The impact of strong assimilation on the perception of connected speech

M. Gareth Gaskell and Natalie D. Snoeren
Department of Psychology, University of York, York, England

Models of compensation for phonological variation in spoken word recognition differ in their ability to accommodate complete assimilatory alternations (such as *run* assimilating fully to *rum* in the context of *a quick run picks you up*). Two experiments addressed whether such complete changes can be observed in casual speech, and if so, whether they trigger perceptual compensation. Experiment 1 used recordings of naïve speakers and found that the presence of following context supporting place assimilation led to an increase in miscommunication rate when listeners were asked to identify the potentially assimilated words. This result was also obtained when trained phoneticians gave their considered judgments of a subset of the stimuli. Experiment 2 examined the extent to which words articulated correctly by naïve speakers (e.g., *rum*) would be perceived as assimilated, and found that compensation for assimilation in these stimuli depended on the type of following phonemic context and the semantic fit with the preceding sentence. These results suggest that place assimilation does involve complete alternations and that the perceptual system can compensate for them in certain circumstances.

Key words: Assimilation, speech perception, spoken word recognition, lexical access, phonological variation

Thanks are due to Rachel Cole-Fletcher, Xiaojia Feng, Liam Maxwell and MaryAnn Noonan for assistance in running the experiments, and to Susie Flint, who carried out pilot work as a BSc project. Philip Quinlan provided guidance on statistical issues, and members of the Psycholinguistic Research group at York provided many thoughtful comments on the research. Special thanks are due to the 8 phoneticians who took the time to take part in the posttest of Experiment 1, and to Paul Foulkes and Richard Ogden, who provided advice on the posttest. This research was supported by a grant from the Medical Research Council and an Economic and Social Research Council fellowship awarded to MGG, and by funding from the Fondation Fyssen awarded to NS.

Speech is the primary means of human communication, and the ability to recognize speech is one area in which the human perceptual system can outperform automatic systems by a considerable degree (Lippmann, 1997; Moore, 2007). In terms of spoken word recognition, a characteristic of the human system that underlies this high-quality performance is an exquisite sensitivity to fine-grained details of speech at segmental and subsegmental levels (Andruski, Blumstein & Burton, 1996; Davis, Marslen-Wilson & Gaskell, 2002). For example, differences of ten or twenty milliseconds in the durations of the /s/ segments in "once paid" vs. "one spade" are picked up by the perceptual system in order to activate the appropriate lexical representations (Shatzman & McQueen, 2006). This ability to make use of subtle cues or cue combinations is imperative given the task at hand, but also comes with a potential cost. Speech is inherently variable, and many relevant fine distinctions can be obscured by the kinds of assimilation and other phenomena that commonly take place during the production of connected speech. For example, assimilation of place of articulation can mean that a word like *sit* can be produced more like *sip* in the context of "sit politely", or like *sick* in the context of "sit carefully". Given that such changes in isolation can have drastic effects on the activation of the corresponding lexical entries (Marslen-Wilson, Moss & van Halen, 1996; McQueen & Viebahn, 2007), researchers need to explain how the perceptual system maintains sensitivity to the subtleties of speech without variability on the production side having catastrophic effects on communication.

In recent years this crucial question has become a focus of research, with studies of vowel deletion (LoCasto & Connine, 2002), unreleased and glottalized stop consonants (Deelman & Connine, 2001; Sumner & Samuel, 2005), and speech rate effects (Janse, 2004), as well as a number of studies of assimilated speech (e.g., Gaskell & Marslen-Wilson, 1996; Gow, 2001; Snoeren, Seguí & Hallé, 2008). These studies have highlighted the distinction between changes in form related to normal connected speech processes in production, and arbitrary deviations or mispronunciations: phonologically regular variations are perceptually acceptable, whereas comparable irregular variants tend to be treated as nonwords, preserving the overall sensitivity to form deviation in the perceptual system. For example, Sumner & Samuel (2005) examined the influence of variants of /t/-final words (e.g., *flute*) on access to their meanings using semantic priming. Glottalized and coarticulated forms of these words facilitated recognition of a semantically related word (e.g., *music*) as well as the citation form, but unnatural variants involving a feature change (e.g., "fluce") showed no priming effect.

In the debate over how this distinction between natural and abnormal changes is realized in the perceptual system, there has been a particular focus on assimilation of place of articulation, particularly in US and British English. Place assimilation often affects word-final alveolar consonants such as /d/, /t/ and /n/, which in connected speech can take on the place of articulation of the following labial or velar consonant, as in the examples of "sit politely" and "sit carefully" above. This description of assimilation in terms of discrete consonants, alongside the fact that place assimilation is asymmetric (coronal consonants take on labial and velar place but not vice versa), has meant that it has typically been described as a phonological phenomenon. For example, Lahiri and colleagues (Lahiri & Marslen-Wilson, 1991; Lahiri & Reetz, 2002) explain the asymmetry of place assimilation in terms of a default status of coronals (Paradis & Prunet, 1991), which according to underspecification theory means that lexical representations of coronal consonants such as alveolars do not contain specific [CORONAL] features. This leaves them more susceptible to gaining the place of the following noncoronal labial or velar consonant in the appropriate context. Abstract phonological views of assimilation can be contrasted with phonetic or articulatory descriptions in terms of overlap of gestures (e.g., Browman & Goldstein, 1992). These alternatives are supported by the finding that many cases of place assimilation leave some residual alveolar gesture alongside the velar or bilabial closure (Nolan, 1992). Acoustically, place assimilations leave patterns of formant change that are sometimes characteristic of alveolars, sometimes characteristic of velars or labials, and sometimes intermediate (Gow, 2002). All these forms of data suggest that assimilation is not a discrete shift from one consonant to another as a phonological view might predict, but is instead a more graded phenomenon. This gradedness might result in phonetic evidence of varying strength for two different consonants (as in place assimilation), or in a continuum running between two different consonants (as in voice assimilation).

Current models of assimilation perception

Several competing theories have addressed the psycholinguistic impact of assimilation in perception. Lahiri and colleagues have argued that underspecification of lexical form can explain assimilation perception as well as production (Lahiri & Marslen-Wilson, 1991; Lahiri & Reetz, 2002). With respect to place assimilation, the asymmetry of representation provided by underspecification is crucial. Coronal segments such as /n/ are lexically underspecified in such a model, meaning that the spoken form “leam” would be acceptable as a token of *lean* because the lexical representation of the /n/ in lean is unspecified for place of articulation, and so does not mismatch the incoming [m]. Conversely, “trin” is unacceptable as a token of *trim*, because the lexical representation of *trim* specifies labial place, which mismatches the incoming speech. This allows underspecification to accommodate asymmetries in the perceptual response to deviant phonemes in mismatch negativity and priming experiments (Eulitz & Lahiri, 2004; Wheeldon & Waksler, 2004). Because underspecification tolerates potentially assimilatory changes regardless of the following context it also predicts that “leam bacon” and “leam gammon” should be equivalent in terms of their activation of the lexical representation of *lean*. This prediction fits with cross-modal priming data from Wheeldon and Waksler (2004), but contradicts a larger number of studies (e.g., Coenen, Zwitserlood & Bölte, 2001; Gaskell & Marslen-Wilson, 1996, 1998, 2001; Gow, 2001, 2002, 2003; Mitterer & Blomert, 2003). These latter studies have demonstrated that contextually “viable” assimilations, in which the place of articulation of the following context matches the place of the assimilated segment (e.g., “leam bacon”) are more acceptable to the perceptual system than unviable ones (e.g., “leam gammon”), in which the same place change could not have occurred through assimilation.

An alternative account was developed in response to these contextual viability effects, highlighting an active process of compensation, rather than accommodation through representation. Gaskell, Hare & Marslen-Wilson (1995) developed a connectionist model that learned to compensate for assimilation through experience of the conditions in which assimilation occurred. The key aspect of this approach is the correlation of an assimilated consonant with a place-matching following context. When these circumstances are observed, the system learns to compensate for assimilation in the surface forms of speech, leading to a “cleaned up” underlying representation. Further empirical studies (Gaskell & Marslen-Wilson, 1996, 1998,

2001) demonstrated the importance of higher-level knowledge (e.g., lexical and sentential context) in the learned compensation process.

A third type of model has stressed the importance of phonetic detail in the accommodation of regular variation. As mentioned above, place assimilation, like many other assimilation types (e.g., Snoeren, Hallé & Seguí, 2006), is a graded phenomenon (Nolan, 1992). In many cases, an assimilated segment will have features in common with both the intended segment and the assimilation context. Thus the [m] in "leam bacon" may have features associated with both alveolar (relating to the final /n/) and labial (relating to the following /b/) place. Gow (2002, 2003) argued that compensation for place assimilation involves a process of feature parsing, in which the two sets of place features in an assimilated segment must be associated with the appropriate underlying segments. The alveolar features of an assimilated /n/ will be associated with the /n/ itself, whereas the bilabial features will be used as evidence for an upcoming bilabial consonant. In the absence of a following bilabial consonant (as in "leam gammon"), the unattributed bilabial features are associated with the nasal consonant itself, allowing a feature parsing account to accommodate viability effects on the perception of assimilated consonants (e.g., Gaskell & Marslen-Wilson, 1998; Gow, 2003). This view of assimilation compensation is attractive in that it makes use of the same feature extraction process that is assumed to operate throughout the perception of connected speech in order to deal with coarticulated segments. Compensation for assimilation, by this view, is simply another case of compensation for coarticulation, in which context-dependent changes in segments do not neutralize phonemic distinctions (Mann & Repp, 1981; Repp, 1983). Mitterer and Blomert (2003) proposed a second low-level account of assimilation compensation, in which compensation occurs at a level that is both perceptual and subsymbolic.

Differential predictions of the current models

The models described above should not be thought of as mutually exclusive (Gaskell, 2003; Snoeren et al., 2008), but nonetheless several testable predictions of the various accounts of compensation for assimilation can be derived. First, the feature parsing account (Gow, 2003) can only operate effectively given partial assimilation. This is because it requires information about the intended place of articulation to be present once the assimilated place has been associated with the following consonant. Consider a case where assimilation is complete, such that the /n/ in *lean bacon* loses all coronal features to become like a typical /m/. In these circumstances perceptually associating the bilabial features solely with the following consonant is implausible, because no place of articulation would remain for the assimilated segment (and indeed all adjacent segments identical in place of articulation would have to be treated the same way, leaving the /n/ in *line drawing* perceptually "placeless"). Therefore, the feature parsing model is weakened if some subset of assimilations are complete or near-complete, and listeners compensate for assimilation in these circumstances.

A related prediction of the feature parsing model is that no compensation for assimilation should occur in cases where the listener encounters two consonants in connected speech that happen to have the same noncoronal place of articulation (e.g., /p#b/ in "ripe berries"). The reasoning is the same: the /p/ in ripe will contain no trace of an alveolar consonant, and so cannot be treated as alveolar. On the other hand, the learned compensation approach, because it accommodates complete assimilatory change, would tend to compensate in this case, depending on the biases available from lexical and sentential context, as well as phonetic factors such as the presence of a release burst in the /p/ and the timing of the following consonant.

Evidence relating to the strength of place assimilation

With reference to the first prediction of the feature parsing model, Nolan (1992) showed that there was a continuum of place assimilation for voiced consonant, with a significant proportion of assimilations leaving no residual alveolar contacts. These strongly assimilated sequences were difficult to discriminate from matched velar or bilabial distractors (e.g., *road collapsed* vs. *rogue collapsed*); although performance was above chance when a two-alternative forced choice (2AFC) comparison was used.

Nolan's study made use of a single phonetician to produce the assimilations, and so it is difficult to draw strong conclusions about the extent of natural variation from this study. Ellis and Hardcastle (2002) used ten naïve English speakers in a combined electropalatography and electromagnetic articulography study of nasal place assimilation. The speakers varied widely in assimilation style, with two producing no observable assimilation, four showing apparently complete assimilation in all cases, and four demonstrating a more variable pattern of responses (either a gradient of assimilation or alternating between no assimilation and full assimilation).

Although these production data suggest that many cases of assimilation are not partial, Gow (2002, 2003) presented perception data suggesting that assimilation does not produce lexical ambiguity. Gow (2002) used recordings of a naïve US English speaker, who read sequences designed to encourage place assimilation. The potentially assimilated words were members of minimal pairs (e.g., *right/ripe*) such that a strong assimilation might lead to lexical ambiguity. Pretesting of the recordings with the speech following a potentially assimilated segment spliced out showed that many of the stimuli were rated as being more like noncoronal-ending words (e.g., *ripe*) than the corresponding coronal-final item (e.g., *right*). Nonetheless, even when these relatively strong assimilations were used as primes in a cross-modal form priming study, there was no significant facilitation of noncoronal-final visual targets, despite the robust priming of the coronal-final items. In other words, even stimuli that when spliced out of context sounded like they ended in noncoronal segments were treated by the perceptual system as if they ended in coronal segments. A very different pattern of priming emerged when the speaker intended to produce the noncoronal-final words (e.g., *ripe*). Here, despite a following context suitable for assimilation (e.g., *ripe berries*) there was no evidence of compensation for assimilation, with only the noncoronal-final targets showing priming.

These data suggest that assimilated speech does not cause ambiguity, in that an assimilated token of *right* accesses only the meaning of *right*, whereas a token or *ripe* placed in an environment suitable for compensation for assimilation access only the meaning of *ripe*. Furthermore, the fact that no compensation for assimilation occurred with intended noncoronal segments appears to confirm the second prediction described above. However, there are a couple of potential weaknesses in this argument. First, only one speaker was used in Gow's (2002) experiment. Ellis and Hardcastle (2002) demonstrated a substantial degree of interspeaker variability in terms of assimilation style. It is possible that priming would be very different in the case where a speaker's assimilation style was "all-or-nothing" or dominated by extreme assimilations. A second point is that the naïve speakers in the above studies were aware that they were being recorded for the purposes of an experiment. There is a strong possibility that speakers, even when asked to speak casually, speak in a clearer way when asked to record sentences for an experiment. A pilot study in our lab (Flint, 2003) confirmed this intuition, showing that when speakers recorded potential assimilations in circumstances similar to those of Gow (2002) their speech was much less ambiguous than if speakers recorded the same sequences as "responses" in a dummy memory task.[1] Therefore, it may be premature to draw strong conclusions from data based on recordings of "casual" speech when a speaker is aware of the nature of the recording.

The present study

Our paper addressed two empirical questions derived from the differential predictions presented earlier:

1) Does speech contain place assimilations that are extreme enough to be treated as normal velar or labial segments?

2) Do normal velar or labial segments trigger compensation for assimilation when articulated in appropriate contexts?

A positive answer to either of these questions would be evidence against a feature parsing account of the perception of assimilated speech and evidence for a learned compensation account. We used a methodology similar to Flint (2003), in which participants unwittingly generated potential assimilations as responses in a memory experiment. These recordings were then used as stimuli to examine the perception of assimilations in different circumstances. In Experiment 1, we addressed the first of our research questions, relating to the incidence of strong, ambiguity-creating assimilation in normal speech. All usable recordings were presented to participants in a speeded-choice task, in which listeners had to decide which of two alternatives (e.g., *ripe/right*) was presented in the preceding speech. A subset of these stimuli was also categorized by a group of trained phoneticians in order to derive expert, considered judgments as to the intended place of articulation of the potentially assimilated consonant. Experiment 2 addressed the second of our objectives, looking at whether intended noncoronal consonants can trigger compensation for assimilation. A second sample of the recordings was presented in a speeded-choice task, with this time only intended noncoronal segments used. The influence of the place of articulation of the following context was assessed by presenting these stimuli with both their original following context and cross-spliced onto alternative contexts.

General Method and Materials (Dummy Experiment)

The main experiments used materials collected in a dummy test in which participants produced sentences as responses to a simple memory task. The intention was to encourage participants to produce given sentences but in a reasonably natural and casual fashion (cf. Flint, 2003). To achieve this, participants took part in a series of trials, in which three numbered words were presented (e.g., "(1) in (2) out (3) up"). After a short delay a sentence frame was presented in which a word was replaced by a number (e.g., "I think a quick run picks you (3) "). The number referred to one of the words previously seen, and participants were instructed to say the full sentence with the correct word added (e.g., "I think a quick run picks you up"). These spoken

[1]. One plausible indicator of casual speech is speech rate. Flint (2003) showed that recording condition had a significant effect on speech duration. That is, stimulus duration in the "memory" condition (mean = 773 ms) was shorter than for stimuli recorded in the "reading" condition (mean = 1279 ms). Moreover, stimulus duration was inversely correlated with confusion rate (Pearson's $r = -0.41$, $p<.05$); thus, the faster, more casual speech yielded more ambiguous assimilations.

responses were recorded digitally via a microphone, which the participants believed was for checking whether the correct word had been inserted.

Participants

The participants were thirty-two members of the University of York participant panel (15 male, 17 female), and were paid for their time. All had English as a first language, and all but one (South African) had UK accents.

Stimuli

The key materials (see Appendix) were 40 sets of sentences previously used in Gaskell and Marslen-Wilson (2001; Experiment 3). These sets were based on word-pairs such as *run* and *rum*, which deviate only on the place of articulation of the final phoneme, and for which the coronal form could assimilate to become like the noncoronal form (see Table 1). The coronal forms had a mean occurrence frequency of 58 per million (median 10 per million), whereas the noncoronal forms tended to be of lower frequency (mean 17, median 5). Of the 40 word pairs, 10 ended in voiced stops, 17 in unvoiced stops and 13 in nasal consonants. The phoneme following the critical word matched the place of either the noncoronal form (providing a viable context for assimilation) or the coronal form (providing an unviable context for assimilation). The key sentential contexts (including preceding and following words) were assessed for any preferential bias towards either member of the critical word pairs by collecting participants' ratings of sentential plausibility (see Gaskell & Marslen-Wilson, 2001). These ratings showed that all conditions were rated as plausible, with no significant differences between conditions in terms of plausibility (mean ratings per condition were between 5.7 and 5.9 on a 7 point scale). Two further conditions: the control sentences from Gaskell & Marslen-Wilson (2001), and alternative versions of the test sentences encouraging assimilation away from the noncoronal item (as in *run gets*) were also used in the dummy experiment but not in the main experiments.

The sentences were presented in the dummy experiment with one word replaced by a numbered space (e.g., "I think a quick run picks you (3) "; replaced word = *up*). The position of the replaced word in the sentence varied between trials, but was always separated from the potential assimilation by at least one word. Two words similar to the replaced word in meaning and syntactic class were selected as foils (e.g., *out*, *in*). To provide more variety in the composition of the sentences, and to provide filler materials for the main experiments, the test sentences were interspersed with 198 filler sentences, again taken from Gaskell and Marslen-Wilson (2001). This made a total of 438 sentences for presentation (40 test sentences x 6 conditions + 198 fillers). To ensure that the speakers did not become aware of the minimal pairs embedded in the experiments and alter their speaking style to compensate, each participant was only presented with one third of the materials. The test sentences were split up into three lists such that for any one item only two of the 6 conditions were encountered by a participant, and the conditions were paired and counterbalanced so that the paired conditions made use of the same test word (e.g., *I think a quick run picks you up* was paired with *I think a quick run does you good*). To avoid contrastive stress effects, the paired test sentences were presented in different halves of the experiment, separated by a short break.

Procedure

Participants were tested individually in a small, quiet testing booth. They were instructed that the experiment examined memory for words, and whether sentence context can affect memory performance. The written stimuli were delivered using

DMDX software (Forster & Forster, 2003) and an Acer laptop computer, which was chosen because of its low acoustic noise output. Responses were recorded using a Sennheiser Microphone linked to the laptop via a Creative Extigy external soundcard. The recording level was set individually for each participant using DMDX, and then 10 practice trials were run, followed by the main experiment divided into two blocks of 73 stimuli. On each trial, three numbered words were presented on the screen one above the other for 3 s. The screen was then blanked for 2 s before the test sentence was presented with a number (1-3) inserted in place of one word. Participants were instructed to report the full sentence, with the word referred to by the number inserted in the correct place. They were informed that their response would be recorded by the computer to check whether they had completed the sentence in the appropriate way. On presentation of the sentence, participants' verbal responses were recorded on the laptop (sample rate 22 kHz, 16 bit precision) for 6 s, and then the next trial began. Trials were presented in a separate pseudorandom order for each participant.

After the experiments, participants were debriefed on the real nature of the experiment, and were told that they could choose to have their recordings deleted. None of the participants had identified any ulterior motive for the experiment before that point, and all participants consented for their recordings to be used in the following experiments.

Stimulus post-processing

Thirty-two participants were recorded on the three lists (between 10 and 11 per list), generating a total of 4992 six-second sound files. Speech files were normalized for peak amplitude, and Adobe Audition 1.0 was used to trim automatically the silent pauses preceding and following the spoken sentences, with some files being retrimmed by hand. The test files were then tagged for clear problems such as dysfluency, poor quality recording (e.g., oversampling), or incorrect responses. The remaining recordings (77%) were then ready for use in the following experiments.

Experiment 1

Experiment 1 examined whether the stimuli recorded as described above would create lexical ambiguity in identification. Following the presentation of a spoken sentence, participants were given two choices as to the identification of the target word (e.g., *run* vs. *rum*). Participants then had to select the appropriate response as quickly as possible. If assimilation creates lexical ambiguity, then this should be reflected in delayed responses and/or greater miscommunication rates. The main experiment involved the large set of sentences from the post-processing described above, which were presented to members of our participant panel. A smaller post-test made use of selected recordings and expert phoneticians in order to provide a more detailed analysis of the properties of strongly assimilated segments.

Participants

Twenty members of the University of York participant panel were tested. The participants were mainly undergraduate students aged between 18 and 25, and took part either for course credit or for a fee.

Stimuli and Design

Experiment 1 made use of participants' recordings described above for the conditions marked in Table 1. These conditions represented the four combinations of two independent variables. Target words were recorded with the intention of producing either a coronal or a noncoronal consonant (the *underlying place* variable), and the following word varied with respect to whether it made coronal to noncoronal place assimilation viable (the *viability* variable). It was expected that production and/or perception of potential assimilation would vary according to the *final consonant type* (voiced stop, unvoiced stop, or nasal), and so stimuli were coded with respect to this variable as well. The sentence *speaker* was used as an independent variable in some subsidiary analyses, reflecting the expectation that speakers would vary in assimilation style. All 1317 recordings of test items that passed the post-processing filter stage described above were used, with only small differences between the key conditions in terms of numbers of stimuli (see Table 2). The stimuli were split into twelve blocks of similar composition, with participants tested in three sessions of four blocks. The written forms of the two target words related to each sentence were used to indicate the two possible responses. Allocation of targets to keys was counterbalanced across two lists such that there was no systematic relation between target word type and response key.

Procedure

Participants were tested individually in quiet testing rooms. On each trial a spoken sentence was presented via Beyerdynamic DT770Pro headphones, followed immediately by the two response choices for that sentence (e.g., *run* vs. *rum*). These were presented in upper case either side of the centre of a computer screen. Participants made timed responses by pressing one of two buttons on a Trust Gamepad to indicate which of the two words they believed they had heard. Response times were measured from onset of the visually presented words using DMDX (Forster & Forster, 2003). After a participant had responded or a 5-s timeout had been reached there was a 450-ms inter-trial interval, followed by the next sentence. Each of the three sessions began with a 20-trial practice session, followed by three blocks of 110 trials and one block of 109 trials. Most participants chose to spread their three sessions over the course of two days, but the only stipulation was that each session break should last at least an hour. The order of presentation of the blocks was counterbalanced across participants, and within each block order of presentation was randomized individually for each participant.

Results

All responses were coded in terms of response choice (coronal or noncoronal-final target) and response time (see Table 3). Because of the wide variation in numbers of RT datapoints between conditions, our analyses focused on response choice.[2]

---

[2]. Analyses of the response times were restricted to conditions in which the speaker's intended target matched the listener's selected target (the closest equivalent of correct responses in this kind of experiment), with separate two-way ANOVAs on each target type. For the coronal targets there was a marginal effect of final consonant type ($F1[2,38] = 2.9$, $p = .065$; $F2[2,676] = 2.6$, $p = .072$), with slightly faster responses to words ending in voiced stops (530 ms) than words ending in nasals (541 ms) or unvoiced stops (544 ms). There was no main effect of viability ($F1 < 1$; $F2[1,676] = 3.4$, $p = .064$), but the interaction of these two variables was significant by participants and marginal by items ($F1[2,38] = 5.4$, $p < .01$; $F2[2,676] = 2.6$, $p = .077$). In cases where participants correctly identified words ending in coronal consonants there was a tendency for participants to be slowed in the assimilatory context for targets with unvoiced stops (29 ms), but not for voiced (-5 ms) or nasal (-10 ms) consonants. Similarly, for noncoronal targets, the main effects of viability ($F1[1,19] = 3.1$, $p = .095$; $F2 < 1$) and final consonant type ($F1 < 1$; $F2[2,518] = 2.1$, $p > .1$) were nonsignificant, but there was a marginal interaction between the two ($F1[2,38] = 6.2$, $p < .01$; $F2[2,518] = 2.5$, $p = .086$). Again this appeared to indicate that a viable context for assimilation delayed correct responses for the unvoiced stop stimuli (28 ms), but not for the nasals (-7 ms) or the voiced stops (1 ms).

Three way ANOVAs on the data using the within-participant (but between-item) variables underlying place, viability and final consonant type revealed the following statistically significant effects. Unsurprisingly, underlying place had a strong effect on response choice ($F1[1,19] = 1543.0$, $p < .001$; $F2[1,1305] = 2366.8$, $p < .001$), with overall 80.2% coronal responses when speakers intended to say words ending in coronal consonants, and 20.2% coronal responses when they did not. There were also main effects of viability ($F1[1,19] = 41.4$, $p < .001$; $F2[1,1305] = 13.7$, $p < .001$) and final consonant type ($F1[2,38] = 11.1.0$, $p < .001$; $F2[2,1305] = 3.9$, $p < .05$), which are more easily interpreted in conjunction with the interactions between the three variables. Most crucially, there was a strong interaction between viability and underlying place ($F1[1,19] = 499.5$, $p < .001$; $F2[1,1305] = 55.0$, $p < .001$). When speakers intended to produce targets with coronal final segments (e.g., *run*), there was a marked effect of viability (a 13.7% shift), with viable contexts for assimilation tending to impair recognition of the intended place of the final consonant. On the other hand, when speakers intended to produce targets with noncoronal final segments there was a smaller (4.5%) shift in the opposite direction. This can be interpreted as a disguising of the intended place of articulation by a context that meets the conditions required for place assimilation.

There were also two-way interactions between final consonant type and both viability ($F1[2,38] = 21.6$, $p < .001$; $F2[2,1305] = 3.4$, $p < .05$) and underlying place ($F1[2,38] = 230.4$, $p < .001$; $F2[2,1305] = 34.5$, $p < .001$), and a further three way interaction between all three variables ($F1[2,38] = 72.1$, $p < .001$; $F2[2,1305] = 8.4$, $p < .001$). These effects can be understood in terms of the differing masking effects of assimilatory context for different types of final consonant (see Figure 1). For voiced stops, there was little or no effect of contextual viability on the recognition of the targets. For nasal consonants there was a strong effect for underlying coronal consonants but not for underlying noncoronal consonants, whereas for unvoiced stops the effects occur symmetrically for both types of target.

Phoneticians' Posttest

A subset of strong assimilations was selected for more detailed investigation. These 36 stimuli had an intended coronal place for the final consonant of the key word, with a viable context for assimilation, and responses to the main test indicated that strong assimilation had taken place (mean coronal scores varied from 0% to 45%, with a mean of 24%). Although these stimuli tended to be misperceived (treated as noncoronal-final) in a speeded choice task, it is possible that some subtle but detectable cues to the true place of articulation of the assimilated consonant remained. Nolan (1992) used considered judgments of expert phoneticians as a means of addressing this question, and found that they were more sensitive than either first-pass responses or electropalatographic data in terms of discerning the true underlying place.

We therefore asked phoneticians to rate these strong assimilations using a procedure similar to Nolan (1992). The 36 strong assimilations were all good quality recordings with the critical consonant being clear (and unreleased in the case of stop consonants). They were intermixed with 18 comparison sentences drawn from the same population of speakers and critical words. The comparison sentences had viable following context for assimilation, were of similar sound quality, and did not contain released stops or extended gaps between the critical consonant and the following segment. Half of these had an intended coronal place for the final consonant of the key word and 100% coronal responses in the main test. In other words, these items were unassimilated or very weakly assimilated, and were expected to be relatively easy for phoneticians to decide on, providing a baseline for the strong assimilations. The remaining 9 items were all intended noncoronals, and acted as catch trials in the phoneticians' test (i.e., they should be treated as noncoronals but could be misperceived as strong assimilations). The mean coronal score for these items was 29% (range 5-75%), similar to the 24% for the strong assimilations. In sum, the weakly assimilated stimuli should tend to be selected as having coronal underlying place, and the underlying noncoronals should be treated on the whole as noncoronals. The question of interest was where the strong assimilations would lie on this continuum.

Eight expert phoneticians took part in the test, selected from the Linguistics Departments of the Universities of York and Sheffield. Five of the phoneticians were native British-English speakers, with one US English, one German and one Greek native speaker. All were fully trained in phonetics and were career phoneticians at varying levels. The phoneticians were knowledgeable about English place assimilation, and were told that they would hear sentences in which different speakers unwittingly generated potential assimilations as responses in a memory experiment. Their task was to decide which of two alternatives of a minimal pair shown on the computer screen was intended by the speaker. First pass and considered judgments were collected, as in Nolan (1992). On each trial in the first section, a sentence was presented once via headphones. A pair of words then appeared on screen and participants had to pick the one that they thought the speaker intended to say, with responses collected via button-press. The second section followed the same overall procedure except that listeners had the option to replay the speech as many times as they wished prior to responding with their considered judgment. In this section participants also rated their response on a 1-9 confidence scale (1 = not confident at all about response given; 9 = highly confident about response). There was no time constraint in either task, and participants were encouraged to take as long as they liked to decide. The experiment lasted roughly 20-25 minutes.

The results of the posttest were analyzed using planned comparisons in by-participants ANOVAs, and showed that the strong assimilations remained highly ambiguous under scrutiny from expert phoneticians (see Table 4). On the first pass, without an option to replay a sentence, strong assimilations were treated as coronal 35% of the time, which was significantly less than the coronal response proportion for weak assimilations (75%; $F[1,7] = 13.2$, $p < .01$) and identical to the noncoronal items ($F[1,7] < 1$). In their considered judgments, with the option to replay the stimuli as often as they liked, the experts were more inclined to identify coronal underlying consonants, especially for the weak assimilations, which were pushed up to near ceiling levels (94%). However, there was no evidence that the strong assimilations became unambiguous: these were responded to as coronals 45% of the time, which was significantly less than for the weak assimilations ($F[1,7] = 67.3$, $p < .001$) but again no different from the underlying noncoronals ($F[1,7] <1$). The confidence ratings confirmed the greater degree of uncertainty generated by strong as opposed to weak assimilations (strong vs. weak: $F[1,7] = 13.6$, $p < .01$; strong vs. noncoronal: $F[1,7] < 1$). Finally, this increased uncertainty for strong assimilations was also reflected in an increased tendency to replay the stimuli (strong vs. weak: $F[1,7] = 9.1$, $p < .05$; strong vs. noncoronal: $F[1,7] = 2.0$; $p > .1$).

Discussion

Experiment 1 examined the extent to which naïve speakers produce strong assimilation of coronal consonants. Given a semantically neutral context and the presence of a lexical competitor matching the assimilated form, such strong assimilations should create lexical ambiguity, leading to increasing miscommunication. The results for the coronal targets suggest that this is the case. Overall, miscommunication rate more than doubled from 13% to 27% when the following segments provided the appropriate conditions for assimilation to occur. Context effects were strongest for nasals (20% viability effect) and reasonably strong for unvoiced stops (17% effect), but minimal for voiced stops (4%). This gradient of assimilation effects fits in with the argument that assimilation tends to be strongest in situations where perceptibility is weakest (e.g., Kohler, 1990; Mitterer, Csépe, Honbolygo & Blomert, 2006; Steriade, 2001).

A possible explanation of the assimilation effect found here is that for some reason viable contexts for assimilation add noise to the communication process, rather than generating tokens that are consistently misperceived. The distributions of response selections for the coronal targets suggest that this is not the case. Each sentence used in Experiment 1 can be assigned a coronal response proportion score, by calculating the percentage of participants who selected the target word ending in a coronal consonant (e.g., *run*). Figure 2 (upper panel) shows the distribution of these scores for the two categories of consonants that showed context effects (unvoiced stops and nasals). Increasing levels of noise would imply that the increase in response frequency for viable as opposed to unviable contexts would be in or around the 40%-60% assimilation range, since these are the most likely outcomes given a random response selection. However, the shift in the response distribution between unviable and viable assimilations goes right across the range of responses, leading to a flattening of the distribution for viable assimilatory contexts. Most importantly, there is a marked effect of viability on the left hand side of the distribution: only 1% of coronal targets with unviable contexts were judged by at least 80% of the participants to be noncoronal, whereas for the viable contexts this figure is 11%. This flattened distribution with a substantial proportion of responses at the noncoronal end cannot be explained in terms of random noise. Instead these stimuli appear to be complete or near-complete assimilations, which despite the presence of a viable phonemic context for assimilation (and a frequency bias favoring the coronal form) are consistently treated at face value by listeners.

Further evidence that strong assimilations occur comes from the analysis of phoneticians' considered responses to selected stimuli. Given that the responses in the main experiment were speeded, it could be argued that any ambiguity caused by strong assimilations is transient, and so could be resolved given a little more time. Equally, it could be that there are subtle traces of the underlying coronal place in the strongly assimilated speech that the perceptual system could make use of after careful or repeated listening. We addressed this question by asking phonetics experts to listen carefully and repeatedly to some of the stimuli that appeared to be strongly assimilated and try to reconstruct the original underlying place. However, after listening to each stimulus as often as they wished, the phoneticians were at chance level (45%; comparison with chance level of 50%: $t[7] = -1.1$; $p > .1$) on identifying the underlying place of the strong assimilations. Furthermore, no individual phonetician stood out from the group as being better able to pick out strong assimilations from true noncoronals. This posttest is very strong evidence that the strong assimilations we chose are functionally identical to their labial and velar counterparts.

We cannot (and need not) go that small step further and state that the strong assimilations must be complete because it is almost impossible to rule out the case that some analysis of the strong assimilations would reveal some cues, however weak, relating to the intended coronal consonant. However, the issue of completeness in production is not important so far as we are concerned, Perceptually, the strong assimilations undeniably create ambiguity, and this ambiguity is not resolved swiftly, as the main experiments demonstrate. Nor is this ambiguity resolved after more time, or with more presentations, even to the expert ear.

The presence of such ambiguity, above and beyond the normal levels of confusion between minimal pairs, is not predicted by Gow's (2002) feature parsing model, which relies on residual cues to the underlying place of an assimilated consonant to resolve ambiguity in all cases. In contrast, the presence of these strong assimilations in speech is not problematic for the learned compensation model of Gaskell and Marslen-Wilson (1996, 2001; Gaskell, 2003) because the model can make use of other cues to compensate for assimilation, alongside any information contained in the assimilated consonant. This sounds like a rather odd conclusion to draw, given that in the extreme cases just highlighted listeners on the whole *did not* compensate for assimilation. However, this was because the conditions were carefully engineered such that sentential context and lexical knowledge could not be made use of. The intention was not to demonstrate that assimilation will commonly cause miscommunication, but was instead to demonstrate that listeners cannot always rely on coarticulatory information in partial assimilations. In more normal circumstances there would be other cues available to prevent miscommunication; Experiment 2 looks at how sentential and phonemic context might perform this role.

The second key aim of the current study was to examine whether listeners compensate for assimilation even in cases where no acoustic cues to coronal place are present (that is, when speakers intentionally produce labial and velar consonants). Experiment 1 on its own cannot answer this question, but it does provide a few clues. Noncoronal targets were predominantly perceived as intended, with little effect of following context on miscommunication. Nonetheless, some effects of context were found, with unvoiced stops showing a 12% increase in miscommunication for contexts that superficially supported assimilation. Furthermore, the phoneticians showed quite a strong tendency to identify underlyingly coronal place for a subset of these stimuli. These results—if they reflect a perceptual effect rather than some production difference between the stimuli—do not fit in with a feature parsing account of assimilation perception, which relies on information about two places of articulation within the same segment (Gow, 2002, 2003). The simple cooccurrence of two noncoronal consonants should not lead to any compensation for assimilation because the first consonant does not contain any coronal information to be separated from the noncoronal information. On the other hand, a weak effect of contextual viability is easier to explain in a learned compensation account (Gaskell, 2003). Despite the overwhelming bias in favor of treated noncoronal consonants on face value, there should be a small tendency to compensate for assimilation when the noncoronal segment is followed by a viable context, as opposed to an unviable one.

However, returning to the response distributions, the case for a true effect of viability for underlying noncoronal targets is less clear (Figure 2, lower panel). Here, the shift in the distribution for unvoiced stops and nasals attributable to the contextual viability is less extreme, with no difference at the right-hand end of the distributions. The distribution of responses for the noncoronal targets could be explained in terms of an addition of random noise in the responses associated with a viable context for assimilation. More concretely, given that the effect is only evident for unvoiced stops, the requirement to articulate two bilabial or velar stops in sequence (as in *ripe berries*) may have decreased the likelihood of a released stop, and reduced the information about place of articulation available to the listener, provoking more guesses.

Experiment 2 tested whether words ending in normal labial or velar consonants (and so not susceptible to assimilation) could be treated as assimilated when followed by a viable assimilatory context. If the effect found in Experiment 1 was a production effect, then it should have left its mark on the articulation of the final consonants of the noncoronal target words (i.e., the final consonants in the viable condition should be less clear tokens of noncoronal consonants). On the other hand, if a perceptual effect is observable, then the following context of the targets should be crucial. To assess these possibilities, Experiment 2 made use of a smaller set of noncoronal stimuli varying in terms of the place of the following context. To evaluate the perceptual influence of the following context independently of any potential influence of articulation of the final consonants, these stimuli were presented both intact and cross-spliced, with cross-splicing position being the onset of the segment following the target word. If assimilation perception mechanisms have no role to play in the identification of these targets, then the only effect of cross-splicing should be a possible delay due to artifacts created by the cross-splicing process. On the other hand if following context does have a role to play in the perception of these stimuli, then splicing *picks* onto the *rum* from *rum does* should make the word more ambiguous (more like *run*), whereas splicing *does* onto the *rum* from *rum picks* should have the opposite effect.

Experiment 2 again used a two-alternative word selection task and assessed the influence of following context in semantically neutral circumstances. However, unlike Experiment 1, a second condition used the same task but with the spliced sentences preceded by sentence contexts providing a semantic bias towards the noncoronal item (e.g., *run*). Gaskell and Marslen-Wilson (2001) showed that the neutral sentences did not show contextual viability effects on priming when presented alone. However, when the same sentences were preceded by sentences with a closer semantic fit to the word ending in a coronal segment (e.g., *It's best to start the day with a burst of activity. I think a quick run/m picks you up*) clear viability effects emerged. Gaskell and Marslen-Wilson used a non-naïve speaker whose intention was to produce stimuli that could be perceived as assimilated (and so could have contained residual coronal cues), Experiment 2 examined whether these effects can still be observed when naïve speakers are required to produce words ending in noncoronal contexts. Viability effects observed in these circumstances would be clear evidence for a probabilistic model of compensation for assimilation (Gaskell & Marslen-Wilson, 2001; Gaskell, 2003).

Experiment 2

Participants

Ninety-four members of the University of York participant panel were tested, none of whom had taken part in Experiments 1. Of these, 46 took part in the semantically neutral subexperiment and 48 took part in the semantically biased subexperiment (see below). The participants were mainly undergraduate students aged between 18 and 25, and took part either for course credit or for a fee.

Stimuli and Design

Experiment 2 used a relatively standard design based around 40 item-sets recorded in the dummy experiment and tested in Experiment 1 (see Appendix). For each item-set, pairs of recordings of the noncoronal target (e.g., *rum*) sentences were selected for presentation in Experiment 2, with participants required to identify this target word. The selected pairs had the same speaker for each member of the pair, and were chosen on the basis of the responses in Experiment 1, avoiding cases where there was good agreement between participants on the underlying consonant of the target word in the viable condition. This selection criterion made effects of following context possible, by avoiding ceiling and floor effects. The selected pairs were all fairly good quality recordings, and cases where there was a clear release of a target-final stop or a clear pause following the target word were avoided. Alignment points for cross-splicing were identified using Adobe Audition, with the alignment point being the onset of the segment following the target word. Sentences for which this was difficult to identify or where splicing caused a perceivable artefact were eliminated and replaced. The final set of recordings came from 16 different speakers, who each contributed to up to 5 test sets.

The selected stimuli were reasonably representative of the full set of materials in Experiment 1, as indexed by their response scores. For the full set of stimuli in Experiment 1, 24% of noncoronal targets with a viable context were identified as coronal, as opposed to 18% for the noncoronal targets with unviable contexts. For the selected stimuli, the equivalent figures are 32% in the viable context condition and 18% again for the unviable contexts. This difference is to be expected if some compensation for assimilation occurs for normal labial and velar segments in appropriate fluent contexts. The key question was whether these effects would remain when following context was cross-spliced, with the speech following the target word swapped between the viable and unviable conditions. This manipulation allowed a purer test of the effect of following context on the perception of the target stimuli. Thus, four versions of each test item were used in Experiment 2, varying in terms of the nature of the context following the assimilated speech (original viability) and presence or absence of cross-splicing.

For one set of participants, these materials were presented as in Experiment 1 with semantically neutral sentential context. For a second set, the same materials were used but with each sentence preceded by a semantically biasing sentence. The biasing materials were taken from Gaskell & Marslen-Wilson (2001), and had a closer semantic fit to the coronal interpretation of the prime word, while not ruling out the non-coronal alternative (see Gaskell & Marslen-Wilson, 2001, p. 338, for ratings data confirming the bias effect). Twelve of these sentences were adapted in order to strengthen their semantic influence (see Appendix). No single word in the biasing sentence had a strong semantic or associative link with the coronal target. The biasing sentences were recorded afresh in a soundproof booth using the first author as a speaker, and transferred digitally to computer. Because the speaker for the first sentence never matched the second sentence speaker, the sentences were less mutually coherent than would be the case in normal utterances. This may have acted to reduce the influence of the semantic bias, but could not have been confounded with any effect of phonemic context observed.

Filler trials were constructed from the same pool of recordings as the test items, with the same mix of speakers, but were tailored to the requirements of the task. Because all 40 test items contained words that were intended to match the noncoronal targets, a similar number of filler trials (35) used targets with coronal final consonants. For all of these items, the alternative choice differed only on the final consonant, with 15 of them being potential assimilated minimal pairs (e.g., ROD/ROB). A further 60 fillers, employed a target monosyllable and alternative word differing only on the final consonant(s). Twenty practice sentences with a similar composition were selected. All of these sentences were sampled from the same group of speakers as the test materials, and for the semantically biasing conditions were preceded by the appropriate sentence from Gaskell and Marslen-Wilson (2001), again recorded afresh. Test stimuli were counterbalanced across four lists within each subexperiment such that all participants encountered the same proportion of trials in each condition and no participant encountered any target more than once.

Procedure

The procedure was mostly identical to that of Experiment 1. However, the experiment was run as a single session, with 20 practice items followed by the test and filler items, divided into 4 blocks. For the semantically biased subexperiment, there were also differences in the composition of each trial, with the target sentences were preceded by a biasing sentence and a 750 ms inter-stimulus interval. In all cases, immediately following the auditory target sentence, two response alternatives were presented in capital letters to the left and right of the centre of the screen and the listener was required to select the appropriate response with either the left or right button on their keypad.

Results

A technical error meant that the data for one participant were lost, leaving data for between 11 and 12 participants in all 8 lists. Responses were coded in terms of response choice (coronal or noncoronal-final target) and response time (see Table 5).

A three-way ANOVA on the response choice data using the within-participant, within-items variables original viability (i.e., before cross-splicing) and cross-splicing (intact vs. cross-spliced following context), and the between-participants, within-items variable semantic bias (presence or absence of preceding biasing sentence), revealed a main effect of semantic bias ($F1[1,91] = 91.7$, $p < .001$; $F2[1,39] = 84.8$, $p < .001$). Without the biasing sentences, only 36% of the critical consonants were (incorrectly) categorized as coronal. However, when the semantic context was a more appropriate for the coronal-final word, this false positive rate rose to 56%. There was also a significant interaction between original viability and cross-splicing ($F1[1,91] = 32.9$, $p < .001$; $F2[1,39] = 20.0$, $p < .001$). This interaction (see Figure 3) showed that the viability of the following context of the target words had a substantial influence on participants' identity judgments. Collapsing the data across levels of semantic bias, participants were 14% more likely in the unspliced condition to identify an intended noncoronal target (e.g., *rum*) as underlyingly coronal if the following context appeared conducive to assimilation. This difference was equivalent in size to the effect found for these materials in Experiment 1 (32% vs. 18%), and is not necessarily a perceptual effect of the phonemic context. Instead, it could be that the selected materials contained less clear final segments in the viable condition than the unviable one. Nonetheless, when through cross-splicing a viable context was replaced with an unviable one there was a 7% reduction in coronal responses ($F1[1,91] = 7.9$, $p < .01$; $F2[1,39] = 7.6$, $p < .01$), implying a reduced tendency to compensate for assimilation when the following context was incompatible. On the other hand, replacing an unviable context with a viable one increased by 12% the tendency to treat the key words as ending in a coronal consonant ($F1[1,91] = 21.6$, $p < .001$; $F2[1,39] = 19.4$, $p < .001$). This interactive pattern is clear evidence that the materials in Experiment 2 evoked perceptual compensation for assimilation.

None of the other main effects or interactions approach significance levels in either analysis, with the exception of the main effect of original viability, which was significant by participants but not items ($F1[1,91] = 6.3$, $p < .05$; $F2[1,39] = 2.5$, $p > .1$).

As in Experiment 1, the response time data were of secondary interest. Response times were longer than in Experiment 1; grand average response times were over a second for both types of response, whereas response times using the same task in Experiment 1 were typically 500-600 ms. There are several potential reasons for the slower response times in this experiment, but perhaps most obviously the level of competition between the two response alternatives was greater, as reflected in the more even split between the two alternatives. This factor also meant that analyses of both types of response had large amounts of missing data and were underpowered.[3]

Discussion

Experiment 2 demonstrated that compensation for assimilation can in some cases lead to miscommunication where no assimilation is present. The key stimuli in this experiment were all produced by the speakers as intended, with final bilabial or velar consonants. However, the following segmental context of the stimuli had a significant influence on the perception of these stimuli, with a clear tendency to compensate for assimilation if the context happened to fit the normal circumstances in which assimilation occurs. This result is important because it demonstrates that context-dependent compensation for assimilation occurs even in cases where the key consonant is not in fact assimilated. Previous demonstrations of compensation for assimilation have either used real assimilated speech (e.g., Gow, 2001, 2003) or have used a non-naïve speaker who intended to produce stimuli that could be assimilated (e.g., Gaskell & Marslen-Wilson, 1996, 2001). For the stimuli used in Experiment 2, naïve speakers read words ending in consonants that do not assimilate place of articulation (i.e., labial and velar consonants). Despite this, listeners compensated for assimilation when the following context matched in place.

This finding is again problematic for the feature parsing model of Gow (2003), which relies on assimilation being partial, and therefore leaving traces of the intended segment in the realization of the assimilated consonant. In Experiment 2, the key segments did not undergo assimilation, and so could not have contained any partial assimilatory residue. One might perhaps counter that noncoronal segments articulated prior to coronal segments (e.g., the /m/ in *rum does*) could have taken on some characteristics of the following consonant, and so generated segments with cues to two places of articulation. However, Experiment 2 also included stimuli in which the noncoronal segments were articulated followed by consonants with the same place (e.g., *rum picks*), and the perception of these stimuli was also affected by contextual

---

3. For the noncoronal response times, which might be considered the correct responses despite being in the minority, there was a main effects of semantic bias ($F1[1,91] = 17.4$, $p < .001$; $F2[1,27] = 102.4$, $p < .001$), with faster responses in the neutral condition, when the semantic context did not bias the competing response. Otherwise, there were no fully significant main effects or interactions, although the effect of original viability reached significance level in the participants analysis ($F1[1,91] = 5.6$, $p < .05$; $F2[1,27] = 3.0$, $p > .1$). This is suggestive of a trend for clearer phonetic evidence in favor of a noncoronal consonant for tokens originally produced in the context of a following coronal consonant, as hypothesized in the Discussion section of Experiment 1. For the coronal response times, again only the main effect of semantic bias reached significance level ($F1[1,82] = 7.2$, $p < .01$, $F2(1,27) = 22.3$, $p < .001$), with slower responses in the semantically biased condition when competition between responses was more balanced (i.e., closer to a 50:50 split between responses).

viability. In combination, these results suggest that compensation for place assimilation can occur even for consonants that have no phonetic evidence of an underlying coronal place.

The influence of semantic context in Experiment 2 was additive with contextual viability (see Figure 3), and viability effects were observed in the absence of a semantic bias towards the word ending with a coronal consonant. This finding clarifies an uncertainty derived from the previous study to make use of these stimulus sentences (Gaskell & Marslen-Wilson, 2001). Gaskell and Marslen-Wilson found no contextual viability effect in the absence of a sentential bias (Experiments 1 and 2), whereas when the biasing sentence was present (Experiment 3), a viability effect was observed. This description suggests that some semantic evidence in favor of a coronal form (e.g., *run*) is required for contextual viability effects to emerge. However, no interaction between these two variables was found in a cross-experiment comparison. The current results suggest that the failure to observe viability effects in the absence of a semantic bias in the 2001 study was most likely due to the difficulty of observing differences between weakly activated words (and hence nonsignificant test-control differences) using cross-modal priming. The more direct comparison allowed in a 2AFC test has shown that contextual viability is still influential even when lexical candidates are relatively weakly supported by their fit with the sentence context.

General Discussion

The two experiments reported in this paper provide a set of insights into the compensation processes underlying the perception of assimilated speech, and particularly the ability of these processes to deal with extreme assimilations. Experiment 1 elicited naïve speakers' assimilations, showing that the presence of assimilatory context causes ambiguity. The rate of miscommunication between naïve speakers and listeners more than doubled when words ending in coronal consonants were followed by a viable context for assimilation, as opposed to an unviable context. Furthermore, examination of the distribution of responses showed that the continuum of English place assimilation is sufficiently extended in some cases to leave a group of listeners in unanimous agreement that a different word was produced and to leave phoneticians unable to discern their true underlying place at levels better than chance. The key point from the experiment is that some assimilated coronal segments are effectively indistinguishable from their noncoronal counterparts. Such extreme assimilations will result in miscommunications unless other contextual factors can mitigate their effects in perception.

Experiment 2 showed that listeners make use of following context and semantic biases to compensate for assimilation, even in cases where no assimilation has occurred: when listeners utter words ending in labial or velar place and the following segment happens to have the same place of articulation. These combined results imply that compensating for assimilated speech is more than the simple extraction of partial cues to the underlying place. Strong assimilations do occur, and the human system is able to deal with them.

The title of Gow's (2002) article was "Does English coronal place assimilation create lexical ambiguity?", and the results suggested not: words such as *right* and *ripe* did not lead to lexical ambiguity when placed in potentially assimilating context. The current data are obviously in conflict with Gow's (2002) finding. It is difficult to be certain about the cause of this difference, but the procedure for eliciting the stimuli is likely to be partly responsible (cf. Flint, 2003). Whereas Gow (2002) used a single speaker, who was aware that the recordings would be used for a speech perception experiment, the current research used multiple speakers who were entirely naïve about the goal of the recording, and would most likely have produced utterances more like spontaneous speech (cf. Mitterer & Ernestus, 2006 for similar conclusion relating to /t/-lenition). The circumstances we created were by no means optimal for studying casual speech: there was also no attempt to induce fast speech, or even to ensure that speakers spoke fluently. It was evident from listening to the recordings that several speakers used an exaggeratedly clear and punctuated style (possibly to aid the experimenter in assessing their responses). Despite this, there was clear evidence of strong assimilation creating lexical ambiguity.

On the understanding that place assimilation does not create lexical ambiguity, Gow (2002) proposed that compensation for place assimilation is a matter of extracting the featural cues to the place of articulation underlying the partial assimilation. Mitterer and colleagues (Mitterer & Blomert, 2003; Mitterer et al., 2006) argue for a similar, low-level account of perceptual compensation. Although the precise mechanism of compensation is less explicit for the latter model, it also suggests that compensation for assimilation relies on relatively weak assimilation, and highlights the role of phonetic detail in the compensation process. The identification of cases in which place assimilation does create lexical ambiguity severely weakens these models. In these cases the best a feature parsing model can do is identify the segment incorrectly at a pre-lexical level irrespective of phonemic context (e.g., a strongly assimilated *run* would be misidentified as *rum*), and then use lexical or semantic information to correct this error at a later stage. This possibility does not fit with the results of Experiment 2, which showed that listeners use phonemic context in their evaluation of canonically uttered bilabial and velar segments irrespective of whether semantic context favored the coronal form. The feature parsing approach would have no problem in explaining how a higher-level semantic bias could cause an overall shift in activations, but it cannot explain how the contextual viability of the following context should also affect response selection, given that feature parsing only applies to partial assimilation in which cues to alveolar place remain. Thus it seems that a feature parsing or similar low-level model does not predict either of the two main findings: that complete place assimilation does occur, creating lexical ambiguity, and that compensation for assimilation occurs for some segments that contain no featural cues relating to underlying coronal place.

Returning to the underspecification model (Lahiri & Marslen-Wilson, 1991; Lahiri & Reetz, 2002), the observation of strong assimilation in Experiment 1 is consistent because a system that deals with dichotomous representations will tend to predict that discrete or complete changes should occur, and that these can be accommodated through underspecification of the assimilated segment. However, the observation of a contextual viability effect in Experiment 2 is more challenging for an underspecified account because following context should be irrelevant to the matching process (Wheeldon & Waksler, 2004). As described in the introduction, this viability effect fits with a large set of studies showing similar effects of context (Coenen, Zwitserlood & Bölte, 2001; Gaskell & Marslen-Wilson, 1996, 1998, 2001; Gow, 2001, 2002, 2003; Mitterer & Blomert, 2003). However, many of the previous findings could perhaps be explained in terms of lower-level compensation for partial assimilation as in feature parsing. By this account, underspecification may only apply in the case of complete assimilations. Indeed, the only published study (Wheeldon & Waksler, 2004) to show no influence of following word context required a naïve speaker to read sentences containing complete phonemic changes (e.g., *wickib prince* vs. *wickib ghost*). The results of Experiment 2 are therefore important in demonstrating that contextual viability is influential even in such extreme circumstances, and once again weakens the argument for a model based on underspecification. We can only speculate as to why we observed these effects when Wheeldon and Waksler did not, but it may be important that our stimuli for Experiment 2 were carefully chosen to ensure that the relative timing of the key consonants was appropriate for a potential assimilation. We do not wish to claim that compensation for assimilation *always* operates when two labial or velar segments co-occur, only if they are articulated in reasonably swift succession.

The probabilistic learned compensation model (Gaskell, 2003; Gaskell & Marslen-Wilson, 2001) has no problem in accommodating the current data because it incorporates a learning component, allowing the observation of phonetically extreme assimilation to trigger an equivalent compensation process. The extent of compensation in this model depends on other factors as well as segmental following context. In particular, the use of preceding sentential context to enhance the likelihood of compensation for assimilation fits well into this account.

However, the learned compensation model has been criticized as being incompatible with some studies suggesting that compensation for assimilation is a language-specific process. A clear prediction of this type of model is that listeners should compensate more for assimilations that they have experienced often (i.e., those that occur in their native language) than those that they have little experience of (that occur in other languages, but not their own). However, the data relating to this prediction are mixed. Some studies have found little or no difference between speakers of different languages in their perception of assimilated speech (Gow & Im, 2004; Mitterer et al., 2006), whereas a second set of studies have found marked divergences (Darcy, Peperkamp & Dupoux, in press; Darcy, Ramus, Christophe, Kinzler & Dupoux, in press; Lee, 2005). The most obvious discrepancy between these studies is in terms of the use of higher level knowledge (Darcy, Ramus et al., in press). For example, Gow and Im (2004) found no evidence of language specific effects when participants were required to monitor for phoneme targets in isolated VCCV sequences. On the other hand, Darcy, Ramus et al. (in press) found clear differences in the way American and French listeners compensated for voicing and place assimilation using a word detection task for stimuli embedded in sentences. Thus, the studies finding only language universal processing used experimental set-ups that encouraged focus on the lower-level of processing, in which language-universal compensation for coarticulation is dominant. On the other hand, the studies showing differences between groups of speakers in terms of their treatment of assimilation focused on the higher-level processes, in which listeners may learn to compensate for assimilation based on the statistical properties of their particular language.

In any case, the main contribution of the current research is to provide support for the higher-level accounts (e.g., Gaskell & Marslen-Wilson, 2001; Gaskell, 2003). Natural English place assimilation can be strong enough to create stimuli that are frequently misperceived as their noncoronal counterparts. Furthermore, words produced in their canonical forms can trigger compensation for assimilation where no such assimilation exists. These findings provide strong evidence against the idea that compensation for assimilation amounts to no more than a matter of separating two sets of features in coarticulated speech.

References

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

Browman, C. P. & Goldstein, L. (1992). Articulatory phonology: an overview. *Phonetica, 49*, 155-80.

Coenen, E., Zwitserlood, P., & Bölte, J. (2001). Variation and assimilation in German: Consequences for lexical access and representation. *Language and Cognitive Processes, 16*, 535-564.

Darcy, I., Ramus, F., Christophe, A., Kinzler, K., & Dupoux, E. (in press). Phonological knowledge in compensation for native and non-native assimilation. In F. Kügler, C. Féry & R. van de Vijver (Eds.) *Variation and Gradience in Phonetics and Phonology.* Berlin: Mouton De Gruyter.

Darcy, I., S. Peperkamp & E. Dupoux (in press). Bilinguals play by the rules: Perceptual compensation for assimilation in late L2-learners. In J. Cole & J. Hualde (Eds.) *Laboratory Phonology 9*. Berlin: Mouton De Gruyter.

Davis, M. H., Marslen-Wilson, W. D. & Gaskell, M. G. (2002). Leading up the lexical garden-path: Segmentation and ambiguity in spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 218-244.

Deelman, T., & Connine, C. M. (2001). Missing information in spoken word recognition: Nonreleased stop consonants. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 656-663.

Ellis, L., & Hardcastle, W. J. (2002). Categorical and gradient properties of assimilation in alveolar to velar sequences: evidence from EPG and EMA data. *Journal of Phonetics, 30*, 373-396.

Eulitz, C., & Lahiri, A. (2004). Neurobiological evidence for abstract phonological representations in the mental lexicon during speech recognition. *Journal of Cognitive Neuroscience,* 16, 577-583.

Flint, S. (2003). *Does place assimilation create lexical ambiguity?* Unpublished master's thesis, University of York, York, UK.

Forster, K. I., & Forster, J. C. (2003). A Windows display program with millisecond accuracy. *Behavior Research Methods Instruments & Computers, 35*, 116-124.

Gaskell, M. G. (2003). Modelling regressive and progressive effects of assimilation in speech perception. *Journal of Phonetics, 31*, 447-463.

Gaskell, M. G., Hare, M., & MarslenWilson, W. D. (1995). A connectionist model of phonological representation in speech perception. *Cognitive Science, 19*, 407-439.

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

Gaskell, M. G., & Marslen-Wilson, W. D. (1998). Mechanisms of phonological inference in speech perception. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 380-396.

Gaskell, M. G., & Marslen-Wilson, W. D. (2001). Lexical ambiguity resolution and spoken word recognition: Bridging the gap. *Journal of Memory and Language, 44*, 325-349.

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.

Gow, D. W. (2003). Feature parsing: Feature cue mapping in spoken word recognition. *Perception & Psychophysics, 65*, 575-590.

Gow, D. W., & Im, A. M. (2004). A cross-linguistic examination of assimilation context effects. *Journal of Memory and Language, 51*, 279-296.

Janse, E. (2004). Word perception in fast speech: artificially time-compressed vs. naturally produced fast speech. *Speech Communication, 42*, 155-173.

Kohler, K. J. (1990). Segmental reduction in connected speech in German: Phonological facts and phonetic explanations. In W. J. Hardcastle & A. Marchal (Eds.), *Speech production and speech modelling* (pp. 69–92). Dordrecht, The Netherlands: Kluwer.

Lahiri, A. and Reetz, H. (2002). Underspecified recognition. In C. Gussenhoven, N. Werner, & T. Rietveld (Eds.), *Labphon 7* (pp. 637-676). Berlin: Mouton de Gruyter.

Lahiri, A., & Marslen-Wilson, W. D. (1991). The mental representation of lexical form: A phonological approach to the recognition lexicon. *Cognition, 38*, 245-294.

Lee, S. (2005). The effect of assimilation contexts in word detection. *Studies in Phonetics, Phonology and Morphology, 11*, 105-124.

Lippmann, R.P. (1997) Speech recognition by machines and humans. *Speech Communication*, 22, 1-16.

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

Mann, V. A., & Repp, B. H. (1981). Influence of preceding fricative on stop consonant perception. *Journal of the Acoustical Society of America, 69*, 548-558.

Marslen-Wilson, W., Moss, H. E. & Van Halen, S. (1996). Perceptual distance and competition in lexical access. Journal of Experimental Psychology: Human Perception & Performance, 22, 1376-92.

McQueen, J.M., & Viebahn, M. (2007). Tracking recognition of spoken words by tracking looks to printed words. *Quarterly Journal of Experimental Psychology, 60*, 661-671.

Mitterer, H., & Blomert, L. (2003). Coping with phonological assimilation in speech perception: Evidence for early compensation. *Perception & Psychophysics, 65*, 956-969.

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., Csépe, V., & Blomert, L. (2006). The role of perceptual integration in the recognition of assimilated word forms. *Quarterly Journal of Experimental Psychology, 59*, 1395-1424.

Moore, R K. (2007). Spoken language processing by machine. In M. G. Gaskell (Ed.) *Oxford Handbook of Psycholinguistuics.* (pp. 723-738), Oxford: Oxford University Press.

Nolan, F. (1992). The descriptive role of segments: Evidence from assimilation. In D. R. Ladd & G. Docherty (Eds.), *Laboratory Phonology II* (pp. 261-280). Cambridge: Cambridge University Press.

Paradis, C., & Prunet, J.-F. (1991). *The Special Status of Coronals: Internal and External Evidence,* Phonetics and Phonology Vol. 2, Academic Press, San Diego.

Repp, B. H. (1983). Bidirectional contrast effects in the perception of VC-CV sequences. *Perception and Psychophysics, 33*, 147-155.

Shatzman, K.B. & McQueen, J.M. (2006). The modulation of lexical competition by segment duration. *Psychonomic Bulletin & Review*, 13, 966-971.

Snoeren, N. D., Hallé, P. A., & Seguí, J. (2006). A voice for the voiceless: Production and perception of assimilated stops in French. *Journal of Phonetics, 34*, 241-268.

Snoeren, N.D., Hallé, P.A. & Seguí, J. (2008). The perceptual processing of partially and completely assimilated words in French. *Journal of Experimental Psychology: Human Perception and Performance, 34*, 193-204.

Steriade, D. (2001). Directional asymmetries in place assimilation: A perceptual account. In E. Hume & K. Johnson (Eds.), *The role of speech perception in phonology* (pp. 219–250). New York: Academic Press.

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.

Wheeldon, L., & Waksler, R. (2004). Phonological underspecification and mapping mechanisms in the speech recognition lexicon. *Brain and Language, 90*, 401-412.

Appendix

Test materials used for the Experiments. The first, italicized sentence was used only in Experiment 2 to supply a biasing semantic context. Biasing sentences marked with an asterisk were adapted from Gaskell & Marslen-Wilson (2001); the remaining materials were unadapted. The second sentence was used in both experiments. The format of each sentence was: Pretarget phrase (coronal/noncoronal target) [viable/unviable context].

1. *The chef swiftly removed the head and tail, then checked the diners' order.* They asked for the (cod/cob) [poached/too late]
2. *After the auction, the ranch-owner was in high spirits.* The sale of the (herd/herb) [provided enough money to keep him going/netted enough money to keep him going]
3. *The ceremony was held in June and the sunny weather added to the air of celebration.* An article about the (bride/bribe) [made the local paper/turned up in the local paper]
4. *The council was worried about the effects on health of the old water pipes.* They got the (lead/leg) [covered immediately/tested immediately]
5. *Fashions are OK, but they can be dangerous.* That new (fad/fag) [causes cancer/tends to cause cancer]
6. *We woke up with a start from a deep sleep.* That terrible (thud/thug) [caught us by surprise/took us by surprise]
7. *The conditions in the outback were difficult for driving.* In the intense heat, the (mud/mug) [cracked up completely/turned to dust]
8. *Kate was a bit worried about the route she was taking.* After a few miles, the (road/rogue) [cut across the desert/turned North across the desert]
9. **The young plants were in a bad state, and Patrick's attempts to protect them didn't help.* When he sprayed it, the (bud/bug) [curled up and died/twisted slightly]
10. **They had played the third and second placed sides, but even worse was to come.* Their final game was against the (lead/league) [cricket team/tennis team]
11. *The head office of the telecommunications company was empty.* The manager was at the (phone/foam) [packaging department/distributors]
12. *We were impressed by her stylish delivery and intonation.* Jane finished off the (scene/seam) [beautifully/deftly]
13. *She was learning about planting her allotment the hard way.* Mary threw the (bean/beam) [promptly on the ground/dutifully on the ground]
14. **The hotel room was shabby, and the bedclothes were thin and old.* It was a rather (worn/warm) [blanket/duvet]
15. *It's best to start the day with a burst of activity.* I think a quick (run/rum) [picks you up/does you good]
16. **His daughter had thrown the shapes all over the room, and was playing with the cube.* Harry found the (cone/comb) [pretty quickly/down on the floor]
17. *Pete was listening to the radio on the way home.* Because he wasn't concentrating, the (turn/term) [passed him by completely/took him by surprise]
18. **The debate on bacon prices drew representatives from all over Northern Europe, but the Norwegian left early.* After the speech, the (Dane/Dame) [planned to leave/decided to leave]
19. *Luckily the group found somewhere big enough to shelter in.* At least the (barn/balm) [protected them from the wind/decreased the effect of the wind]

20. **Lee was experimenting with printing a range of abstract shapes.* The T-shirt had a (line/lime) [print on it/drawing on it]
21. *The union action was well supported.* Unfortunately, the (ban/bang) [caused havoc/destroyed the equilibrium]
22. *She would use her spare time at night to try and get fit.* Sandra (ran/rang) [carefully to avoid upsetting the neighbors/ten times last month]
23. *The staff canteen is handy for snacks.* A large (bun/bung) [costs a pound/tends to fill the gap]
24. *It would normally turn up at feeding time, rubbing her ankles and looking hungry.* Julie saw the (cat/cap) [by the front door/next to the post-box]
25. *The oven was switched on just before dinner time.* Nick started to (heat/heap) [bowls on the stove/dinner plates on the stove]
26. *He was irritated by the newly uncovered dust in the doorway.* Stephen put the (mat/map) [back where he found it/down for a moment]
27. *The children were guided to the desks in front of the blackboard.* They had to (sit/sip) [meekly while the party carried on/daintily while the party carried on]
28. *His arguments were always elegant and entertaining.* Michael used his (wit/whip) [brilliantly/discerningly]
29. *County councils have improved their safety measures for winter road conditions when bad weather is forecast.* They appear to (grit/grip) [motorways much better/newer roads much better]
30. **The weather had worsened considerably, and snow flurries had begun to fall as he drove across the country.* Four hours of intermittent (sleet/sleep) [brought out the worst in Mark/destroyed Mark's good humor]
31. **There was plenty of fruit to choose from, and it would have been easy to select the wrong ones.* They picked the (right/ripe) [berries for the pie/nectarines for the punch]
32. **It's easy to see the building where the soldiers are stationed, right at the top of the hill.* When you reach the (fort/fork) [go left/don't turn left]
33. *They were obviously in the mood to get drunk.* The customers had most of the (port/pork) [guzzled early on/demolished early on]
34. *The competition on the banks of the river was well attended.* The new fish (bait/bake) [got tested yesterday/didn't work very well]
35. **Growth had been slow for some time, but then the economy of the region picked up at the end of the year.* The council found the (late/lake) [growth surprising/display interesting]
36. **Rob sipped his pint and aimed at the treble twenty section of the board.* The (dart/dark) [gradually fell/descended slowly]
37. *The apple looked so appealing that he couldn't help himself.* Paul took a (bite/bike) [guiltily/deliberately]
38. **Philip's nephew was trying to catch shrimps in a rockpool with his new fishing pole.* The little boy's (net/neck) [got cut by the blade/nearly got cut by the blade]
39. *Ben usually went straight to the bus stop after a late shift at work.* He would (wait/wake) [grumpily at ten o'clock/daily at ten o'clock]
40. *It's eyes were barely visible beneath the layers of dirt.* The filth from the ancient house left the (rat/rack) [grimy/dirty]

Table 1. Example test sentences. The key word is underlined.

Sentence Type	Example
Word-final coronal:	
Unviable context	I think a quick run does you good
Viable context	I think a quick run picks you up
Word-final noncoronal:	
Unviable context	I think a quick rum does you good
Viable context	I think a quick rum picks you up

Note: the viability variable refers to the extent to which the segment following the key word provides a suitable context for place assimilation of a coronal form.

Table 2. Distribution of stimuli tested in Experiment 1 according to intended (underlying) target-final consonant, contextual viability and type of assimilation.

	Underlying coronal		Underlying noncoronal	
	Viable	Unviable	Viable	Unviable
Type of assimilation				
Voiced stop	87	89	83	76
Nasal	105	113	99	101
Unvoiced stop	148	142	133	141
Total	340	344	315	318

Table 3. Proportions of coronal responses and response times in Experiment 1. Figures in brackets are standard error values.

	Proportion of coronal responses		Coronal RT (ms)		Noncoronal RT (ms)	
	Viable	Unviable	Viable	Unviable	Viable	Unviable
Underlying coronal						
Voiced stop	.88 (.01)	.93 (.01)	524 (20)	535 (22)	536 (34)	579 (32)
Nasal	.66 (.01)	.86 (.01)	626 (32)	584 (27)	523 (20)	522 (20)
Unvoiced stop	.66 (.02)	.83 (.01)	542 (23)	540 (22)	555 (27)	648 (33)
Underlying noncoronal						
Voiced stop	.15 (.02)	.15 (.01)	565 (21)	595 (26)	524 (24)	531 (26)
Nasal	.23 (.02)	.22 (.02)	553 (25)	536 (22)	583 (27)	618 (29)
Unvoiced stop	.29 (.01)	.17 (.01)	575 (26)	555 (25)	540 (22)	512 (21)

Note: RT = response time.

Table 4. Expert phoneticians' ratings of selected stimuli in Experiment 1. Figures in brackets are standard error values.

	Weak assimilation	Strong assimilation	Underlying noncoronal
First pass			
Coronal response proportion	.75 (.08)	.35 (.04)	.35 (.06)
Considered Judgment			
Coronal response proportion	.94 (.02)	.45 (.04)	.47 (.09)
Confidence rating (1-9)	7.97 (.45)	5.83 (.65)	5.95 (.74)
Number of presentations	2.39 (.27)	2.83 (.33)	3.28 (.58)

Table 5. Proportions of coronal responses and response times in Experiment 2.
Figures in brackets are standard error values.
Note: RT = response time.

	Proportion of coronal responses	Coronal RT (ms)	Noncoronal RT (ms)
Semantically neutral-unspliced			
Originally viable	.42 (.02)	914 (55)	852 (66)
Originally unviable	.28 (.03)	905 (53)	763 (32)
Semantically neutral-spliced			
Originally viable	.35 (.03)	864 (76)	800 (38)
Originally unviable	.38 (.03)	826 (35)	795 (35)
Semantic bias-unspliced			
Originally viable	.62 (.02)	1066 (63)	1110 (76)
Originally unviable	.48 (.03)	1081 (62)	985 (45)
Semantic bias-spliced			
Originally viable	.55 (.03)	1115 (59)	1168 (77)
Originally unviable	.62 (.03)	1080 (54)	1083 (84)

Figure 1. Response choice data for Experiment 1. Error bars represent standard error.

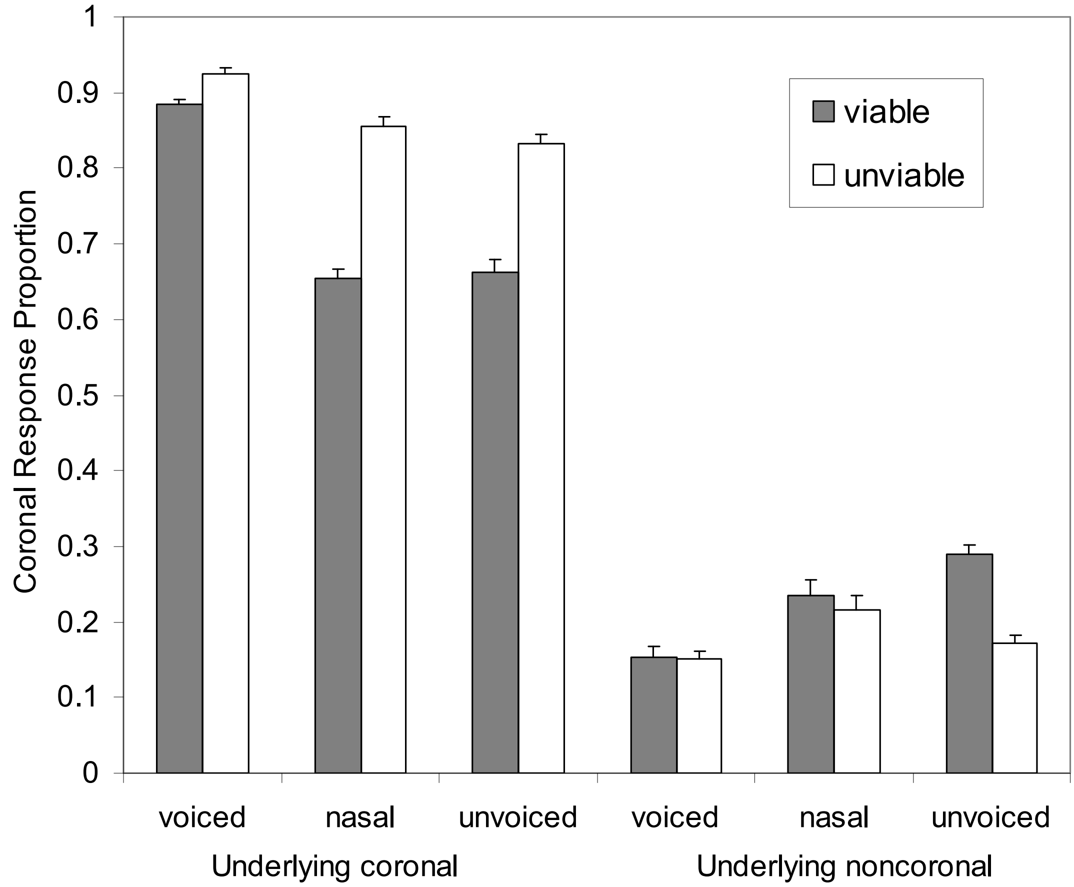

Figure 2. Distribution of Experiment 1 item response choices for combined unvoiced stop and nasal stimuli for underlying coronal targets (upper panel) and noncoronal targets (lower panel). The y-axis measures the proportion of stimuli in each condition falling within the given bands.

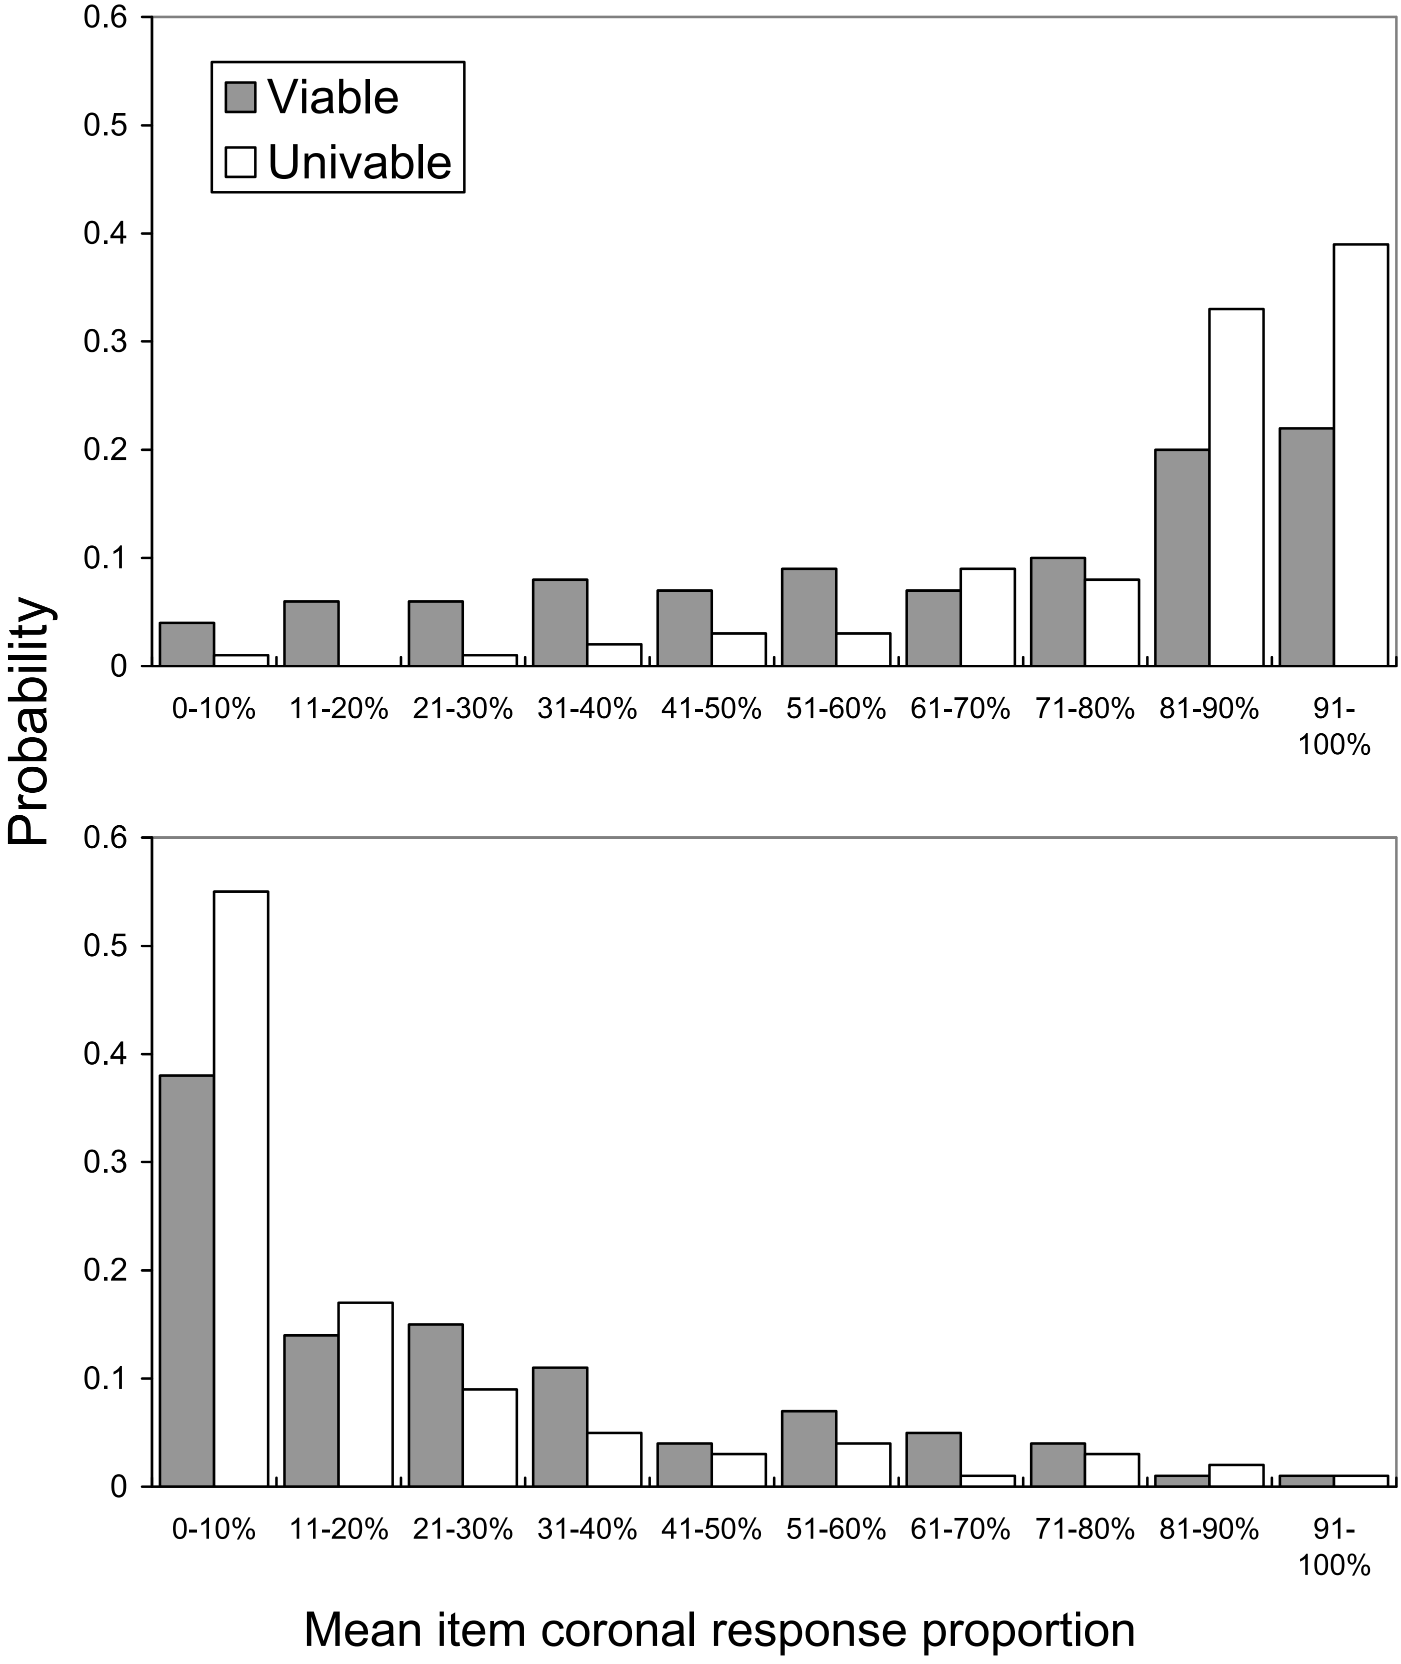

Figure 3. Response choice data for Experiment 2. Error bars represent standard error.