**Individual variation in laboratory sound change**

In many languages, a voiceless stop preceding an approximant may develop into an affricate. The narrower degree of constriction of an approximant may encourage aspiration of a stop to develop into frication, which can become associated with the stop+approximant environment. Historically, in American English, two sound changes have occurred which resulted in stop affrication: /tr/ -> [ʧr] and /tj/ -> [ʧj], as in the words *trick* and *mature*. Another stop+approximant cluster remains which exhibits variation, but has not yet developed a sound change into an affricate: /tw/. The palatal nature of /j/ and /r/ have encouraged a retracted affricate to develop in this environment, but /w/ has two points of articulation, velar and labial, so it is not a foregone conclusion that a resulting affricate must be palatal; /ts/ is also a possible affricate in this environment (c.f. German /t/->[ts] as in *Zwillinge*, or Japanese /t/ -> [ts] in front of /ɯ/). However, the parallel trajectory of /tr/ and /tj/ would encourage a palatal affricate, and lip-rounding accompanying /w/ would create a lower-pitched frication, which could sound more retracted. At this point in time, both variants are acceptable and in use by a subset of the American English speaking population; although, the retracted affricate [ʧw] seems to be more common than the front affricate [tsw]. The variation is something that most people do not notice, and has not yet reached the status of marker or indicator.

Making use of the two variant pronunciations, an experiment designed to elicit sound change in the laboratory combined the methodologies of perceptual learning and convergence in order to give English speaking participants the opportunity to change their production and perception of the consonant cluster /tw/. Generally, this experiment was successful: most participants developed a shift in both perception and production of /tw/ in the direction of the pronunciation variant they were trained on. Most also developed a change in their perception of the related phonological environment /tu/. There were, however, 9 out of 28 participants who did not exhibit significant pronunciation change (convergence) during the shadowing task. Generally, these non-convergers also showed a less robust effect of perceptual learning, as evidenced by lower acceptance rates than convergers for words containing the pronunciation variants during a lexical decision task. Their acceptance rates of other trained and untrained filler words was equivalent to the convergers, so it was not merely a general attention difference. Three of the non-convergers exhibited only a very small change during shadowing, which increased to significant levels after the lexical decision task. Three other non-convergers seemingly ignored the variation and maintained a more-or-less consistent pronunciation of their own. Still three others rebelled against the training variant, making their productions even less like the trainers than they were before the experiment.

This paper explores those non-convergers' production and perceptual responses to the experiment, relative to the convergers, and examines factors that may have influenced their non-convergence. One of the big factors was the original starting point of their pronunciation. All three of those who used both pronunciation variants before the experiment displayed convergence. From this result, we can hypothesize that flexibility positively affects convergence. However, 5 out of 9 of the non-convergers used only one or the other affricated variant before training, suggesting that just because a talker uses an innovative variant, does not mean that he or she will converge in a given situation.

In addition to numerical descriptions of the non-convergers' performance, there are descriptive elements that add to the analysis. One non-converger expressed extreme displeasure during and after the experiment, saying that "one guy was very annoying." This participant's acceptance rates for words containing the training variant were lower than average. It is difficult to say whether the pronunciation made the talker "annoying" or if it was some other quality; however, there were many comments about the "strange" and "wrong" pronunciation of the talkers in the front training condition, while none of the participants expressed any disapproval of the stimuli in the retracted training condition. Participants who heard the [tsw] pronunciation variant as [sw] were disrupted because of the apparent mismatch between the <tw> spelling and the perceived [sw] pronunciation. Obviously, phonemic perception of a pronunciation variant influences perceptual learning, and it is probable that two different mechanisms, one conscious and one not, that are used to learn a phonetically different pronunciation versus a phonemically different pronunciation. Other factors that can be tied to convergence and non-convergence and robustness of perceptual learning, such as dialect background, gender, attitude, phonetic and phonological properties of the training variants, and experimental design are discussed in the paper.

A probabilistic model of sound change might make the following assumptions: Phonetic variation of a single sound may be statistically regular, with a distribution of its occurrence in all of the words in a language, by all the speakers of a language, that is centered over a prototypical acoustic or articulatory space, with variant pronunciations clustered normally around the center. But not every speaker of a language talks to every other speaker of a language or uses all of the words, in all the same proportions; therefore, every speaker of a language will have a pattern of variation that differs from the population norm. These patterns evolve over time as the pronunciation variants that are heard and spoken are closer or further from the prototype, or skew in one or more directions. If these assumptions are true, it is also true that sound systems are constantly changing and evolving. However, without additional constraints, this hypothetically simplistic base model would over-predict the rate of sound change. We can hypothesize some of these necessary constraints based on the observed outcomes of talkers and listeners in experimental conditions, especially when they do not perform as expected.