

Perceptual distribution of merging phonemes

VALERIE FREEMAN

*University of Washington**

1 Introduction

This study seeks to map the perceptual vowel space of front vowel phonemes undergoing merger before voiced velars in Pacific Northwest English (PNWE). In production, most speakers spectrally merge /ɛg, eg/ at a point between their non-prevelar counterparts /ɛ, e/, but the height of /æɡ/ is more variable. With variable production in the speech community, a question of perception arises: do Northwesterners maintain the same category boundaries for prevelar front vowels as non-prevelars, or are the prevelars merged in perception as they often are in production? This study addresses the question by mapping the perceptual space of front vowels in prevelar vs. precoronal contexts. Stimuli were created to synthesize an initial /b/ followed by 24 front-vowel formant value combinations (F1, F2) with no offglide or coda transitions. Twenty Northwestern subjects were told that each stimulus was the first part of a word that had been cut off in the middle, and they indicated which word they heard with a button press. In the first three blocks of randomized stimuli presentation, the word choices were in the shape /b_d/: *bad, bid, bayed, bed, bead*; the second three blocks used the same randomly-presented stimuli (unbeknownst to subjects), but the word choices were /b_g/: *bag, big, bagel, beg, beagle*. This design forces lexical access during the task, as subjects must imagine they are hearing words, not contextless phonemes.

The paper is organized as follows: Section 2 presents background information on the merger in production, followed by predictions for perception. Section 3 describes the experimental design, stimuli creation, and response procedures. Section 4 presents results, and Section 5 concludes with discussion and future work.

2 Motivation

2.1 Merger in Production

Recent sociophonetic studies of Pacific Northwest English (PNWE) (e.g., Freeman 2014; Riebold 2014; Wassink et al. 2009) have described “low-front pre-velar raising/merger,” a sound change in progress involving the raising and upgliding of the low-front vowels /æ, ɛ/ and the lowering of mid-front /eɪ/ before voiced velars /ŋ, g/. Raising is advanced and stable before the velar nasal: /æŋ, eɪŋ/ are merged at a location in F1xF2 space between non-prevelar /ɛ, eɪ/ (e.g., *length, Lang* [lɛŋθ, lɛŋ]) (Freeman 2015). Before the voiced stop, /ɛg, eɪg/ are also merged at this intermediate location, so that words like *beg, vague* rhyme [bɛɪg, vɛɪg]. However, the height of /æɡ/ is more variable across speaker groups and speaking styles,

*Portions of this work were supported by NIH R01 DC60014. Special thanks to John Riebold, Dan McCloy, Richard Wright, and members of the UW Phonetics Lab and 2013-14 phonetic perception seminars.

with higher positions seen in middle-aged speakers, males, and casual styles. In studies of Washingtonians, for example, middle-aged speakers (late 30s to early 60s) show the highest positions across three generations, with some males showing near-complete overlap with the higher prevelars, so that words like *bag* rhyme with *beg*, *vague* [bɛɪg] (Freeman 2014; Riebold 2014). With lower and more variable positions for both /æɡ/ and /ɛɡ/ found in older generations (Reed 1961; Wassink and Riebold 2013), the advancement seen in middle-aged speakers suggests a change progressing over time, perhaps toward full merger of all three vowels /æ, ɛ, ɛɪ/ before voiced velars. This progression seems to have continued for /ɛɡ, ɛɪɡ/, which remain merged at [ɛɪg] in younger adults (currently 20s to mid-30s), but /æɡ/ again shows increased variation and lower positions, suggesting that the last prevelar to join the tide may also be more sensitive to social meaning (Freeman 2014).

2.2 Predictions for Perception

A thorough investigation of phonological merger must examine both production and perception; it cannot be assumed that the two are identical. For example, in cases of near-merger (cf. Di Paolo 1988), speakers produce different variants but judge them the same, and in cases where merger is sensitive to stigma and style-shifting (cf. e.g., Labov 1994), speakers produce merged variants but judge them as different. This study aims to determine whether some version of these situations hold in PNWE, or whether perception mirrors production. Since the two parts of the PNWE merger appear to be treated differently in production, they are examined separately in perception here, beginning with the following hypotheses:

H1: Perception of /ɛɡ, ɛɪɡ/ mirrors production. As /ɛɡ, ɛɪɡ/ show (near-)complete merger in production for nearly all PNWE speakers examined in previous studies, they are expected to be merged in the same location in perception as well. In this study, stimuli with F1xF2 values matching those of the merged /ɛɡ, ɛɪɡ/ in production will be judged as belonging to either class, with no clear bias toward either option.

H2: Perception of /æɡ/ is more varied, as in production. Given the wide variation in production of /æɡ/ present in the speech community, Northwesterners are expected to accept variation in perception as well. This variation could take the form of judging a wider variety of stimuli as acceptable realizations of /æɡ/, or there could be wider variation between subjects.

3 Experimental Design

Stimuli were created (Sec. 3.1) to synthesize 24 front-vowel formant combinations following an initial /b/. Subjects were told (Sec. 3.2) they were hearing words “cut off” in the middle, and they indicated which word with a button press in two test conditions. In the /b_d/ condition, all five options were lexical items of the form /b_d/: *bad*, *bid*, *bayed*, *bed*, *bead*. In the /b_g/ condition, all options are lexical items of the form /b_g/: *bag*, *big*, *bagel*, *beg*, *beagle*. Unbeknownst to subjects, the same stimuli were played in both conditions, each repeated randomly in three blocks.

The /b_d/ condition is intended to map listeners’ percepts of unmerged front vowels in PNWE. The /b_g/ condition should map the acceptability of each F1xF2 production as a

candidate for membership in each phonemic vowel undergoing merger, or in other words, which phonemes listeners expect as possible intentions of PNWE speakers. Where there is overlap between phonemes in production, responses should show greater variation, and full merger should be indicated by equal or random assignment of stimuli to each of the merged phonemes. Because there is relatively less overlap in production of the front vowels before /d/, the /b_d/ condition should show more clearly delineated responses, while the greater overlap in production before /g/ is expected to cause greater variation, indicating competing options and lower confidence.


The task is designed to access listeners' lexicons by priming them to expect and respond with real words. Many studies have shown that listeners can distinguish vowel sounds in tasks designed to avoid lexical knowledge, but the merger in progress seems to be below the level of social awareness, meaning that speakers are generally not aware of the change (Freeman 2014). Without appealing to the lexicon, listeners may respond at a more abstract level, thinking of prototypical vowels rather than observed realizations. Including word-initial /b/ in the stimuli is intended to facilitate lexical access by making the stimuli more word-like, rather than asking subjects to imagine that the heard vowel has been extracted from a word, a situation that does not occur naturally.

3.1 Stimuli

Purely synthetic stimuli were created in order to be able to control all parameters of the signal. A male speaker of PNWE was selected to provide a model for the vowel space and F1xF2 values for the synthetic stimuli. The model speaker was a male Caucasian second-generation Seattleite in his mid-50s chosen from several recorded in a previous production study (cf. e.g., Freeman 2014; Wassink et al. 2009). The phonemes in his front vowel space show less overlap than other speakers', and their configuration is fairly linear, which simplifies a model of raising on an F1xF2 slope. For the model vowel space, shown in Figure 1, midpoint formant values of word list tokens in non-nasal, non-liquid contexts were measured in Praat (Boersma and Weenink 2013) and plotted in F1xF2 space using the R package PhonR (McCloy 2015). In addition, low-front pre-/g/ tokens (e.g., *egg*, *beg*; *bag*) were plotted separately to determine the locations for these contexts, which are raised for this speaker. Note that /εg/ tokens like *vague*, *bagel* were not available for this speaker when stimuli were created, but later analysis showed substantial overlap with /εg/ (Freeman 2014).

Ellipses of two standard deviations around the means of the model speaker's vowels were used as a guide for the areas to be represented in the synthetic stimuli. Figure 2 shows the F1xF2 values for the created stimuli (black dots) overlaid with the model speaker's ellipses. Stimuli were set at even intervals with 75 Hz between each value in F1 and 150 Hz between each in F2. These values approximate those separating stimuli in Johnson et al.'s (1993) method of adjustment task, which used step sizes of 0.37 Bark, "slightly larger than the just-noticeable-differences reported by Flanagan (1957)" (p. 57). Using Hertz rather than Bark was judged sufficient for this experiment because the Hertz-Bark relationship for the affected vowel space is roughly linear (cf. Ladefoged 1996). Some values that fit inside or very close to the model ellipses were not used (e.g., 400x1800, 475x1650); plots of the model speaker's entire vowel space showed overlap in these central areas with back or central vowels such as /ʊ/ and /ʌ/, and stimuli created with these values auditorily sounded too central, as well.

Figure 1: Model speaker's front vowel space. Word list token midpoints (F1x F2 in Hz) in non-liquid, non-nasal alveolar and pre-g contexts, with ellipses of 2 standard deviations around vowel means.

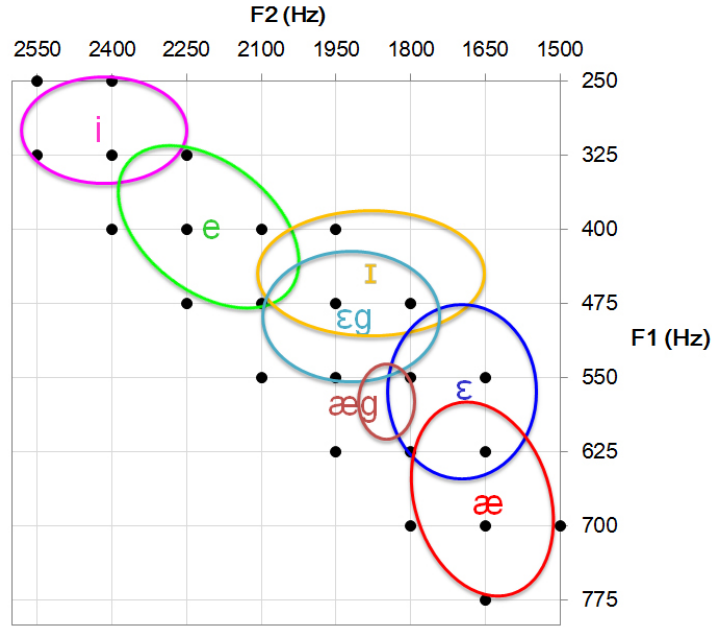


Stimuli on the front edge of the model vowel space were included, even though the model speaker showed no tokens in this range, in order to allow for the possibility that raised vowels merge in more front locations than unaffected vowels, as suggested by the fronted locations of the prevelar model tokens.

Stimuli were created at a sampling rate of 11,025 Hz in SynthWorks (Scicon R&D Inc. 2004), a Klatt-based synthesizer, on a Macintosh desktop computer. Each stimulus begins with a synthesized /b/ consisting of a 10-ms release and 40-ms voiced transition to the 120-ms steady-state vowel. (The 40-ms transition duration is in line with values reported by Walsh and Diehl (2007) for the percept of a bilabial stop release.) Thus, the total duration for each stimulus is 170 ms, which was judged to sound natural but short for words said in isolation, in line with the scenario to be given to subjects that they will hear a single word “cut off” in the middle. Duration is held constant across stimuli to avoid introducing any durational cues that may covary with underlying phonemic quality or merger application. Vowel formant values remain steady over their duration so as to avoid any cues toward gliding that could similarly bias subjects’ lexical decisions. Pitch (f_0) begins at 100 Hz and rises to 110 Hz linearly over the duration of each stimulus. This pattern matches that of the model speaker, whose pitch rose slightly as he read the target words with focus intonation in a carrier phrase. Flat and slightly falling pitch contours were also applied but rejected because they were judged to sound very unnatural. Falling contours are also undesirable because they often occur phrase-finally, potentially biasing subjects toward perceiving codas; with no formant transitions, the coda could be perceived as a glottal stop, often a component or allophone of alveolars but not velars, which could bias or confuse lexical decisions.

F3 was calculated as in Johnson et al. (1993), following the formula given for front vowels

Figure 2: Stimulus grid. Black dots indicate F1xF2 values (Hz) of synthesized stimuli; ellipses follow the distributions of the model speaker’s vowels, as in Figure 1.



in Nearey (1989:2095), with the output rounded to the nearest 10 Hz:

$$(1) \quad F3 \text{ (Hz)} = 0.522 \cdot F1 + 1.197 \cdot F2 + 57$$

Bandwidths for the first three formants (B1, B2, B3) were calculated following the formulas used by Johnson et al. (1993) to approximate the model values given in Klatt (1980), with the outputs rounded to the nearest 5 Hz:

$$(2) \quad B1 \text{ (Hz)} = 29.27 + 0.061 \cdot F1 - 0.027 \cdot F2 + 0.02 \cdot F3$$

$$(3) \quad B2 \text{ (Hz)} = -120.22 - 0.116 \cdot F1 + 0.107 \cdot F3$$

$$(4) \quad B3 \text{ (Hz)} = -432.1 + 0.053 \cdot F1 + 0.142 \cdot F2 + 0.151 \cdot F3$$

F4 and F5 were fixed at 3500 and 3700, respectively, both with bandwidths of 200 Hz, over the entire duration of each stimulus. For the first 10-ms frame corresponding to the /b/ release, values were set as recommended for /b/ in Klatt (1980): F1, F2, F3 at 200, 1100, 2150 Hz, respectively, with respective bandwidths of 60, 110, 130 Hz. Amplitude of voicing (AV) was fixed at 60 dB, as described by Klatt (1980) as typical for a full vowel, except in the first 10-ms frame, the /b/ release, which was set at 20 dB to create a release-burst percept. For the same purpose, the amplitudes of aspiration (AH) and frication (AF) were both set at 60 dB in the first frame, with AH fixed at 20 dB and AF at 0 dB thereafter. Amplitude of the bypass (AB) was set to 63 dB in the first frame, following the suggested values for /b/ in Klatt (1980), and 0 dB thereafter.

Formant transitions for the initial /b/ were created as follows. Just after release (in the second 10-ms frame), F1, F2, F3 were fixed at the lowest values used for each formant in the

target vowels, 250, 1500, and 2150, respectively. Formant values in the next 10-ms frame were calculated as 80% of the distance in Hz to the formant values of the following steady-state vowel. The remaining two frames were interpolated linearly. Bandwidths for the first three formants remained at 60, 110, and 130, respectively, in the frame after the /b/ release, and were then interpolated linearly to the values calculated at the onset of the steady-state vowel. This method successfully gave the auditory percept of a /b/, whose formant values begin low and rise quickly after release before gradually reaching those of a following vowel. Values following the release were similar to those reported in Kewley-Port (1982) and those of the model speaker.

The open quotient (OQ) and spectral tilt (TL) were manipulated toward a slightly breathy voice quality to improve the overall auditory naturalness of the stimuli, which otherwise sounded rather robotic. Sample values were found by trial-and-error, relying on the experimenter’s auditory percepts to increase the naturalness of the voice quality and to reduce variation between stimuli in auditory percepts of pitch. This resulted in a pattern that aligned with the diagonally-sloping shape of the sample vowel space; the pattern was then regularized to yield the following values. From the frame following the /b/ release, the open quotient was set to 55% for all stimuli except the three on the right edge of the F1xF2 diagonal slope seen in Figure 2 (i.e., 400x1950, 475x1800, 550x1650), which remained at an OQ of 50%. Spectral tilt remained at 0 dB for this edge and stimuli on the next diagonal line (that connecting 250x2400 to 700x1500). TL for the next diagonal (connecting 250x2550 to 700x1650) was set at 10 dB, and TL for the left-most edge was set to 15 dB.

All other parameters remained at the defaults set in SynthWorks.

Finally, the files created in SynthWorks were exported as WAV files and imported into Praat, where any residual pops were removed from the end of the vowel by setting the waveform to zero after the last zero-crossing within the vowel waveform’s final periodic cycle. The resulting stimuli sound reasonably natural over circumaural headphones, although some apparent pitch differences remain.

3.2 Response Procedures

Subjects were 20 native PNWE speakers with normal hearing who grew up in Washington, Oregon, or Idaho. Table 1 shows the distributions of subjects by gender and age group. Although reasonably balanced, there are relatively fewer middle-aged subjects and a slight over-representation of older females.

Table 1: Distribution of subjects by age and gender.

Ages	Female	Male	Total
18 – 29	4	3	7
30 – 59	2	3	5
60 – 75	6	2	8
<i>Totals</i>	12	8	20

After a brief demographic questionnaire and hearing screening, subjects were seated in a sound-attenuated booth at a computer screen with circumaural headphones and an ioLab

Figure 3: Button box response options, /b_d/ condition.




Figure 4: Button box response options, /b_g/ condition.



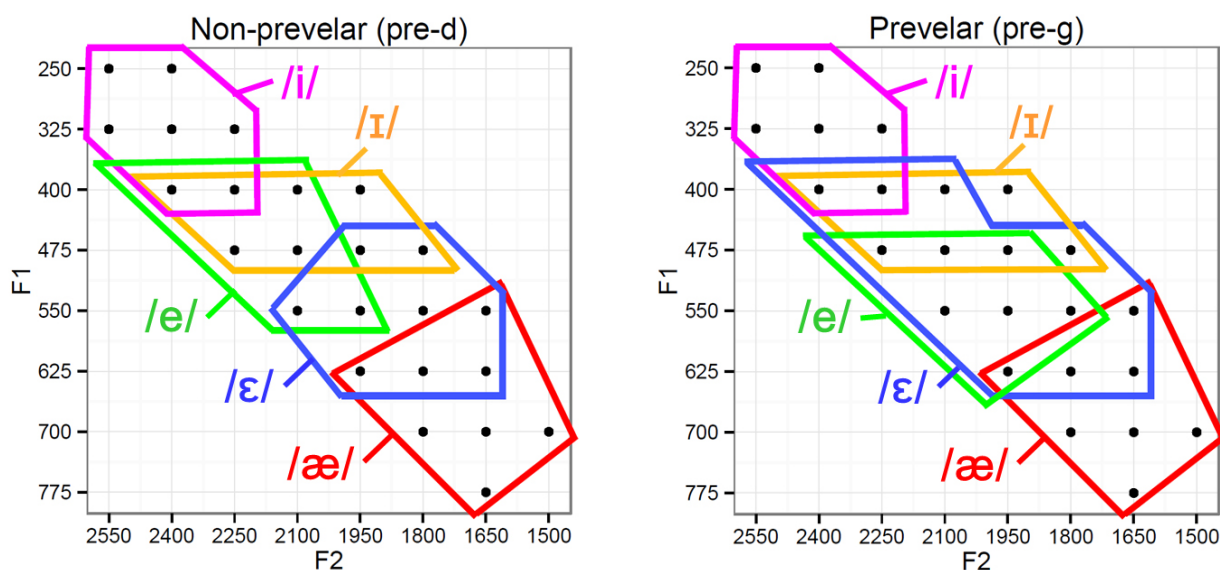
Systems button response box, initially labeled with /b_d/ response options *bad*, *bid*, *bayed*, *bed*, *bead*, as pictured in Figure 3. PsychoPy (Peirce 2014) was used to present stimuli and instructions and collect responses. The experimenter was present in the booth during a short training phase used to familiarize subjects with the stimuli and response procedures. The following instructions appeared on the screen and were read aloud by the experimenter:

You will hear a computerized voice saying words that have been cut off in the middle. After each one, press the button below the word you heard.
We'll start with a few to practice. First get familiar with the word choices.
When you're ready to start, press any lighted button. (The screen will be gray while you listen.)

The experimenter then asked subjects to read the response options on the button box aloud to ensure they were familiar words and to become comfortable with their locations. Once begun, the training phase consisted of three of the stimuli played in random order through the headphones, each beginning 600 ms after the previous button response. These stimuli were chosen for their ease of discriminability: the highest/most front vowel (250x2550 Hz), exemplary of /i/, the lowest (775x1650 Hz), exemplary of /æ/, and an intermediate node judged by the experimenter to sound exemplary of /ε/ (550x1950 Hz). For all subjects, this was sufficient training to become comfortable with response procedures. The experimenter reminded subjects to respond “as fast as possible while still being accurate” and to choose the “first word that came to mind” when unsure.

The testing phase began after the experimenter left the booth, following the same presentation procedures as in the training phase. Three stimulus blocks were presented, each including all 24 stimuli in independently-generated random orders, for a total of 72 presentations. Instructions then appeared on the screen asking subjects to pause, and the experimenter returned to the booth to exchange the response option labels on the button box for /b_g/ words *bag*, *big*, *bagel*, *beg*, *beagle*, as shown in Figure 4. The experimenter explained that the instructions for the next set were the same, “but now you have to choose which of *these* words you heard,” and again asked subjects to read the options aloud. The experimenter then left the booth, and the second testing condition proceeded exactly as the first, with all 24 stimuli presented in three randomized blocks. Crucially, subjects were not told that the stimuli were the same in both conditions, and when asked open-ended questions

Figure 5: Responses, /b_d/ condition (left), /b_g/ condition (right). Outlines mark stimuli with at least 20% of responses selecting words with the indicated vowels.



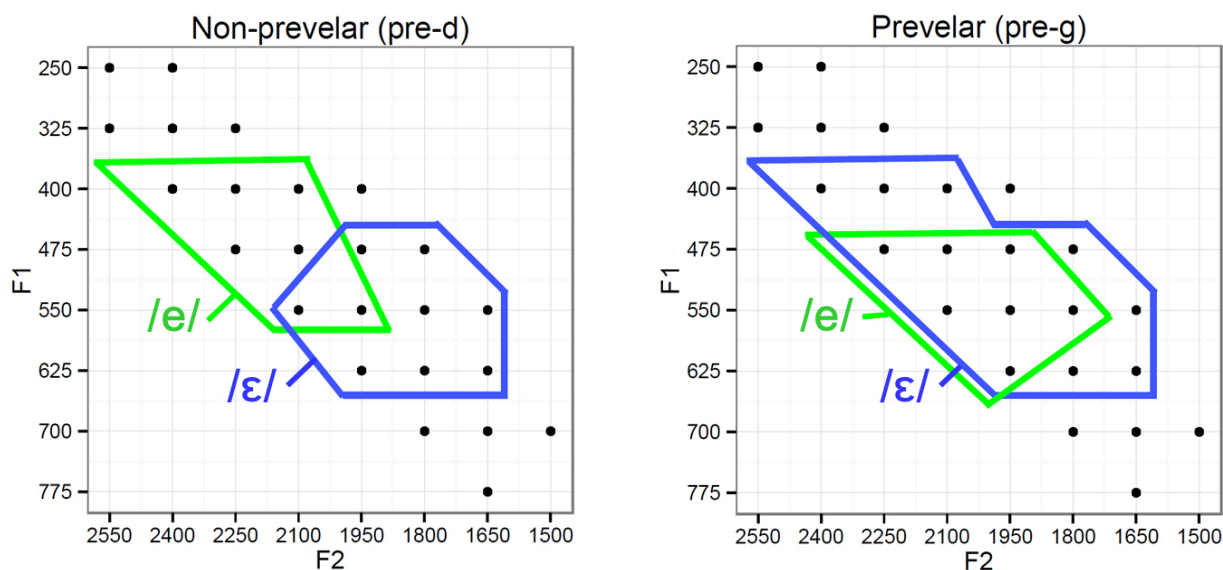
about the experience, only one asked if they were. (Others' responses indicated they did not suspect this, e.g., by saying they heard many/not many of a certain /b_g/ word.)

4 Results

Responses in the /b_d/ control condition mirror production in PNWE, as expected. Figure 5 (left panel) shows the distributions of responses of all subjects pooled together; outlines mark stimuli (black dots) with at least 20% of responses given as the word representing the indicated vowel. (So, for example, at least 20% of responses to the stimuli with an F1 of 625 Hz were “bad” and at least 20% were “bed.”)

In the /b_g/ test condition (Figure 5 right panel), the high vowels /i, ɪ/ show the same responses as in the /b_d/ condition, also as expected, since these vowels are unaffected before /g/ in production. As predicted, the mid-vowel responses differ between conditions. Figure 6 highlights these responses, also shown in Figure 5. The distribution of /ɛg/ responses (right, blue) is expanded upward and forward in F1xF2 space to include the entire distribution of /ed/ responses (left, green) as well as all /ɛd/ responses (left, blue). The distribution of /eg/ responses (right, green) is shifted downward and backward from that of /ed/ to fall entirely within the distribution of /ɛg/. Unexpectedly, the distribution of /æg/ (red, Figure 5) does not differ substantially from /æd/ with all subjects pooled; however, as individuals, about a third of subjects show an upward expansion of /æg/ compared to /æd/.

Figure 6: Mid-vowel responses, /b_d/ condition (left), /b_g/ condition (right). Outlines mark stimuli with at least 20% of responses selecting words with the indicated vowels.



5 Discussion

Both hypotheses were partially supported. For H1, responses for / ϵ g/, ϵ g/ do overlap substantially, with / ϵ g/ shifted downward, but responses for / ϵ g/ expanded upward rather than shifting. Although this pattern was not precisely predicted, it is consistent with merger and collapse of the two small word classes. That is, there are only a handful of / ϵ g/-class words (e.g., *vague*, *plague*, *bagel*, *pagan*, *vagrant*, *fragrant*, *flagrant*), some of which are uncommon, and none of them have minimal pairs with words in either / ϵ g/, ϵ g/ class. While the / ϵ g/-class is also small, it is larger and its members more common than / ϵ g/, making it a better candidate to represent a collapsed / ϵ g/, ϵ g/ class. In other words, when subjects were forced to choose between two members of a single category, they more often chose the more frequent member. However, these speculations should be confirmed with future work involving more members of each word class with word/phoneme frequency and familiarity considered.

Contrary to H2, results for all subjects pooled showed no shift in location or area for / ϵ g/ compared to / ϵ d/. However, there was increased variation in responses between subjects, with about a third showing an upward expansion of / ϵ g/ responses. Thus, it may be that some subjects responded as predicted, accepting wide variation for / ϵ g/. It is also possible that subjects responded as they themselves would pronounce the word choices. In anticipation of this hypothesis, subjects were recorded reading a word list which includes more than one word in each class; the relationship between subjects' own productions and judgments will be examined in a follow-up study.

In short, the reduction in perceptual distinctions among / ϵ g/, ϵ g, ϵ g/ further supports the characterization of these prevelar vowels as spectrally merged or merging in PNWE, and the variation between subjects is consistent with variable production in the speech community.

The study design encouraged lexical access (rather than abstract phonemic representations as in many standard phonological perception designs) by creating “partial-word” stimuli and telling subjects they were hearing pieces of words.

In future work, natural stimuli will be used first to determine whether listeners distinguish prevelar words without external contextual cues, and then to examine the acoustic features in production that predict listener classifications. In other designs, various aspects of synthetic or edited natural stimuli will be manipulated: vowel duration to investigate whether the shorter duration of /ɛg/ reported in production (Freeman 2014) distinguishes it from the other prevelars, formant slopes to examine the contributions of upglides, and pitch and voice quality in simulation of different talkers, emotions, conversational contexts, etc. Reaction times for subjects’ responses were also collected in this study and may be examined in future work as measures of confidence. In other designs, a sorting task will be used to allow subjects to repeat and compare stimuli before judging their class membership.

6 References

- BOERSMA, PAUL and DAVID WEENINK. 2013. Praat: Doing phonetics by computer. <http://www.praat.org/>.
- DI PAOLO, MARIANNA. 1988. Pronunciation and categorization in sound change. In Kathy Ferrara, Becky Brown, Keith Walters and John Baugh, eds., *Linguistic Change and Contact: NWAV-XVI*. pp. 84–92. Austin: Department of Linguistics, University of Texas.
- FLANAGAN, JAMES. 1957. Estimates of the maximum precision necessary in quantizing certain dimensions of vowel sounds. *Journal of the Acoustical Society of America* 29:533–534.
- FREEMAN, VALERIE. 2014. Bag, beg, bagel: Prevelar raising and merger in Pacific Northwest English. *University of Washington Working Papers in Linguistics* 32.
- FREEMAN, VALERIE. 2015. The prevelar vowel system in Seattle. Poster presented at the American Dialect Society (ADS) Annual Meeting, Portland, OR, Jan. 8–11.
- JOHNSON, KEITH, EDWARD FLEMMING and RICHARD WRIGHT. 1993. The hyperspace effect: Phonetic targets are hyperarticulated. In Kenneth de Jong and Joyce McDonough, eds., *UCLA Working Papers in Phonetics*. volume 83. pp. 55–73.
- KEWLEY-PORT, DIANE. 1982. Measurement of formant transitions in naturally produced stop consonant-vowel syllables. *Journal of the Acoustical Society of America* 72(2):379–389.
- KLATT, DENNIS H. 1980. Software for a Cascade/Parallel Formant Synthesizer. *Journal of the Acoustical Society of America* 67:971–995.
- LABOV, WILLIAM. 1994. *Principles of Linguistic Change, Internal Factors*. volume 1. Malden, MA: Wiley-Blackwell.

- LADEFOGED, PETER. 1996. *Elements of acoustic phonetics*. The University of Chicago Press. second edition.
- MCCLOY, DANIEL R. 2015. *phonR: tools for phoneticians and phonologists*. R package version 1.0-1.
- NEAREY, TERRANCE M. 1989. Static, dynamic and relational properties in vowel perception. *Journal of the Acoustical Society of America* 85:2088–2113.
- PEIRCE, JONATHAN. 2014. PsychoPy (version 1.80.03). <http://www.psychopy.org/>.
- REED, CARROLL E. 1961. The pronunciation of English in the Pacific Northwest. *Language* 37(4):559–564.
- RIEBOLD, JOHN M. 2014. The Ethnic Distribution of a Regional Change: /æɡ, ɛɡ, eɡ/ in Washington State. Paper presented at New Ways of Analyzing Variation (NWAV 43), Chicago.
- SCICON R&D INC.. 2004. SynthWorks (Version 8.5B for OSX). <http://www.sciconrd.com/synthworks.aspx>.
- WALSH, MARGARET A and RANDY L DIEHL. 2007. Formant transition duration and amplitude rise time as cues to the stop/glide distinction. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology* 43(3):603–620.
- WASSINK, ALICIA BECKFORD and JOHN M RIEBOLD. 2013. Individual variation and linguistic innovation in the American Pacific Northwest. Paper presented at the Chicago Linguistic Society (CLS 49) Workshop on Sound Change Actuation.
- WASSINK, ALICIA BECKFORD, ROBERT SQUIZZERO, RACHEL SCHIRRA and JEFF CONN. 2009. Effects of Style and Gender on Fronting and Raising of /æ/, /e:/ and /ɛ/ before /g/ in Seattle English. Paper presented at New Ways of Analyzing Variation (NWAV 38), Ottawa.