listeners based their responses overall on the acoustic (i.e., durational) characteristics of the stimuli, were quite insensitive to the sentence meaning and certainly were not

inclined to report /r/ if there was none in the orthographic representation, the L2 listeners mimicked none of these patterns. They made relatively little use of the acoustic information, they were significantly influenced by sentence meaning, and they were far more likely to report that an /r/ was present in word-onset position when the orthography represented none in the preceding word-coda position. Although Dutch English-learners have British English as target variety (Gussenhoven and Broeders, 1997), and this form of English is readily available in their daily listening experience, they cannot process this particular British casual speech process, not found in their L1, in a native-like manner.

We conclude, therefore, that the casual speech patterns of the L1 influence interpretation of casual speech processes in an L2. A process familiar from the L1 can be processed easily; an unfamiliar L2 process with no corresponding L1 experience is very hard to adapt to, in the sense that it is hard to learn to process it in the way that L1 listeners do. This finding has significant implications for the understanding of speech, since, as outlined above, L1 listeners are very good at adjusting spoken-word recognition to take account of the effects of casual speech processes. We predict that German listeners to Dutch should be able to judge whether they are hearing Dutch *kas* or *kast* with similar efficiency to native Dutch listeners (cf. Janse et al., 2007; Mitterer and McQueen, 2009). The same should in fact be true for any L2 listener confronted with a casual speech process that is already familiar from the L1. Dutch listeners to English, however, will be unable to mimic the efficiency with which British English listeners distinguish an intrusive /r/ from an intended /r/. This could lead them into misrecognition of words; indeed, our further investigations have shown that this happens (Tuinman et al., 2007). Again, we predict that this pattern will hold in similar fashion for other cases of L1-L2 mismatch in a casual speech process.

However, recall that the L2 listeners also deviated from L1 performance in the match case in one respect: they were more likely than the L1 listeners to report that a /t/ was present when the /t/ would constitute a morphological inflection. As production evidence indicated, morphologically conditioned /t/ is precisely the kind of /t/ that is less likely to undergo reduction in their L1. This pattern therefore raises the intriguing possibility that listeners not only can deal easily with a casual speech process of the L2 that matches an L1 process, but they are even sensitive to any subtle differences in its distribution of occurrence across the languages, and can accommodate their perceptual responses accordingly. This fine-grained sensitivity to the distributional patterning of casual speech processes seems to us certainly worthy of future investigation, both in L1 and in L2 listening. We hope, too, that many further studies of the role of casual speech processes in L2 listening will be undertaken, in particular with respect to the implications of the L1-L2 mapping for successful communication.

5. Acknowledgements

This research was supported by a doctoral grant from the Max Planck Society and by a Fulbright travelling fellowship, both to the first author. We are grateful to Holger Mitterer for advice on all aspects of the project, and to the audience at New Sounds 2010 for helpful input.

REFERENCES

Alphen, P. M. van, McQueen, J. M. 2006. The effect of voice onset time differences on lexical access. *Journal of Experimental Psychology: Human Perception and Performance* 32. 187–196.

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

Baayen, R. H. 2008. *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.

Best, C. T., Tyler, M. D. 2007. Nonnative and second-language speech perception: Commonalities and complementarities. In: Munro, M. J., Bohn, O.-S. (eds.). *Second language speech learning* Amsterdam: John Benjamins.

Brouwer, S., Mitterer, H., Ernestus, M. 2008. How acoustically reduced forms activate the lexicon: Evidence from eye-tracking *Laboratory Phonology 11*. Wellington, New Zealand.

Collins, B., Mees, I. M. 1999. *The phonetics of English and Dutch* (4th ed.). Leiden: Brill.

Darcy, I., Peperkamp, S., Dupoux, E. 2007. Bilinguals play by the rules: Perceptual compensation for assimilation in late L2-learners. In: Cole, J., Hualde, J. (eds.). *Papers in Laboratory Phonology 9*. Berlin: Mouton de Gruyter. 411–442.

Dutoit, T., Pagel, V., Pierret, N., Bataille, F., Vreken, O. van der. 1996. *The MBROLA Project: Towards a Set of High-Quality Speech Synthesizers Free of Use for Non-Commercial Purposes*. Paper presented at the 4th International Conference on Spoken Language Processing (ICSLP 1996), Philadelphia.

Ernestus, M., Baayen, H., Schreuder, R. 2002. The recognition of reduced word forms. *Brain and Language* 81. 162–173.

Flege, J. E. 1995. Second language speech learning: Theory, findings, and problems. In: Strange, W. (ed.). *Speech perception and linguistic experience: Issues in cross-language research*. Baltimore: York Press. 233–277.

Gaskell, M. G., Marslen-Wilson, W. D. 1996. Phonological variation and inference in lexical access. *Journal of Experimental Psychology: Human Perception and Performance* 22. 144–158.

Gow, D. W. 2001. Assimilation and anticipation in continuous spoken word recognition. *Journal of Memory and Language* 45. 133–159.

Gow, D. W. 2002. Does English coronal place assimilation create lexical ambiguity? *Journal of Experimental Psychology: Human Perception and Performance* 28. 163–179.

Gussenhoven, C., Broeders, A. 1997. *English pronunciation for student teachers.* Groningen: Wolters-Noordhoff.

Guy, G. R. 1980. Variation in the group and the individual. In: Labov, W. (ed.). *Locating language in time and space.* New York: Academic Press. 1–36.

Janse, E., Nooteboom, S. G., Quené, H. 2007. Coping with gradient forms of /t/-deletion and lexical ambiguity in spoken word recognition. *Language and Cognitive Processes* 22. 161–200.

Kemps, R. J., Ernestus, M., Schreuder, R., Baayen, R. H. 2004. Processing reduced word forms: The suffix restoration effect. *Brain and Language* 90. 117–127.

Klatt, D. H. 1980. Software for a cascade/parallel formant synthesizer. *Journal of the Acoustical Society of America* 67. 971–995.

LoCasto, P., Connine, C. M. 2002. Rule-governed missing information in spoken word recognition: Schwa vowel deletion. *Perception and Psychophysics* 64. 208–219.

Mitterer, H., Csépe, V., Blomert, L. 2006. The role of perceptual integration in the perception of assimilation word forms. *Quarterly Journal of Experimental Psychology* 59. 1395–1424.

Mitterer, H., Ernestus, M. 2006. Listeners recover /t/s that speakers reduce: Evidence from /t/-lenition in Dutch. *Journal of Phonetics* 34. 73–103.

Mitterer, H., McQueen, J. M. 2009. Processing reduced word-forms in speech perception using probabilistic knowledge about speech production. *Journal of Experimental Psychology: Human Perception and Performance* 35. 244–263.

Polivanov, E. 1931. La perception des sons d'une langue étrangère. *Travaux du Cercle Linguistique de Prague* 4. 79–96.

Racine, I., Grosjean, F. 2000. Influence de l'effacement du schwa sur la reconnaissance des mots en parole continue. *L'Année Psychologique* 100. 393–417.

Snoeren, N., Segui, J., Hallé, P. 2008. Perceptual processing of partially and fully assimilated words in French. *Journal of Experimental Psychology: Human Perception and Performance* 34. 193–204.

Spinelli, E., McQueen, J. M., Cutler, A. 2003. Processing resyllabified words in French. *Journal of Memory and Language* 48. 233–254.

Sumner, M., Samuel, A. G. 2005. Perception and representation of regular variation: The case of final /t/. *Journal of Memory and Language* 52. 322–338.

Tuinman, A. 2006. Overcompensation for /t/-reduction in Dutch by German-Dutch bilinguals, *Laboratory Phonology 10.* Paris, France.

Tuinman, A., Mitterer, H., Cutler, A. 2007. Speakers differentiate English intrusive and onset /r/, but L2 listeners do not. In: Trouvain, J., Barry, W. J. (eds.). *Proceedings of the 16th International Congress of Phonetic Sciences.* Dudweiler, Germany: Pirrot. 1905–1908.

Warner, N., Jongman, A., Sereno, J., Kemps, R. J. 2004. Incomplete neutralization and other sub-phonemic durational differences in production and perception of Dutch. *Journal of Phonetics* 32. 251–276.

Warner, N., Weber, A. 2001. Perception of epenthetic stops. *Journal of Phonetics* 29. 53–87.