The Implementation of Laryngeal Contrast in Korean as a Second Language

Charles B. Chang*

University of California, Berkeley

1 Introduction

Much research in second-language (L2) speech has investigated how L2 learners acquire laryngeal categories that differ from the laryngeal categories of their first language (L1). However, most of this work has concentrated on languages with two laryngeal categories that differ between L1 and L2 in terms of the same primary cue: voice onset time, or VOT (e.g. French and English: Caramazza *et al*. 1973, Flege 1987; Spanish and English: Flege and Eefting 1988; Italian and English: Flege *et al*. 1995; Portuguese and English: Major 1996). In the present study, I examine how L2 learners come to produce a laryngeal contrast that requires the use of a second phonetic dimension in addition to VOT—namely, the three-way Korean laryngeal contrast among lenis, fortis, and aspirated stops, which in initial position differ primarily in terms of VOT and fundamental frequency (f_0) onset (cf. Han and Weitzman 1970, Kim 2004, *inter alia*). How do learners use (or not use) f_0 onset in conjunction with VOT to realize this three-way contrast?

In this paper, I describe the range of variation in phonetic spaces that learners construct for this novel laryngeal contrast and show how these differ from the results of cross-linguistic perception studies on English speakers hearing Korean. The paper is organized as follows: Section 2 provides some background on L2 speech acquisition, including the concept of cross-language equivalence classification; Sections 3 and 4 describe the design of a longitudinal production experiment and present the results of the last week of testing; and Section 5 discusses the implications of the results and summarizes the major conclusions.

2 Background

Foreign accent has been documented in a wide variety of L2 acquisition studies (for a broad overview, see Major 2001). Some of these focus on the effect of learner age on perceived foreign accent (e.g. Flege *et al*. 1995), while others are concerned with identifying which parts of an L2 phonology are the ones most likely to be produced as accented (e.g. Yamada 1995, Broselow *et al*. 1998,

Leather 1996). Two common conclusions of these studies are that "earlier is better" with respect to age of acquisition, and that aspects of the L2 phonology that are the hardest for learners to acquire in a native-like manner tend not to be those that are very different vis-à-vis L1, but rather those that bear similarities to L1 (i.e. those parts of the L2 phonology with L1 counterparts that stand to seriously "interfere" in L2).

While much of the literature provides evidence of L1 transfer in L2 production, it has also been shown that an adult learner's L2 tends to show traits attributable to neither L1 nor L2. These traits parallel features of L1 acquisition, and so are thought to be universal in nature, partly specified by an innately endowed Universal Grammar (cf. White 1989). Historically, this sort of data motivated a move away from simply analyzing learner language in terms of L1 transfer and instead toward looking at it as a semi-autonomous system often referred to as "interlanguage" (cf. Selinker 1972). Interlanguage is thought of as the learner's dynamic system for L2 that incorporates L2 structure, L1 structure, and universally preferred structure—in proportions that vary depending on a number of factors including the time point in the acquisition process (e.g. the amount of L1 influence is generally found to decrease as learners become more proficient in L2, cf. Major 1987).

A prediction that can be made on the basis of these ideas is that—all other things being equal—learners of the same L1 background learning the same L2 will show consistency in the type of foreign accent they manifest, since they are biased by the same L1 and a language-independent Universal Grammar. These theoretically identical biases should lead to the same mappings of L2 sounds to L1 categories, or L1-L2 "equivalence classifications" (Flege 1987). Indeed, studies of L2 perception by naïve listeners (i.e. those who do not know the L2) have found relative consistency of perceptual patterns, such that several types of L2 category mapping possibilities have been identified, as summarized in the Perceptual Assimilation Model (Best 1994).

With regard to L2 perception of Korean, two studies in particular have examined how L1 English speakers interpret Korean word-initial stop consonants. Francis and Nusbaum (2002) found that before training, L1 English speakers (naïve listeners who were not learning Korean) mostly relied on differences in VOT (and co-varying differences in rate of amplitude change) to distinguish the three laryngeal categories, but after training, seemed to use both VOT and f_0 onset differences (along with co-varying differences in the clarity of formant structure at vowel onset) to distinguish them (however, see Shin 2007 for differing results with trained learners of Korean). The perceptual data show, moreover, that after training, English speakers' perception approximates that of native Korean speakers, who break up the two-dimensional phonetic space in the manner shown in Figure 1 (cf. Kim 2004).

While Francis and Nusbaum's (2002) perception study utilized identification and difference rating tasks, Schmidt's (2007) cross-linguistic

perception study instead had subjects—also L1 English speakers with no knowledge of Korean—label Korean sounds as the perceptually closest English sound and rate the similarity of the English sound to the Korean sound. Her results show that subjects overwhelmingly labeled Korean lenis stops and aspirated stops as English voiceless stops and Korean fortis stops as English voiced stops. However, the Korean categories differed in terms of how similar to English categories they were perceived as being: aspirated stops were rated as more similar to English stops than lenis or fortis stops were. This suggests that for L1 English learners of Korean, the "default" equivalence classifications of Korean and English stops are aspirated-voiceless, lenis-voiceless, and fortis-voiced, but that the strength of the cross-language category identification varies across category pairings.

Figure 1. Schematic of the Korean laryngeal contrast in terms of voice onset time and fundamental frequency onset (based on Kim 2004).

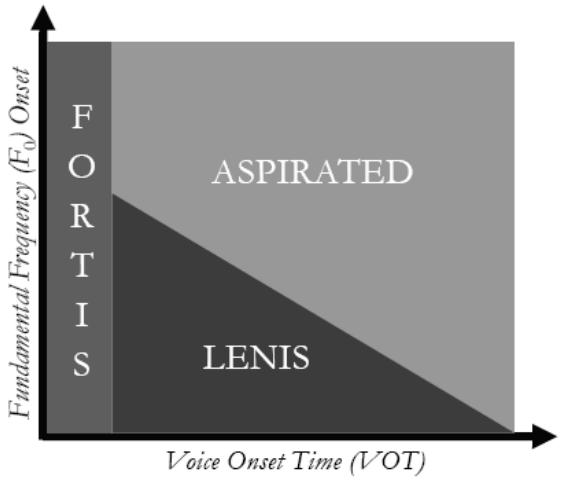

L2 production, however, differs from L2 perception in a number of ways. Even on just a practical level, while it is possible for an L2 perception task to be completed by L2-naïve subjects, it is not possible for an L2 production task to be completed in the same way: in order to speak, subjects need to know what sounds they are aiming for. Thus, studies that have looked at L2 production of Korean have necessarily examined people actively learning the language, including its sound categories. In one such study, Kim and Lotto (2002) found that intermediate Korean learners (most of whom were L1 English speakers) produced distinctions between the three stop types using VOT, but not closure duration or f_0 onset. Shin's (2007) study of elementary Korean learners resulted in similar findings with L1 English learners, who tended to rely just on VOT to

produce the contrast. On the other hand, learners whose L1 was a tone language (e.g. Mandarin, Cantonese, Taishanese) were found to use f_0 as a cue more often than the L1 English learners.

Taken together, the results of studies of L2 perception and production of Korean suggest that L1 English speakers, and perhaps speakers of non-tone languages more generally, can be trained to use f_0 in perception, but nevertheless tend to utilize VOT rather than f_0 to distinguish the Korean laryngeal categories in production. This pattern of production contrasts with that of mature native speakers, who use both dimensions, as well as with that of children acquiring Korean as L1, who separate the lenis and aspirated categories in f_0 well before they separate them in VOT (cf. Jun 2006).

Although Schmidt's (2007) cross-linguistic perceptual findings show consistency in the way learners assimilate Korean categories to English categories, they make no predictions regarding how learners will distinguish the lenis and aspirated categories that are both assimilated to the voiceless category of English. Kim and Lotto (2002), as well as Shin (2007), suggest that learners mainly use VOT to distinguish these categories in production; however, the amount of VOT overlap between learners' lenis and aspirated stop productions—even within one place of articulation—is so large that it is unclear whether learners are actually producing a reliable three-way contrast in VOT.

Another reason to reexamine L2 learners' production of this contrast is the existence of a conflict in cues contributing to cross-language equivalence classification. If we were to pair the Korean and English laryngeal categories on the basis of phonetic similarity (specifically, in terms of similarity in VOT and f_0), aspirated stops would be paired with voiceless stops, since these categories are both long in VOT and high in f_0 onset. However, it is unclear how lenis stops and fortis stops should be classified, since each of these categories resembles voiced stops in one way and voiceless stops in another way. Lenis stops are relatively long in VOT like voiceless stops, but low in f_0 like voiced stops; fortis stops, on the other hand, are short in VOT like voiced stops, but high in f_0 like voiceless stops. Thus, linking lenis and fortis stops to English categories is not straightforward, given that most English speakers show sensitivity to the f_0 difference between voiced and voiceless stops (cf. Haggard *et al*. 1970).

In the present study, I reexamine how L1 English late learners of Korean produce the Korean laryngeal contrast, focusing on an L1 and L2 that do not share the same orthography to avoid the confound of orthographic equivalence present in the majority of studies on L2 voicing categories. The main research question is the following: given little to no explicit phonetic instruction, how successful are late learners of Korean at producing Korean laryngeal categories like native speakers? We will see if, using VOT and f_0 onset, learners manage to produce a full three-way contrast, as well as if they are consistent in their interlanguage phonetic spaces. Finally, we will make some generalizations about the nature of learners' deviation from the native Korean phonetic space.

3 Methods

The longitudinal production experiment took place in a quiet dormitory room and was conducted weekly starting from one week into the language class participants were taking. Every week participants completed a reading task in which they saw a Korean stimulus (spelled in Korean orthography) and read it aloud. Stimuli were presented a total of four times, once each in four randomized blocks following a practice session of five items. Each item was presented on screen for 1.5 seconds and then replaced by a picture of a green traffic light to cue the participant to produce the item. Audio was recorded via a head-mounted condenser microphone for two seconds starting at the time point at which the green light appeared on screen, and the inter-stimulus interval from the end of this recording to the presentation of the following item was one second. All stimuli presentation and audio recording was done in DMDX 3.2.6.3 (Forster 2008) on a laptop computer.

The set of Korean stimuli consisted of 22 Korean monosyllables representing most of the phonemic contrasts in the language. The stimuli were generally of the form CV to make them as easy as possible for novice learners to read. The same set of stimuli was used in every week of the study.

Participants were 26 late learners of Korean (4 males, 22 females; 21–26 years old), native speakers of American English with no prior exposure to Korean taking a six-week course of intensive Korean immersion instruction at the time of the study. On average these learners received four hours of instruction a day, for a total of approximately 82 hours of instruction by the end of the program (roughly equivalent to one semester of college-level Korean). In exit questionnaires, participants reported that class time constituted the majority of their experience with Korean, both in terms of listening and speaking.

Participants' recordings were acoustically analyzed in Praat 5.0.26 (Boersma and Weenink 2008). Manual measurements of VOT and f_0 onset were taken on learners' productions of the nine critical items beginning with plosives (3 laryngeal categories x 3 places of articulation). VOT was measured off a wide-band Fourier spectrogram with a Gaussian window shape (window length: 5 ms; dynamic range: 50 dB; pre-emphasis: 6.0 dB/oct) as the time at voicing onset minus the time at the stop burst. To obtain stable measurements of f_0 onset, the average wavelength of the first three regular glottal periods in the vowel was calculated from the waveform and converted into a frequency value. Initial periods were skipped if they were irregular (e.g. more than 33% longer or shorter than the following period); however, tokens requiring more than five periods of the vowel onset to be skipped were discarded.

In order to put male and female learners on the same f_0 scale, raw f_0 measurements were standardized to *z*-scores by learner, by subtracting the learner's mean f_0 over the duration of the study and dividing by the square root of the learner's variance in f_0 over the duration of the study.

4 Results

Below we will first examine the performance of native Korean speakers on the production task, followed by the performance of L2 learners on the same task at the end of the study. In all plots, lenis productions are plotted in white circles, fortis productions in black squares, and aspirated productions in gray diamonds.

Figure 2. Scatter plots of native Korean speakers' productions.

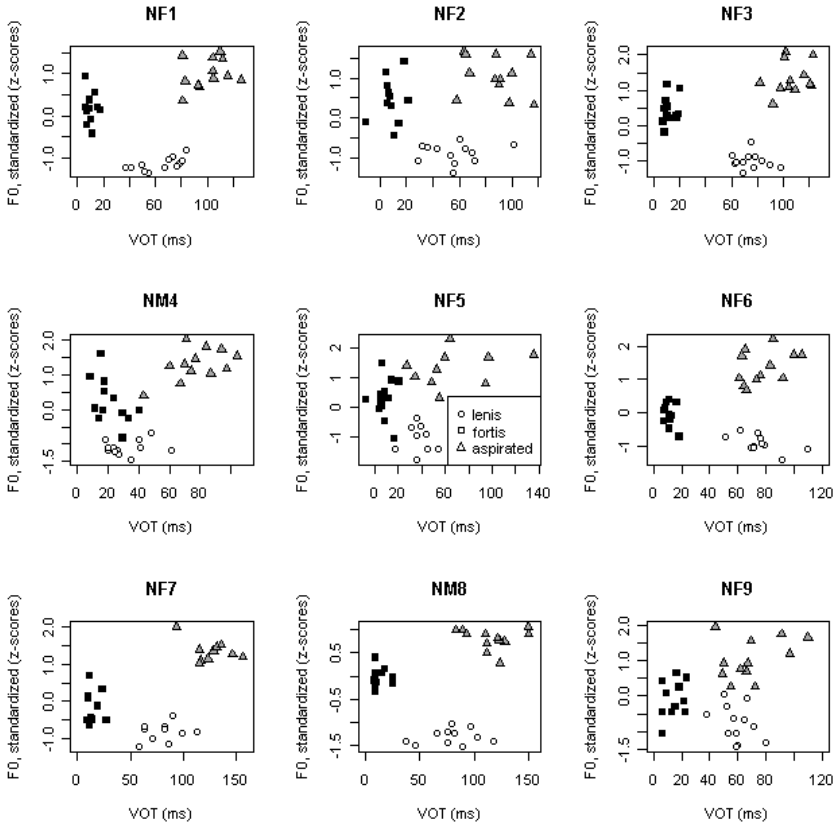

The phonetic spaces of native Korean speakers—learners' Korean teachers and resident assistants in their dormitory—are generally consistent with Kim (2004) in terms of how the Korean laryngeal categories are realized with respect to VOT and f_0 onset. Fortis stops are produced with short VOT and an elevated f_0 onset; lenis stops are produced with longer VOT and a low f_0 onset; and aspirated stops are produced with the longest VOT and the highest f_0 onset (cf.

Figure 2). For most native speakers, lenis and aspirated stops overlap considerably in VOT, and fortis and aspirated stops overlap considerably in f_0, but none of these categories overlap in both dimensions. Thus, VOT and f_0 are necessary and sufficient cues for distinguishing the three laryngeal types.

Figure 3. Representative scatter plots of Week 5 productions in learner group A1 (top left: LM23), A2 (top right: LF54), and A3 (bottom: LF28, LF46).

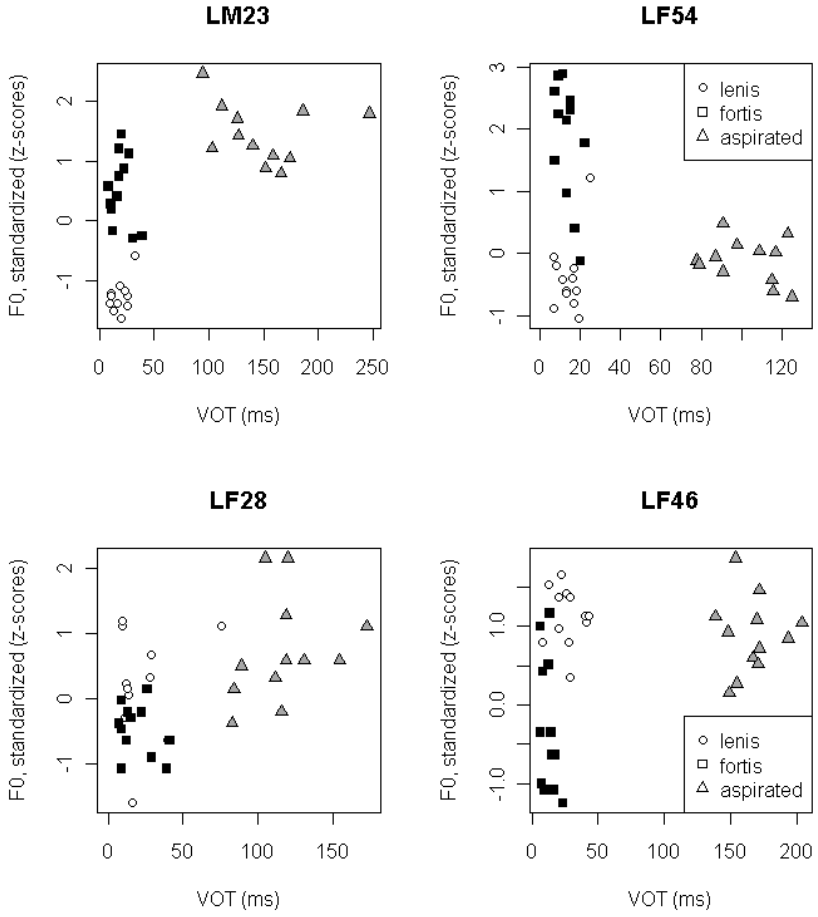

The phonetic spaces of L2 learners look markedly different. One of the most common patterns is found in Groups A (n=7) and B (n=2), where learners essentially produce two two-way contrasts, each in one dimension. In subgroup A1, lenis and fortis stops are both produced with short VOT and are contrasted

on f_0, while fortis and aspirated stops tend to be produced with similar f_0 and are contrasted on VOT (cf. Figure 3, LM23). Subgroup A2 is similar, except that aspirated stops are produced with a relatively low f_0 onset in the range of the lenis stops rather than an elevated f_0 onset in the range of the fortis stops (cf. Figure 3, LF54). Subgroup A3 (cf. Figure 3, LF28 and LF46) resembles subgroup A2, except lenis and fortis stops have been switched in the f_0 dimension: here, lenis stops are produced with higher f_0 than fortis stops, though lenis and aspirated stops are still produced in the same f_0 range.

In Group B, fortis and lenis stops are produced in the same f_0 range and are distinguished on the basis of VOT, while lenis and aspirated stops are produced in the same VOT range and are distinguished on the basis of f_0 (cf. Figure 4).

Figure 4. Representative scatter plots of Week 5 productions in learner group B.

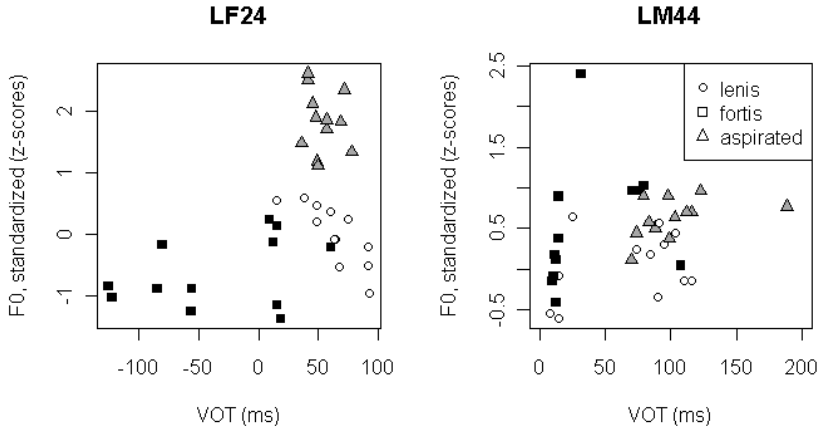

In Group C (n=7), learners produce a three-way contrast either using VOT, using f_0, or using both dimensions. The learners in subgroup C1 (e.g. LF25, cf. Figure 5) make use of both VOT and f_0 to make the contrast, producing fortis stops with short VOT and low f_0, lenis stops with longer VOT and higher f_0, and aspirated stops with the longest VOT and highest f_0. However, the learners in subgroup C2 (e.g. LF52, cf. Figure 5)—like the learners described in Kim and Lotto (2002) and Shin (2007)—rely just on VOT to make a three-way contrast. In contrast, the learner in subgroup C3 (LF04, cf. Figure 5) relies almost entirely on f_0 to make the contrast, producing all three categories in the short-lag VOT range and distinguishing between them by producing lenis stops with the lowest f_0, aspirated stops with intermediate f_0, and fortis stops with the highest f_0.

Moving on to learners who fail to keep the three categories apart with these cues, we see in Group D (n=8) learners who just produce a two-way contrast and, within this group, nearly all possible types of merger. Learners in subgroup D1

merge lenis and aspirated stops in the long-lag VOT region (cf. Figure 6, LF01 and LF29), while learners in subgroup D2 merge lenis and fortis stops in the short-lag VOT region (cf. Figure 6, LF06). In addition, the learner in subgroup D3 even merges fortis and aspirated stops in the long-lag VOT region (cf. Figure 6, LM43).

Finally, in Group E (n=2) learners do not keep any of these categories distinct from the others in terms of VOT and/or f_0 onset, and instead produce all of them over the same wide phonetic space.

Figure 5. Representative scatter plots of Week 5 productions in learner group C1 (top: LF05, LF25), learner group C2 (bottom left: LF52), and learner group C3 (bottom right: LF04).

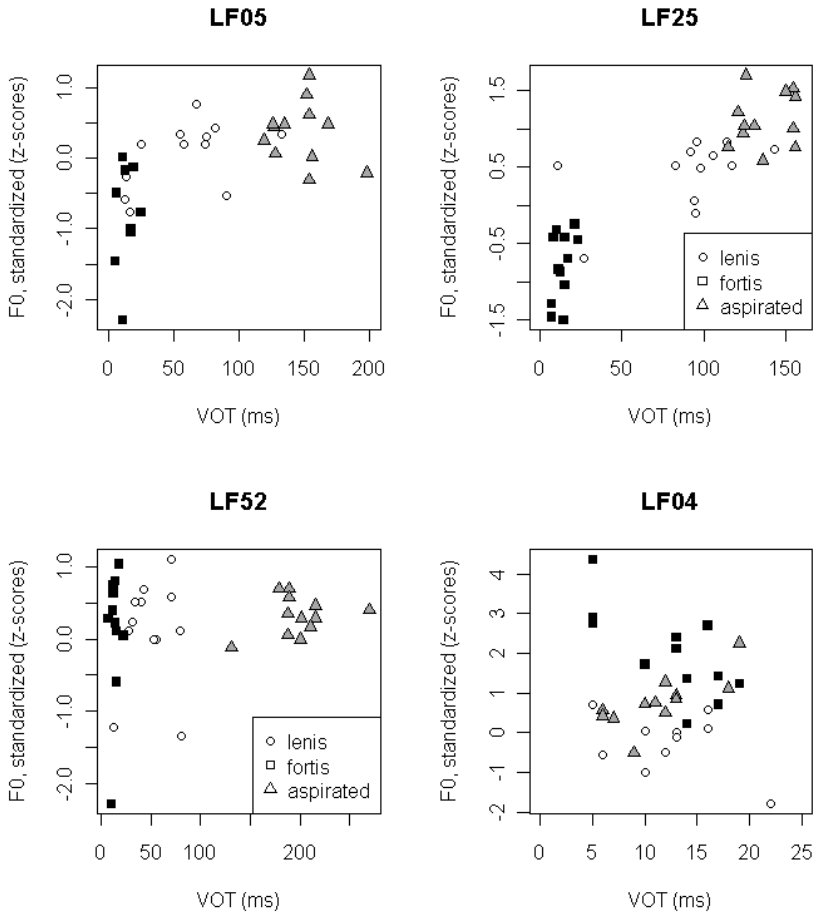

Figure 6. Representative scatter plots of Week 5 productions in learner group D1 (top: LF01, LF29), learner group D2 (bottom left: LF06), and learner group D3 (bottom right: LM43).

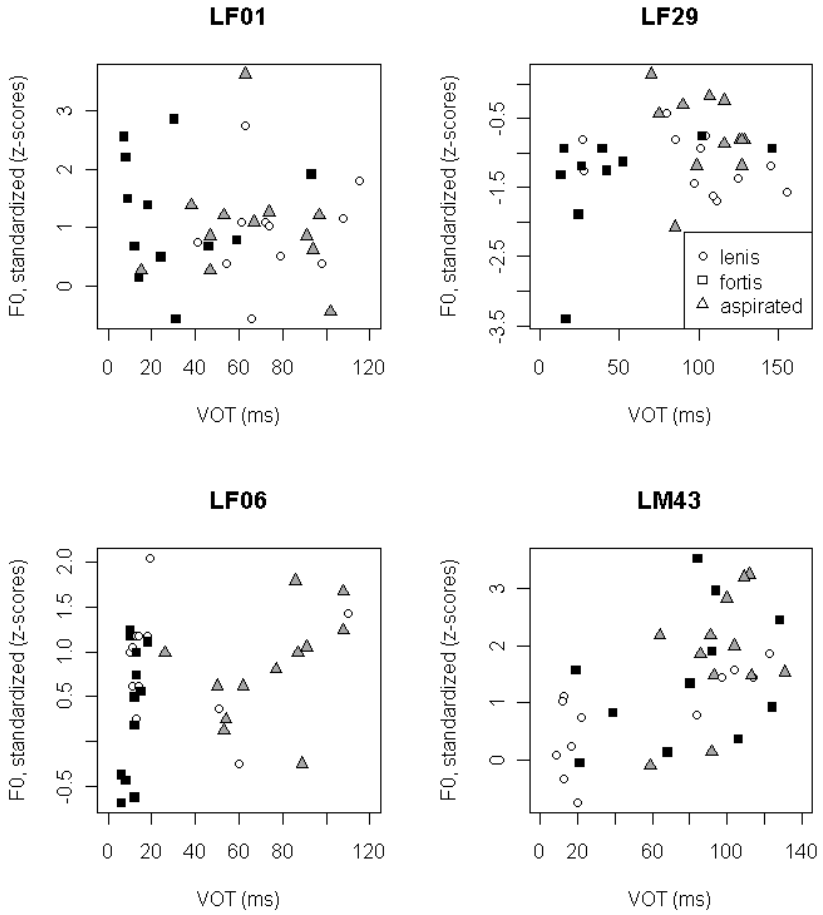

5 Discussion and Conclusions

To summarize, there is wide variation in learner success at restructuring the L1 phonetic space of two laryngeal categories into an L2 phonetic space of three laryngeal categories that resembles native Korean. Some learners fail to produce a three-way contrast, merging two or more categories with different degrees of overlap, but many learners produce three distinct L2 categories. In short, while

incomplete L2 contrast is found, more commonly we see learners finding some way to implement the full L2 contrast.

In addition, there is a dichotomy in the phonetic spaces of learners who produce a three-way contrast (the "full distinguishers") and those who only produce a two-way contrast (the "partial distinguishers"). In both groups, some learners appear to identify lenis stops as a category similar to voiced stops—full distinguishers separating fortis stops from lenis stops on the basis of f_0 onset, and partial distinguishers combining fortis and aspirated stops into a category similar to voiceless stops. However, in both groups there are other learners who identify fortis stops as the voiced-like category. Here the full distinguishers separate lenis stops from fortis stops on the basis of VOT and/or f_0, while the partial distinguishers combine lenis and aspirated stops into a voiceless-like category. These findings are consistent with the predictions of the ambiguous cross-linguistic category correspondences described above. Despite the fact that these learners have the same L1 background, when they are given no explicit instruction on how to produce the Korean laryngeal contrast, they interpret the contrast in multiple ways, resulting in disparate phonetic spaces of Korean that are all quite different from the native Korean phonetic space.

We are left then to wonder: why is there so much variation? If we ignore the influence of affective variables (e.g. motivation, language attitudes, cultural identification), which, as suggested by background questionnaires, do not differ across the groups delineated above in any clear way, we are left with four possible explanations for the variation in learner production.

First, one might posit that beginning L1 English learners of Korean rely heavily on transliterations of Korean, and that variation in production just reflects variability of transliterations (e.g. lenis stops are transliterated alternately as <p, t, k> or <b, d, g> depending on the system). While such variation in transliteration may be related to some apparent differences in cross-language equivalence classifications across learners, it cannot fully account for the variety of phonetic spaces seen in this study. For one, only two laryngeal categories can be orthographically distinguished with wholly different graphemes. Consequently, it is not clear how one can derive from this primarily two-way orthographic variation the multifaceted acoustic variation seen above—in particular, the numerous differences in whether and how the subphonemic cues of VOT and f_0 are used in distinguishing the three L2 laryngeal categories (e.g. LF25 vs. LF52, cf. Figure 5). Moreover, certain production patterns do not line up with transliteration. It can be seen above (cf. Figure 4, LF24) that the fortis series, which is very infrequently transliterated with graphemes for voiced stops, is nonetheless produced as prevoiced (i.e. with voicing prior to stop release) by some learners. This indicates that even if variability in transliteration is responsible for some of the variability in production, it cannot account for all of it.

Second, one could attribute variation in production to variation in input. After all, learners had different Korean teachers, and there are differences among the teachers with respect to the precise organization of their phonetic spaces, though the general pattern is the same (cf. Figure 2). However, inspection of differences among learners along with their class affiliations does not reveal the sort of correspondences one would expect if learners' differences in production simply reflected similar differences in their respective teachers' production. For example, learners LM23 and LF54 were in the same class taught by the same teachers, yet still differ from each other in production: LM23 produces aspirated stops with high f_0, while LF54 produces them with low f_0 (cf. Figure 3), even though their teachers (NF1 and NF7) both produce them, as expected, with high f_0. These facts suggest that even if some differences among learners are rooted in input disparities, input cannot be the whole story.

Third, it is possible that VOT and f_0 are weighted differently across participants in terms of their importance in distinguishing English voiced and voiceless stops. This variability in cue weighting could lead to variation in L2 production, in that learners would not necessarily be biased towards the same schemas of L1-L2 equivalence classifications, but potentially towards different ones. The fact that there is some variability among English speakers with respect to how sensitive they are to f_0 as a cue to the English voicing contrast (cf. Haggard *et al*. 1970) is consistent with this explanation—an interesting possibility that should be tested more thoroughly.

Fourth, it could be that learners utilize explicit strategies to achieve L2 contrast that may or may not be based on actual L2 input patterns (such strategies being likely to differ between individuals). In fact, strategy does seem to account for what at least some learners do. For instance, learner LF52 (a learner in Group C2, where a three-way contrast is produced in the VOT dimension only) expressed in study debriefings that she thought the contrast just had to do with aspiration, and so she ignored pitch. Her coming to the table with that sort of strategic bias largely accounts for why she started producing a three-way VOT contrast in Week 1 of the language program and continued to do so through Week 5, failing to use f_0 at all time points in this study.

On a final note, it should be noted that the results of this production study differ substantially from those of the perception studies described above. Relatively few learners produce the L2 laryngeal categories with a phonetic space that one might predict from cross-linguistic perception results, or with a phonetic space resembling that of native speakers (cf. Figures 1–2). Moreover, among the majority of learners with phonetic spaces that depart significantly from the native phonetic space, there is a large amount of variation in the organization of the phonetic space, in contrast to the high degree of consistency seen in the perceptual performance of listeners in Schmidt (2007). Some possible sources of this variation have been discussed here, but more work is needed to tease apart their effects.

Notes

* Research represented in this paper has been supported in part by a National Science Foundation Graduate Research Fellowship, a UC Berkeley Center for Korean Studies Graduate Student Fellowship, and travel grants from the UC Berkeley Department of Linguistics. I am grateful to Keith Johnson, members of the UC Berkeley Phonology Lab, and the audience at HISOKL 2009 for their helpful comments and feedback; to the Fulbright Korean-American Educational Commission for its assistance with study logistics; and to all study participants for their perseverance during a hot and humid summer. Any errors are my own.

References

Best, Catherine T. 1994. The Emergence of Native-Language Phonological Influences in Infants: A Perceptual Assimilation Model. In *The Development of Speech Perception: The Transition from Speech Sounds to Spoken Words*, eds. J. Goodman and H. C. Nusbaum. Cambridge, MA: MIT Press, 167–224.

Broselow, Ellen, Su-I Chen, and Chilin Wang. 1998. The Emergence of the Unmarked in Second Language Phonology. *Studies in Second Language Acquisition* 20(2).261–280.

Caramazza, Alfonso, Grace H. Yeni-Komshian, Edgar B. Zurif, and Evelyn Carbone. 1973. The Acquisition of a New Phonological Contrast: The Case of Stop Consonants in French-English Bilinguals. *Journal of the Acoustical Society of America* 54(2).421–428.

Flege, James Emil. 1987. The Production of "New" and "Similar" Phones in a Foreign Language: Evidence for the Effect of Equivalence Classification. *Journal of Phonetics* 15(1).47–65.

Flege, James Emil, and Wieke Eefting. 1988. Imitation of a VOT Continuum by Native Speakers of English and Spanish: Evidence for Phonetic Category Formation. *Journal of the Acoustical Society of America* 83(2).729–740.

Flege, James Emil, Murray J. Munro, and Ian R. A. MacKay. 1995. Effects of Age of Second-Language Learning on the Production of English Consonants. *Speech Communication* 16(1).1–26.

Francis, Alexander L., and Howard C. Nusbaum. 2002. Selective Attention and the Acquisition of New Phonetic Categories. *Journal of Experimental Psychology: Human Perception and Performance* 28(2).349–366.

Haggard, Mark, Stephen Ambler, and Mo Callow. 1970. Pitch as a Voicing Cue. *Journal of the Acoustical Society of America* 47(2B).613–617.

Jun, Sun-Ah. 2006. Phonological Development of Korean: A Case Study. In *Harvard Studies in Korean Linguistics XI*, eds. S. Kuno *et al*. Seoul: Hanshin Publishing Company, 17–33.

Kim, Midam. 2004. Correlation between VOT and F0 in the Perception of Korean Stops and Affricates. In *INTERSPEECH-2004*, 49–52.

Kim, Mi-Ran Cho, and Andrew J. Lotto. 2002. An Investigation of Acoustic Characteristics of Korean Stops Produced by Non-Heritage Learners. In *The Korean Language in America*, Vol. 7, ed. J. J. Ree. American Association of Teachers of Korean, 177–188.

Leather, Jonathan. 1996. Interrelation of Perceptual and Productive Learning in the Initial Acquisition of Second-Language Tone. In *Second-Language Speech: Structure and Process*, eds. A. James and J. Leather. Berlin: Mouton de Gruyter, 75–101.

Major, Roy C. 1987. A Model for Interlanguage. In *Interlanguage Phonology: The Acquisition of a Second Language Sound System*, eds. G. Ioup and S. H. Weinberger. Cambridge, MA: Newbury House Publishers, 101–124.

Major, Roy C. 1996. L2 Acquisition, L1 Loss, and the Critical Period Hypothesis. In *Second-Language Speech: Structure and Process*, eds. A. James and J. Leather. Berlin: Mouton de Gruyter, 147–159.

Major, Roy C. 2001. *Foreign Accent: The Ontogeny and Phylogeny of Second Language Phonology*. Mahwah, NJ: Lawrence Erlbaum Associates.

Schmidt, Anna Marie. 2007. Cross-Language Consonant Identification: English and Korean. In *Language Experience in Second Language Speech Learning: In Honor of James Emil Flege*, eds. O.-S. Bohn and M. J. Munro. Amsterdam: John Benjamins Publishing, 185–200.

Selinker, Larry. 1972. Interlanguage. *International Review of Applied Linguistics in Language Teaching* 10(3).209–231.

Shin, Eurie. 2007. How Do Non-Heritage Students Learn to Make the Three-Way Contrast of Korean Stops? In *The Korean Language in America*, Vol. 12, ed. H.-S. Wang. American Association of Teachers of Korean.

White, Lydia. 1989. *Universal Grammar and Second Language Acquisition*. Amsterdam: John Benjamins Publishing.

Yamada, Reiko A. 1995. Age and Acquisition of Second Language Speech Sounds: Perception of American English /r/ and /l/ by Native Speakers of Japanese. In *Speech Perception and Linguistic Experience: Issues in Cross-Language Research*, ed. W. Strange. Baltimore: York, 305–320.

University of California, Berkeley
Department of Linguistics
1203 Dwinelle Hall #2650
Berkeley, CA 94720, U.S.A.
e-mail: cbchang@cal.berkeley.edu