Sounds Shifty: Gender and Age Differences in Perceptual Categorization During a Phonetic Change in Progress

Paul De Decker

1 Introduction

Sociolinguistic descriptions of sound change offer a rich testing ground for models of change. As is often the case, though, models are constructed well after the changes they seek to explain have been brought to completion. This paper departs from such lines of explanation and tests the adequacy of a model of phonetic change on a vocalic shift in progress.

Under direct investigation is a change in the phonetic quality of the low front vowel /æ/, underway in many English speaking areas of Canada (Clarke, Elms, and Youssef 1995, Boberg 2005, De Decker 2002, Hagiwara 2005, Labov, Ash, and Boberg 2006). As part of the larger Canadian Shift (CS), the vowel in words like *sack* retracts to a more centralized position in the vowel space, approaching the area traditionally occupied by the vowel in words like *sock*. The view taken here is that an internally-induced phonetic change such as /æ/-retraction occurs when those leading the change have altered their phonological representations to match that found in production. In other words, younger speakers who are socially located at the leading edge of the change will have a mental definition, identifiable as a range of tolerance that is more centralized than the ranges for older Canadian speakers or speakers of other dialects not involved in the CS.

A strong demonstration of the effect production differences have on perceptual norms is shown by Willis (1972) who found significant cross-dialectal differences in the categorization of acoustic stimuli by listeners in Fort Erie, Ontario and Buffalo, New York. Those from Buffalo, where the pronunciation of /æ/ is described as raised, or tensed (Labov, Yaeger, and Steiner 1972), categorized a continuum of synthetic /ɛ/~/æ/ stimuli as /æ/ higher in the continuum than subjects from Fort Erie, where in Ontario /æ/ is lax, lowered and retracted (Labov, Ash, and Boberg 2006). In addition, Buffalo-based subjects, who produced more fronted variants of /a/ (compared to speakers in Ontario), identified perceptual boundaries between /æ/ and /a/ further front in the vowel space whereas subjects from Ontario placed the boundary further back.

A cross-dialect comparison may be an obvious choice for demonstrating differences in perceptual norms. However, the results are strikingly similar to those found by Janson (1986), who looked at a change in progress within a single dialect of Swedish in which the long vowel /a:/ was considered to be retracting, sounding more like the long vowel /o:/ than the traditional /a:/. A younger group of listeners consisting of 15 Stockholmers between the ages of 13 and 18 perceived a larger range of the continuum as variants of /a:/. In contrast, an older group of subjects aged 33 to 70 perceived the same stimuli as containing more variants of /o:/. These findings suggest that age-based differences in classification exist within a single community.

This paper extends these findings into an account of how phonetic change might proceed below the level of social awareness (Labov 1994). It stipulates that when listeners of successive generations re-analyze a newer range of phonetic variation found in innovative forms into their own perceptual norm, younger generations will tolerate a range of variation that is gradiently larger from that of the previous generation. This largely unmonitored tendency in turn promotes phonetic shifts and drifts such as the Swedish vowel change and, in the present case, /æ/-retraction, to proceed over successive generations without listeners being consciously aware of it. A perceptual experiment was designed to test this hypothesis.

2 Methodology

A sample of subjects (N=39) listened to and categorized word stimuli in which the second formant frequency (F2) of the vowel was resynthesized to approximate a continuous range of forms between the low front vowel /æ/ (found in the word *sack*) and the low back vowel /a/ (found in the word *sock*). The voice of a 55 year-old female from Tillsonburg, Ontario was used for the resynthesis procedure. This informant was recorded while reading a series of words containing the inventory of monophthongs found in Canadian English using a Roland Edirol R-09 recorder with a lavaliere microphone attached to the informant's shirt collar, approximately 10 inches from her mouth. The recording took place in a quiet room in her home. The informant's production of the word *sack* was then resynthesized using a script by Paul Warren. Nineteen stimuli forms were produced ranging in F2 between 2006 Hz and 1259 Hz. F1 was held constant at around 1000 Hz. These stimuli were into web pages and subjects accessed the experiment pages from their homes via the internet. All stimuli were presented in a randomized sequence and the subject could move through the tasks at his/her own pace. They were given the task of clicking on a button to hear a stimulus and then selecting from a drop-down menu to the right of the sound button one of three options they thought the stimulus form sounded most like: 1. sounds like *sack,* 2. sounds like *sock,* 3. could be either *sack* or *sock.* Each of the 19 different stimuli were presented three times.

Monolingual, English-speaking Ontarians who have lived their entire lives in the province were recruited for this experiment. Recruitment was restricted to individuals who grew up speaking English in Canada, beginning at the level of kindergarten onwards. It was crucial that they spoke no languages other than English in order to ensure that any factors influencing their perception were not due to interference from another language or dialect. All subjects were between the ages of 15 and 68 and recruited using a friend of a friend method starting with a contact known to the investigator. An age-span of roughly 50 years enabled an apparent-time analysis of perceptual categorization and a test of the hypotheses concerning the progress of phonetic change over generations of listeners within one speech community.

3 Analysis: Rate of Dominant Category Assignment for Each Stimulus

Rates of /æ/-assignment were calculated, and a multivariate analysis was run to determine the overall likelihood that listeners selected the word *sack* (henceforth /æ/) as their category choice. The category that occurred most frequently for any given stimulus (and supported by a probability of greater than .50) was taken as the dominant category. The rates of dominant category assignment for each stimulus were further examined to determine the relative effects of Gender (Male vs. Female) and Age Group on category assignment.

4 Hypotheses

It is expected that all listeners, regardless of age or gender, will assign /æ/ at higher rates to those stimuli with the highest second formant frequencies, whereas stimuli with the lowest second formant frequencies will be categorized as /æ/ at increasingly lower rates. This is consistent with the acoustic properties of production, where variants of /æ/ are characterized as having the highest F2 (and F1) frequencies (Ladefoged 2005). In contrast, variants of /a/ are produced with low F2 and high F1 (ibid.). However, because the progress of /æ/-retraction is transmitted socially, differences due to the listener's age and gender are also expected. These differences in categorization will be reflected in the rates at which stimuli are labeled /æ/ *within* the category range and the distance to which /æ/ is assigned throughout the entire continuum.

The hypothetical data presented in Figure 1 show the rate at which each stimulus was categorized as /æ/ by gender of the listener, and then overall. /æ/ appears as the dominant category for stimuli #1 through #5 since they were categorized as /æ/ more often than /a/. This means that

the *category range* for /æ/ is #1 to #5. Using this as a guide, two hypotheses are presented.

4.1 Hypothesis 1: Concerning Rates of Assignment within the Category Range

It is predicted that female listeners will be less likely to categorize /æ/ at the onset of the category range given that /æ/ is more centralized in the speech of females. Therefore, females are more likely assign /æ/ at the offset, as shown. A similar expectation is held for Age Group. Likewise, younger listeners are less likely to assign /æ/ to stimuli at the onset of the category range but more likely to exhibit higher rates towards the offset, relative to older listeners.

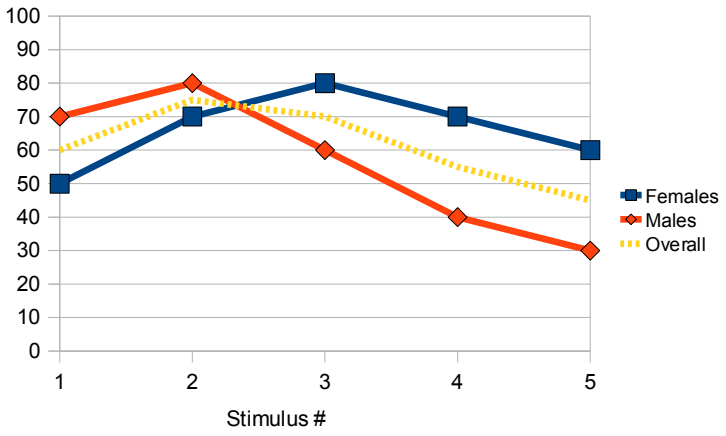

Figure 1: Rate at which each stimulus was categorized as /æ/ (hypothetical data).

4.2 Hypothesis 2: Concerning Rates of Assignment Beyond the Category Range

As a consequence of categorizing /æ/ at higher rates towards the offset of the category range, females and younger age listeners are predicted to show higher rates of dominant category assignment beyond the *overall* category range compared to males and older listeners, respectively. This is also illustrated in Figure 1. We see that although assignment of /æ/ as the dominant category ended at stimulus #4, females are shown to extend /æ/ as the dominant category to a stimulus further along the continuum, in this example stimulus #5.

Both Hypothesis 1 and Hypothesis 2 predict a younger, female-led change in category assignment where the higher F2 stimuli are favored at lower rates and lower F2 stimuli at higher rates. As in the domain of production, males and older listeners will lag behind. It should also be noted that both hypotheses assume a relative *shift* in category ranges (i.e., movement of the left and right edge together), though a second option is possible: categories may show no evidence of shifting but rather an *extension* of the right edge boundary. The results presented below will enable us to determine which option is in progress for this sample of listeners.

5 Results

Subjects were divided into the 3 age groups shown in Table 1: teenagers and those in their twenties (15–29), respondents in their thirties and forties (30–49), and respondents aged 50 to 68 (50+). Table 1 also identifies the number of tokens categorized by each age group and gender for each stimulus number (i.e., stimuli #1, #2, etc.). In total 4,446 tokens (234 responses per stimulus x 19 stimuli) are included in this analysis.

5.1 Overall Patterns

Figure 2 below shows that dominant category assignment follows a monotonic pattern of decline from the start to the end of the continuum so that the rate of /æ/ responses decreases as F2 decreases. The category range for /æ/ was found to span from stimulus #1 to #7 in which /æ/ was assigned as the dominant category (i.e. more often than /a/). Figure 3 shows the effect of the listener's gender on differences in categorization and Figure 4 illustrates the effect of the age group each listener belonged to.

	Male		Female	
Age Group	Subjects	# of Tokens	Subjects	# of Tokens
15–29	5	30	8	48
30–49	7	42	7	42
50+	5	30	7	42
Total	17	102	22	132

Table 1: Subjects and tokens per stimulus by Gender and Age Group.

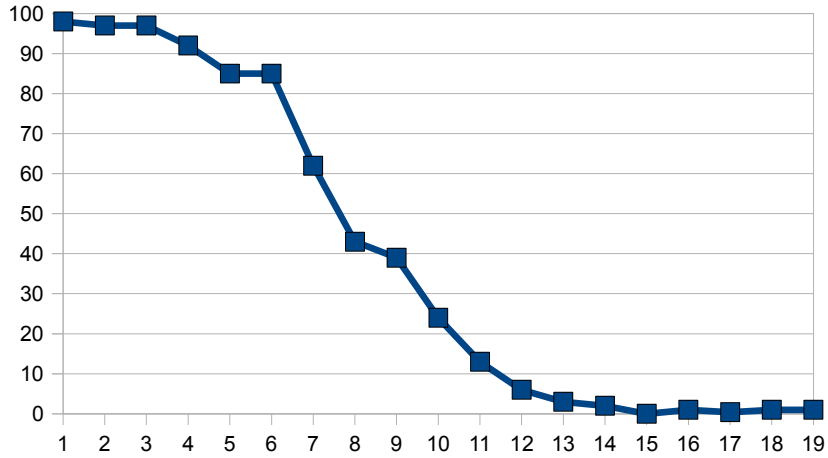

Figure 2: /æ/-categorization across the central continuum.

5.2 Effect of Gender on Rates Within and Beyond the Category Range

Figure 3 shows that both male and female listeners demonstrate roughly equal rates of categorization for stimuli #1–#4 and again from #14–#19, while significant differences were found for stimuli #7, #8, #10, and #11 in which gender significantly influences rates of categorization.

While females show appreciably higher rates of within-category assignment (i.e., #1–#7), Gender is significant at only stimulus #7. This is shown in Table 2, in which a bolded probability value indicates a factor weight that favors /æ/-categorization.

The hypothesis that females would be more likely to assign /æ/ at higher rates towards the end of the category range is supported. However, no difference was found in assignment at the onset of the category range. Both males and females are equally likely to categorize these stimuli as /æ/.

The second hypothesis is strongly consistent with the results. Table 3 shows significant effects due to the listener's gender in which females continue to favor the categorization of stimuli as /æ/ past the overall category boundary. In fact, females were found to exhibit favoring probabilities and higher rates of /æ/-categorization for stimuli #8, #10, and #11, compared to males who disfavored /æ/-categorization.

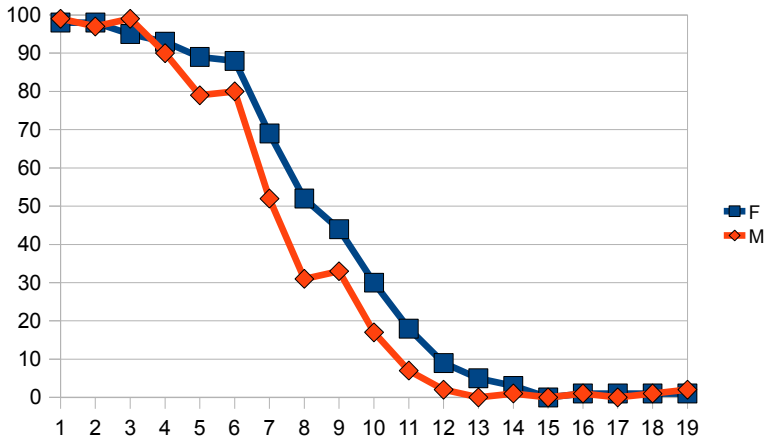

Figure 3: /æ/-categorization by Gender.

	# of tokens	(%)	Probability
Overall	234	***62%***	***.64***
F	132	69%	**.57**
M	102	52%	.40
			Significance = .030

Table 2: Rate and probability of /æ/-categorization by Gender for Stimulus #7.

A pivotal point in the continuum in which males and females show the greatest divergence is at stimulus #8 (females = 52%, males = 31%). At this point, females assigned /æ/ as the dominant category more than half the time whereas males had already switched over to /a/ as the dominant category, assigning /æ/ at a much lower rate. This suggests that females have a mental definition of /æ/ for which the right edge is more centralized than that of the male subjects.

	# of tokens	#8	#10	#11
Overall	234	43% (.43)	24% (.24)	13% (.12)
F	132	**52% (.59)**	**30% (.58)**	**18% (.62)**
M	102	31% (.38)	17% (.39)	7% (.35)
Significance		.002	.02	.01

Table 3: Rates and probabilities of /æ/-categorization for stimuli #8, #10, and #11 by Gender.

5.3 Effect of Age Group on Rates Within and Beyond the Category Range

Figure 4 illustrates age-related patterns of categorization. Listeners were broadly similar throughout the entire continuum. However, Age Group was found to significantly influence rates of assignment within the category range at stimuli #4 , #6, and #7. At these points, the youngest group was found to favor /æ/-categorization the most while the older groups were found to disfavor it.

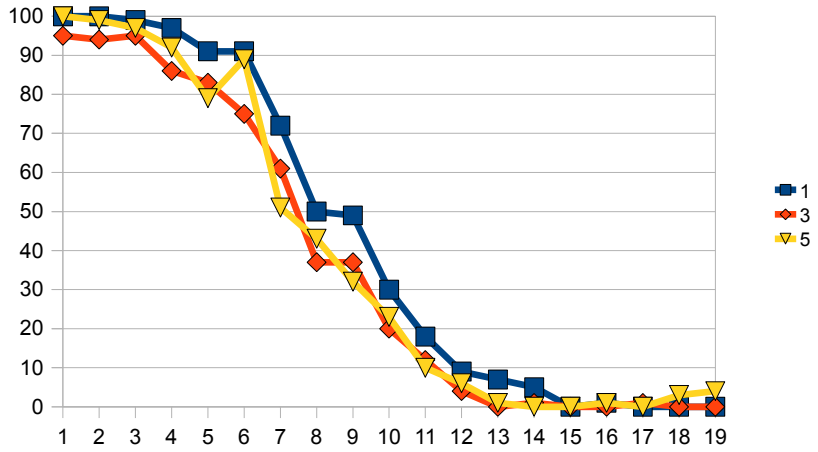

Figure 4: /æ/-categorization by Age Group.

	# of tokens	#4	#6	#7
Overall	234	**92% (.95)**	**85% (.87)**	**62% (.64)**
1	78	**97% (.74)**	**91% (.63)**	**72% (.62)**
3	84	86% (.32)	75% (.33)	61% (.49)
5	72	92% (.44)	**89% (.57)**	51% (.38)
Significance		.035	.028	.027

Table 4: Rates and probabilities of /æ/-categorization for stimuli #4, #6, and #7 by Age Group.

An interesting trend is reported for stimuli #4 and #6 where /æ/ was assigned at higher rates by the oldest age group and lower rates by the middle age group (see Table 4). In both cases, however, the rates and probability weights are lower than for the youngest age group. The favoring effect shown by those age 50+ for #6 places listeners in this age group with the youngest, an unexpected finding. Nonetheless, the general trend is consistent with the first hypothesis as successively lower rates of /æ/-assignment were found towards the offset and the youngest listeners were more likely to categorize these stimuli as /æ/.

No age groups assigned /æ/ as the dominant category past the overall category range, though the youngest group did assign /æ/ to stimulus #8 at a rate of 50%. Despite this, the /æ/-/a/ boundary seems not to be affected by age. All 3 groups assigned /æ/ as the dominant category up to stimulus #7. These data do not show a difference in apparent time consistent with a view that a re-analysis of the category boundary is underway. However, given the unexpected rates for #4 and

#6, it is worth considering how males and females performed within each age group. Figures 5 through 7 below illustrate patterns of /æ/-assignment for each age group, broken down by gender.

5.4 Gender Differences within Each Age Group

5.4.1 Youngest Age Group

In the youngest group, those aged 15-29 (shown in Figure 5), males and females seem to agree on category assignment for stimuli #1 to #6 and then again from #14 to #19. That is, #7 through #13 show some differences according to gender. Table 5 below shows significant effects due to gender at #7 with females favoring /æ/-categorization over males. Therefore, the within-category hypothesis is supported for this age group: females exhibited higher rates towards the end of the category range (up to stimuli #7). Significant effects were also found for stimuli #8 through #10. These results are also consistent with the second hypothesis as they show that females are more likely to assign /æ/ beyond the category range than males of the same age.

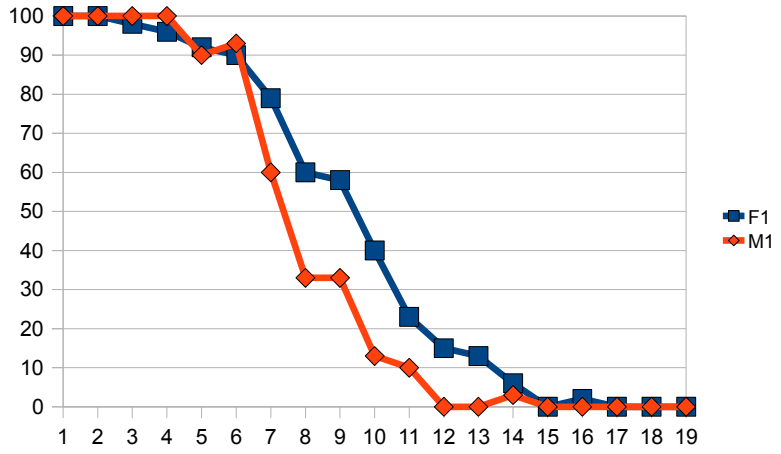

Figure 5: /æ/-categorization by Gender in the youngest age group.

	# of tokens	#7	#8	#9	#10
Overall	56	**72% (.77)**	50% (.50)	49% (.50)	30% (.27)
F1	48	**79% (.61)**	**60% (.61)**	**58% (.60)**	**40% (.64)**
M1	30	60% (.33)	33% (.34)	33% (.35)	13% (.29)
Significance		.047	.019	.034	.010

Table 5: Rates and probabilities of /æ/-categorization for stimuli #7 to #10 among the 15-29 year olds.

5.4.2 Middle Age Group

In the middle age group, those listeners aged 30–49, males and females are broadly consistent in their rates of /æ/-categorization for stimuli #1 to #3 and #13 to #19. At the beginning of the continuum where stimuli are marked by higher F2 values (show in Figure 6), males rate slightly higher than females, though Goldvarb results do not indicate a significant effect due to gender.

That is, males were not found to favor /æ/-categorization over females for these stimuli. Likewise, no statistically significant effects were found between males and females for stimuli #11 to #19.

Table 6 reveals significant findings for stimuli #5 and #7. Here, females were more likely than males to categorize these items as /æ/. This is consistent with the within-category hypothesis: females exhibited higher rates of /æ/-assignment towards the lower end of the category.

Concerning the second hypothesis, females appear to significantly favor /æ/ as the dominant category at stimulus #8 at a rate of 55% whereas males were found to disfavor at a rate of 19%. In fact, females continued to favor /æ/ into the continuum at #11. This suggests a more centralized boundary for /æ/ for female listeners at this age. In contrast, the last stimulus at which males heard /æ/ as the dominant category is at #5.

Therefore, in this age group males fall to the left of the overall category boundary and females to the right. This pattern is consistent with the beyond-category hypothesis. To compare with the youngest age group, the female boundary is 1 stimulus less centralized for these females. The male boundary, though, is 3 stimuli less centralized than that of the youngest males. It would seem, then, that gender differences are more robust in this age group and that as age increases, the boundary is less centralized.

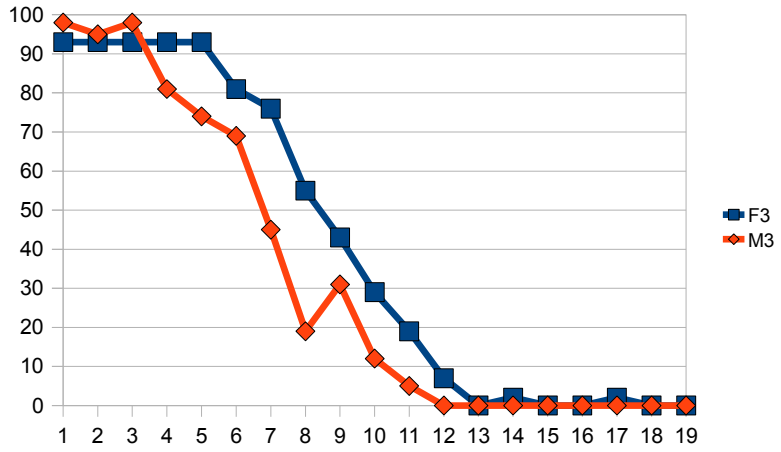

Figure 6: /æ/-categorization by Gender in the middle age group.

	# of tokens	#5	#7	#8	#11
Overall	84	**83% (.86)**	**61% (.63)**	37% (.35)	12% (.10)
F3	42	**93% (.68)**	**76% (.68)**	**55% (.69)**	**19% (.68)**
M3	42	74% (.32)	45% (.32)	19% (.31)	5% (.32)
Significance		.017	.010	.001	.040

Table 6: Rates and probabilities of /æ/-categorization for stimuli #5, #7, #8, and #11 among the 30–49 year olds.

5.4.3 Oldest Age Group

Finally, females and males in the age group composed of subjects aged 50 years and older are remarkably similar throughout the entire continuum. No significant gender effects were found.

The point at which /æ/ was no longer assigned as the dominant category appears to be at stimulus #7. Therefore, the oldest listeners in the sample have a perceptual norm for /æ/ that is identical to the overall norm and most similar to the males of the youngest group. These findings add to the pattern that with age comes a less centralized boundary. Put another way, boundary centralization is a function of youth: the younger the subject, the more retracted the boundary (with exceptions to the males who have the most fronted boundaries, regardless of age).

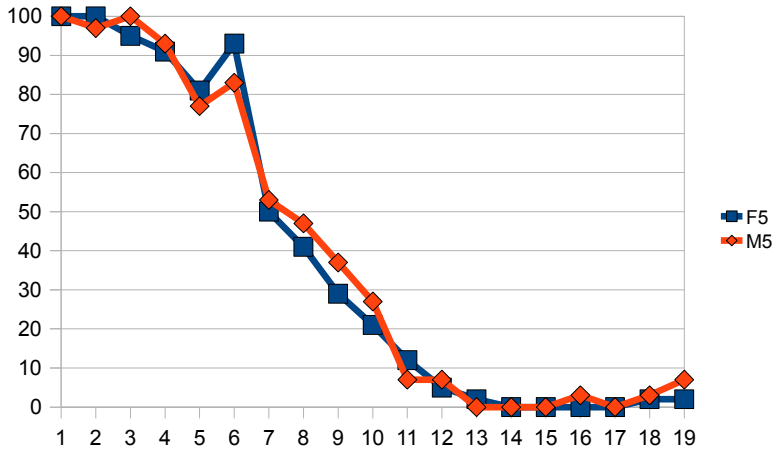

Figure 7: /æ/-categorization by Gender in the oldest age group.

6 Summary and Discussion

The data presented here reflect a listener's knowledge of phonetic variation in the production of /æ/ and /a/ forms. The change in /æ/ is considered to be below the level of social awareness, yet, like patterns found in the speech of Canadians, significant differences in perceptual categorization are also reported. Figure 8 summarizes these differences. This representation shows a centralization process occurring across age groups. Those with the least centralized /æ/-range are males between the ages of 30 and 49. They were found to favor stimuli #1 through #5. Next, both male and female subjects aged 50+ exhibit an /æ/-range consisting of stimuli #1 through #7. Males in the youngest age group join the older listeners here. Females aged 30–49 were found to favor stimuli #1 through #8 pushing the right edge 1 stimulus further than the oldest group. Finally, the most centralized range belongs to the youngest females in the sample who stretched to all the way to #9. In all age groups but the oldest, the boundary is more centralized for females than males.

There is no evidence to suggest that a *shift* in the overall category range is underway (all listeners exhibited comparable high rates of categorization at the left edge), but the apparent-time data strongly show an *extension* of the right edge among younger listeners. In fact, this centralization process appears to be a gendered phenomenon, that is, it is highly regular and consistent among the females. Males, on the other hand, show a mismatch between age and boundary position, though the unifying feature is that they show the least centralized boundaries.

Based on these results, one would predict an asymmetry in perceptual preferences: all listeners in this sample would find forms at the onset of the /æ/-range relatively unmarked, though older listeners should consider forms that the most advanced younger speakers produce (i.e., extending beyond the offset) as marked-sounding. These younger speakers, however, should judge them to be equally fine as the forms at the onset.

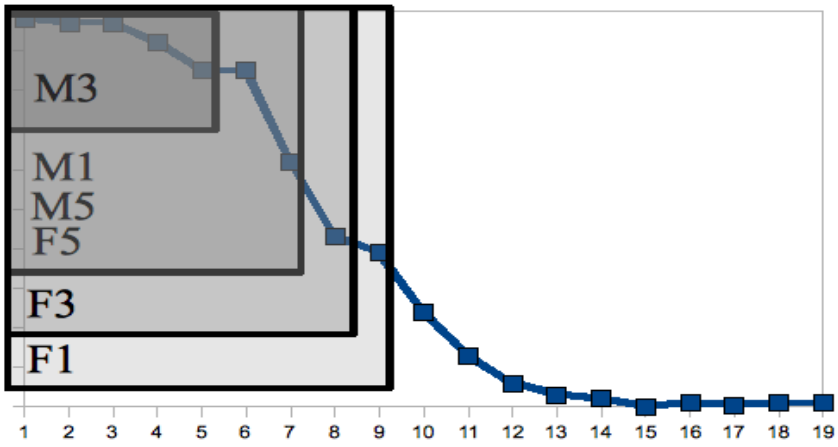

Figure 8: Ranges of /æ/-categorization by Gender and Age Group.

This paper has shown that in addition to small perturbations in the pronunciation of /æ/, there is, on some level, a psychological reflex running along with it. The particular reflex seems to be an extension and not a shift in the mental definition of /æ/ for some listeners in Ontario. This is consistent with the view of perceptual re-analysis presented above and offers an alternative to models that consider sound change as resulting from errors imposed by the listeners, i.e., misperception (cf. Ohala 1981, 1993a, 1993b, 1993c).

References

Boberg, Charles. 2005. The Canadian Shift in Montreal. *Language Variation and Change* 17:133−154.

Clarke, Sandra, Ford Elms, and Amani Youssef. 1995. The Third Dialect of English: Some Canadian evidence. *Language Variation and Change* 7:209−228.

De Decker, Paul. 2002. Hangin' and retractin': Adolescent social practice and phonetic variation. *U. Penn Working Papers in Linguistics 8.3: Papers from NWAV 30*, ed. D.E. Johnson and T. Sanchez, 59−73.

Hagiwara, Rob. 2005. Vowel production in Winnipeg. Paper presented at Canadian English in the Global Context, University of Toronto.

Janson. Tove. 1986. Sound change in perception: An experiment. In *Experimental Phonology*, ed. J. Ohala and J. Jaeger, 253−260. New York: Academic Press.

Labov, William. 1994. *Principles of Linguistic Change: Internal Factors*. Malden, MA: Blackwell.

Labov, William, Sharon Ash, and Charles Boberg. 2006. *The Atlas of North American English: Phonetics, Phonology, and Sound Change*. New York: Mouton de Gruyter.

Labov, William, Malcah Yaeger, and Richard Steiner. 1972. A Quantitative Study of Sound Change in Progress. Report on National Science Foundation Contract NSF-GS-3287. U.S. Regional Survey, Philadelphia.

Ladefoged, Peter. 2005. *A Course in Phonetics.* 5th Edition. Boston: Thomson

Ohala, John. 1981. The listener as a source of sound change. *Papers from the Parasession on Language and Behavior*, ed. C.S. Masek, R.A. Hendrick, and M.F. Miller, 178−203.

Ohala, John. 1993a. Sound change as nature's speech perception experiment. *Speech Communication* 13:155−161.

Ohala, John. 1993b. The phonetics of sound change. In *Historical Linguistics: Problems and Perspectives*, ed. C. Jones, 237−278. London: Longman.

Ohala, John. 1993c. Coarticulation and phonology. *Language and Speech* 36:155−170.

Willis, Clodius. 1972. Perception of vowel phonemes in Fort Erie, Ontario, Canada, and Buffalo, New York: An application of synthetic vowel categorization to dialectology. *Journal of Speech and Hearing*

Research 15:246–255.

Department of Linguistics, SN3050B
Memorial University of Newfoundland
St. John's, Newfoundland, Canada, A1B 3X9
pauldd@mun.ca