Volume 14
Issue 1 *Proceedings of the 31st Annual Penn Linguistics Colloquium*

Article 3

April 2008

Addressing the actuation problem with quantitative models of sound change

Adam Baker
University of Arizona

Addressing the actuation problem with quantitative models of sound change

Abstract

Computational models are presented that evaluate different theories of sound change, particularly with regard to the actuation of change. Standard phonologization of coarticulation models predict counterfactual across-the-board change (cf. Weinreich, Labov, and Herzog 1968). Models that simulate a sigmoidal trajectory of change are more empirically appealing, but also are very sensitive to initial conditions. It is proposed that herein lies the solution to the actuation riddle. Sound change arises when a linguistic leader (Labov 2001) perceives an incidental correlation of social and phonetic variables, and adopts her speech to the "change." This simple incident leads to an entire sound change. We expect sound change to arise with the same frequency as these spurious correlations. The (presumed) infrequency of such correlations offers a schematic solution to the actuation problem.

Addressing the Actuation Problem with Quantitative Models of Sound Change

Adam Baker*

1 Introduction

This paper seeks to address the actuation problem by using computational models of speakers to assess the factors that influence the initiation of sound change. The question raised by the actuation problem is simple: why does sound change occur? Conversely, why does sound change not occur? The problem was first formulated by Weinreich et al. (1968), who raise the issue in a critique of Neogrammarian theory, as expounded by Paul (1880). A solution to the actuation problem is presented here, albeit in somewhat schematic form.

The paper's structure is tripartite. First, phonologization-of-coarticulation models (henceforth, POC models) are examined, and rejected for their inadequacy in dealing with the actuation problem. Then a different kind of model is introduced, based on Labov (2001) and Rogers (1962), that produces sigmoidal trajectories of change. These models are used to determine a locus for the solution to the actuation problem. It is proposed that sound change arises when linguistic leaders observe incidental correlations between social and phonetic variables. If these leaders adopt the change, it spreads to others, initiating a sound change. Since incidental correlations are not expected to occur very frequently, sound change is not expected to be very frequent. This explains why sound change is possible, but occurs relatively infrequently.

2 Inadequacy of Phonologization of Coarticulation Models

Phonologization of coarticulation models have a number of advantages. Since the earliest systematic descriptions of sound change (e.g. Whitney 1867/1896: 69), linguists have considered change to be the gradual accommodation of language to speakers' needs. The POC theory of change has several advantages. It is intuitively appealing, congruent with present-day ideas about markedness, and current in the literature (e.g. Ohala 1983, Pierrehumbert 2001). It also has the distinct advantage of being easily modeled.

*Thanks to Bill Labov, Tony Kroch, Don Ringe, and Christina Villafaña Dalcher for helpful comments. Diana Archangeli, Peter Richtsmeier, Alina Twist, and Yuko Watanabe all offered helpful comments in the preparation of the talk.

To create a simple POC model, we create a model of individuals and connect them in a social network. In the very simple models represented here, individuals speak only one word, and the word is defined by a single phonetic parameter. A speaker's production target relates to what the speakers produce (detailed below). Each speaker also has a prestige rating. The following formula is used to match a simple intuition about prestige: few speakers have high prestige, while many speakers have average prestige. The prestige Π of the i^{th} speaker is given by the expression:

$$\Pi_i = e^{-i/12.5} \tag{1}$$

A speaker's production P is the sum of the production target, random noise, and a coarticulatory bias. These terms are shown in the equation below. s_t is a speaker's production target at time t. $0.1N(0,1)$ is random noise: 0.1 times a normally-distributed random variable with mean zero and variance 1. -0.02 is the coarticulatory bias. Productions are forced to remain between -1 and 1.

$$P_t = s_t + 0.1N(0,1) - 0.02 \tag{2}$$

At every time step of the simulation, individuals speak to people they know, and update their production target based on what they hear, according to the formula below. P_t is the speaker's own production. O_t is the average of the productions the speaker heard, weighted by the prestige of those speakers.

$$s_{t+1} = 0.9s_t + 0.1(0.5P_t + 0.5O_t) \tag{3}$$

This last equation provides the crucial production-perception loop. Since speakers update their targets based on what they hear, their production targets can be affected (over time) by the coarticulatory bias.

Speakers are finally connected in a social network. A network was constructed following the general intuition that more prestigious people have more social connections. A speaker was connected to another speaker with probability 0.2Π; this means that more prestigious speakers are likely to have more social connections. Then, connections were made reciprocal (i.e. if speaker i was connected to speaker j, then speaker j was connected to speaker i). This last step made the correlation between prestige and the number of connections a speaker has weaker than it would have been otherwise, but a correlation was still observable in plots of prestige against number of connections.

Figure 1 shows the result of 100 runs of this model for 500 time steps. In 100 of 100 runs, sound change occurs every time, immediately. In fact this same result obtains if a thousand or a million runs are conducted. If the constants in the above equations are modified—i.e. if the size of the coarticulatory bias is changed to another non-zero number, or any of the other numbers are modified in such a way that terms would not drop out of the equations—then the direction or rate of change is altered, but change will always occur.

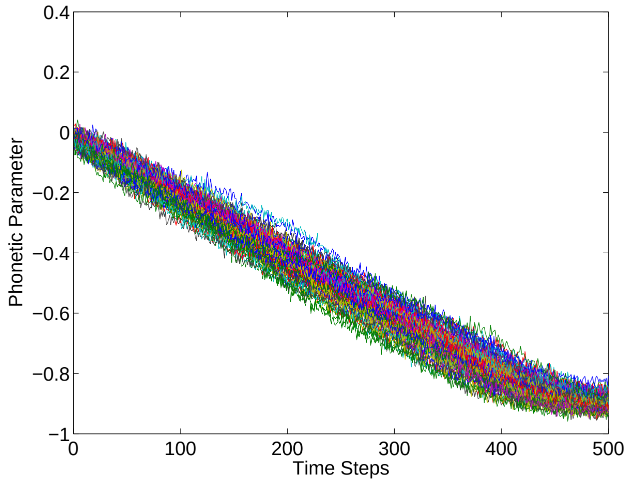

Figure 1: Invariable sound change produced by a phonologization of coarticulation model.

A POC model would then be appropriate for describing language change if sound change always occurred whenever possible, in every language, and at every time. Since this is not what is observed in sound change, POC models are clearly inadequate. No clearer formulation of this problem exists than that of Weinreich, Labov, and Herzog:

> But if the pursuit of ease is the cause of sound change in idiolects, the fundamental questions arise: why do not speakers go about it more quickly, and why do Language Customs split in that some speakers set out on a particular ease-seeking path whereas others retain their less comfortable pattern? . . . For even when the course of a language change has been fully described and its ability explained, the question always remains as to why the change was not actuated sooner, or why it was not simultaneously actuated wherever identical functional properties prevailed. The unsolved actuation riddle is the price paid by any facile and individualistic explanation of language change. It creates the opposite problem—of explaining why language fails to change. (1968: 111–112)

Consequently we must reject POC models entirely, as inconsistent with the empirical data. The next section will provide an alternative model that is empirically more satisfying than the POC model, and that offers indications of the locus of the solution to the actuation problem.

3 Modeling the Sigmoidal Progression of Change

Several researchers have noted a characteristic sigmoidal progression to language change (Bailey 1973, Kroch 1989, Labov 1994, 2001). Change begins slowly, accelerates to a peak rate-of-change midcourse, and then tapers off as the change is completed. The presence of a sigmoidal curve is not at all surprising in linguistic change, since such curves are found frequently in diffusion of innovation research (Rogers 1962), the sociological study of adoption of technology or practices. We can obtain insight into linguistic change by considering change to be the diffusion of a linguistic innovation throughout the community.

The following model is based on the linear model of sound change provided by Labov (2001). It differs from Labov's model, however, in the following ways: more speakers from each age are modeled, there are modifications to the birthing parameters of women, change is considered to be phonetically abrupt rather than phonetically gradual, and speakers have the option of adopting the sound change or not. Only women are modeled, since they are often considered to be the active agents of sound change (Labov 2001).

The model begins with a population of individuals with a uniform age distribution, between 0 and 85. Only half of women give birth, but each gives birth to two daughters. The age at which a woman gives birth is variable: her age at her first daughter's birth is $25 + round(5N(0,1))$. The age of the second birth is the age of first birth, plus either 2, 3, or 4 years (the interval is selected at random). At birth, individuals are assigned random traits, with one exception: they take their linguistic habits from their mothers. If the mother has adopted a sound change, the new baby will as well.

Each individual is also assigned an "innovativeness" rating. This is a concept taken from diffusion of innovation research. An innovativeness score of 0 means than an individual will never change.[1] A score of 1 means that an individual has no resistance to change. Innovativeness scores I were assigned randomly with the formula $0.5 + 8N(0,1)$, which was intended to create a low value of 0, and high value of 1, and a peak at 0.5 (with a normal distribution).

[1] "Propensity to change" might be a more intuitive term for this concept, since it has nothing to do with initiating an innovation; but Rogers (1962) uses "innovativeness."

Any values that fell outside of this range were changed to the closer of 0 or 1.

Time steps of the simulation are given the interpretation of years. Each year, everyone ages, women who are of the appropriate age give birth, women aged 85 die, and everyone considers adopting the sound change. The final step is of course the one critical to the applicability of the model to sound change.

Three factors influence the probability of an individual's adopting the sound change. First, there is an age restriction, in that only those aged 4 to 16 are able to change (variants on this restriction are discussed below); we can define this as a variable $r \in \{0, 1\}$, where r is 1 if the individual is in this age range, and 0 otherwise. The probability of changing is also proportional to the prevalence of the change in the community, and to the individual's innovativeness. Therefore, the probability of adopting the sound change is given by the equation:

$$P_{\text{adoption}} = r \times p \times I \tag{4}$$

Figure 2 shows a typical simulation in which all three factors are taken into account. The significant observations to be made are: (1) the progression of sound change is sigmoidal, and (2) the change lasts ≈ 100 years, an empirically plausible time scale. Significantly, the behavior of the population is longer than the behavior of an individual.

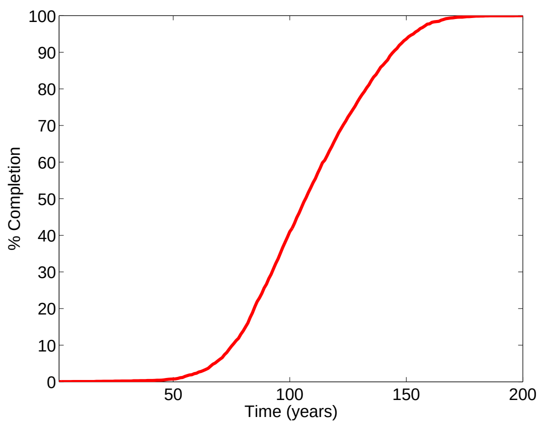

Figure 2: Prevalence of the change across time as individuals change as a function of prevalence and innovativeness, with an age restriction.

It turns out that the moderate rate of change is a direct result of the age restriction. Figure 3 shows a typical simulation without an age restriction on when individuals can change. In this case change occurs (approximately) within a single generation.

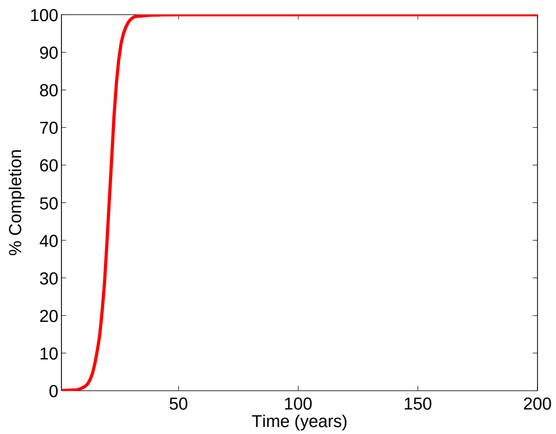

Figure 3: Identical model parameters as Figure 2, but without an age restriction.

The model with an age restriction provides a better match to the empirical data. Nevertheless, research *does* indicate that adults are capable of changing their pronunciations, albeit to a lesser degree than do children (e.g. Yaeger-Dror 1994, 1996). Two piecewise equations are presented below that provide for alternative age restrictions. These alternate definitions of r plug right into equation (4), since their values are in $[0,1]$. In (5), the ability to change tapers off linearly after age 16; in (6) the drop-off is exponential (this creates more rapid drop-off).

$$r_{\text{soft}} = \begin{cases} 0 & A \leq 3 \\ 1 & 4 \leq A \leq 16 \\ 1 - A/68 & A > 16 \end{cases} \tag{5}$$

$$r_{\text{hard}} = \begin{cases} 0 & A \leq 3 \\ 1 & 4 \leq A \leq 16 \\ \exp(-4A/68) & A > 16 \end{cases} \tag{6}$$

The effects of these different age restrictions are shown respectively in Figures 4 and 5. These graphs show intermediate rates of change, indicating that it is the overall ability of the population to change that is the critical factor in modulating the rate of language change.

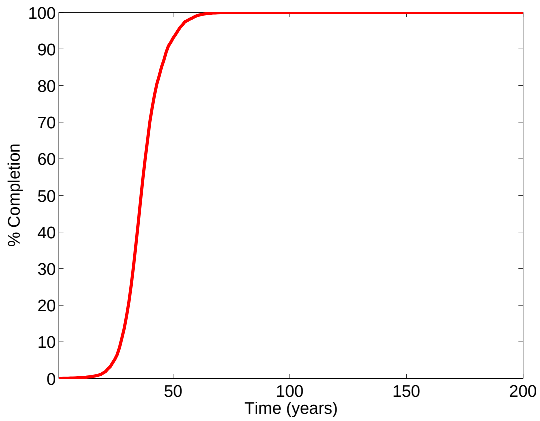

Figure 4: Result obtained with r_{soft}.

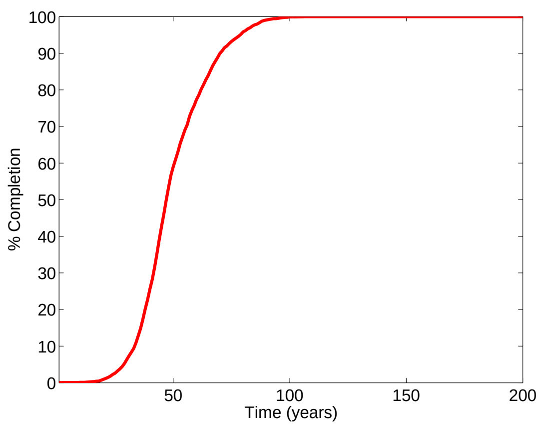

Figure 5: Result obtained with r_{hard}.

We next examine the role of the other two factors that affect an individual's probability of changing. Figure 6 shows a typical simulation, with an age restriction, where individuals consider only innovativeness. This is clearly an inadequate model, since the progression of the change is not sigmoidal but linear. Figure 7 likewise shows the result where only the prevalence of the change is considered. The trajectory of the change in this graph is not quite sigmoidal, but it is close. It is unknown whether this graph is more or less consistent with empirical data than one which included innovativeness as a factor.

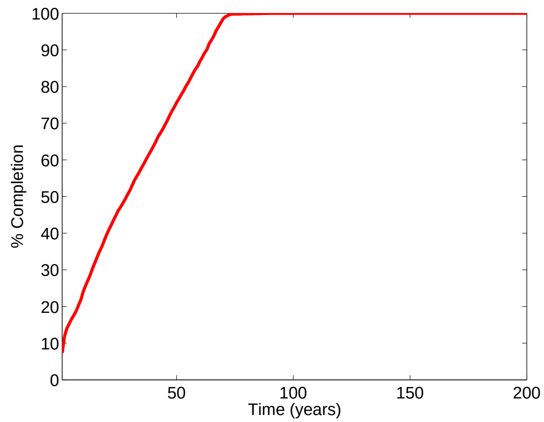

Figure 6: Result where individuals' probability of changing is proportional only to innovativeness (only those aged 4-16 can change).

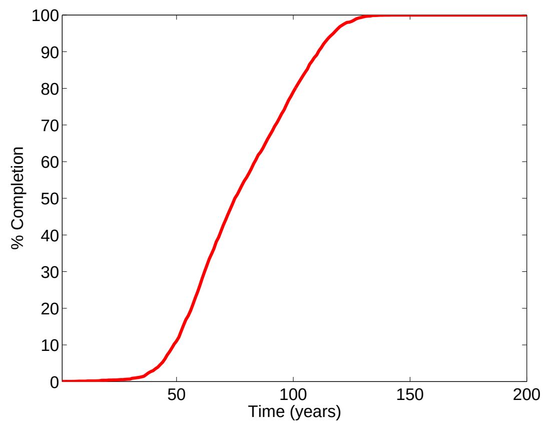

Figure 7: Result where individuals' probability of changing is proportional only to prevalence (only those aged 4-16 can change).

One final empirical check of the model is presented before moving on to the actuation problem. One of the empirical results of Labov (2001)'s linear model of sound change is that it replicates the finding that children and adolescents tend to be less advanced in sound change than young adults. The reason for this is that children initially adopt the linguistic habits of their mothers. If a child's mother is 26 years old at the time of her birth, then by the time the child enters day school, she will have the speech of a 30-year-old. It requires time in a social environment (e.g. between ages 4 and 16) for the child to advance to the level of young adults. Figure 8 shows the prevalence of the sound change by age, at four time steps in the simulation. It is apparent that for the first (lower) lines, there is a peak at $\approx$ 17 years (subsequent time increments have a ceiling effect, concealing the peak). This matches the empirical data, as well as the Labov (2001) model.

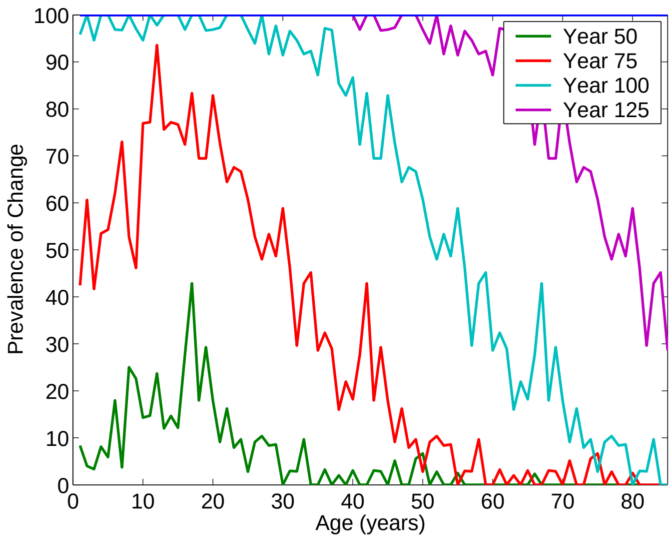

Figure 8: Prevalence of change by age, at four time steps.

4 Looking at Actuation through a Sigmoidal Lens

The sigmoidal progression of linguistic change has strong implications for the actuation of change. The graphs presented thus far have come from simulations that involve 1,000 or 3,000 individuals, and in each of the simulations, only one speaker was selected to have the sound change at the outset. Nevertheless, without fail, this one speaker's linguistic habit propagated through the entire speech community. It is emphasized that this sensitivity to initial conditions is not a mere artifact of a particular model, but is rather a property of the sigmoidal progression of sound change in general. All sound changes start from zero prevalence and, as it were, beat the odds to succeed.

Do models that replicate sigmoidal growth offer any advantage over POC models, when addressing the actuation problem? They do, but indirectly. The extreme sensitivity of sigmoidal growth to initial conditions indicates where the solution to the actuation problem must lie: in the linguistic leaders, the very earliest initiators of linguistic change. At no other time interval is there a chance to affect the outcome.[2]

Regarding the imitation of sound change, two particularly salient features of linguistic leaders are these, drawn from Labov (2001): linguistic leaders are status-conscious, and they have many social connections. We may safely assume that these individuals are keeping track of what they hear, keeping track of the social status of the people they hear it from, and checking for connections between the two. These skills are, after all, necessary to attain minimal competence as a member of a speech community, i.e. to speak in a way that befits one's (possibly idealized) social status.

We then have a small group of people constantly evaluating others' utterances and monitoring for significant correlations between social status and phonetic variables. One of the elementary results of statistics is that a correlation need not be present in a set of data for one to be observed. Occasionally, the social and phonetic variables will fall into a correlation, not because the correlation is part of the social structure of a speech community, but because of sampling error. If a linguistic leader detects such a trend, it seems reasonable to suppose that she will adopt the "change" herself. At that point, of course, the mechanisms outlined above will propagate the change throughout the speech community. An initial mistaken conclusion becomes a self-fulfilling prophecy.

The frequency with which sound change occurs is therefore a direct consequence of the frequency with which such incidental correlations arise, and that linguistic leaders act upon them. Just how often this is expected to happen

[2]Note that this is for a very, very simple social situation. In certain real world situations it is easy to think of explanations for why particular linguistic changes would not "take off." For example, [s] $>$ [θ] seems unlikely to occur in English, since this pronunciation is associated with childish or disordered speech. Lexical items like "loo" and "lift" (instead of *bathroom* and *elevator*) seem unlikely to become more prevalent in America because of their widely-known association with British speech. *Librarian* $>$ *libarian* is disfavored by many people for orthographic reasons. The list could go on for a long time, but these kinds of examples do not in themselves constitute a general solution to the actuation problem.

is unknown, and is a major weakness in the current proposal. In effect, the question is: what is a linguistic leader's p-value? At what point will she reject the null hypotheses and identify a spurious correlation as genuine? Mathematical answers to this question are possible, since the distribution of correlation coefficients of random vectors is well understood. But, absent a wealth of data concerning internal psychological processes, it seems unlikely that any such determination of parameters would be anything but stipulative and therefore uninsightful. Moreover, it is simple to imagine plausible scenarios that would add orders of complexity to the calculations. Suppose a leader thinks she has identified a significant linguistic variable, but then decides to listen again instead of adopting the change immediately. Naturally, the odds of confirming a spurious correlation through subsequent observation are much decreased, which would then lead to the expectation of even less frequent sound change.

However the question is ultimately addressed, it seems inevitable that the answer will be complex. If for no other reason, different linguistic changes seem to occur at different rates. New lexical items are propagated through a speech community within the space of years (at most), while sound changes and syntactic changes can stretch across centuries (Kroch 1989, Labov 2001).

Until the question of how often such misperceptions would arise can be examined empirically, a tentative (and not very satisfying) answer may well have to be: not very often.

5 Conclusions

This paper has examined the actuation of sound change, using computational models of speech communities as a guide. Phonologization of coarticulation models make counterfactual predictions about the prevalence of sound change, and should be rejected for that reason. Computation simulations of this Neogrammarian idea confirm the critique of Weinreich et al. (1968). A different kind of model, based on the work of Labov (2001) and Rogers (1962), replicates the sigmoidal trajectory of language change. An emphasis on such models—or on sigmoidal curves in general—suggests where the solution to the actuation problem must lie: with the very earliest initiators of sound change. It is proposed that linguistic leaders occasionally misidentify linguistic variables, attributing to significance to an incidental correlation of phonetic and social variables. If such a leader adopts the change into her own speech, it would initiate a sequence of sound change that would eventually affect the entire community.

Much work remains in developing these ideas. The major hole in the proposal is the lack of a rigorous way to determine how likely people are to misidentify sound changes. Many other inadequacies of the model could be identified as well: its inability to handle phonetically gradual change (at least in the current implementation), its reliance on a certain interpretation of the time step, as well as various inadequacies of the population model. Nevertheless, it is significant that even simple models of speech communities can be used to discredit inadequate theories of language change, and, for models with more credible empirical results, to focus attention on those stages of change that are crucial to its initiation.

References

Bailey, Charles-James N. 1973. *Variation and Linguistic Theory*. Washington, D.C.: Center for Applied Linguistics.

Kroch, Anthony S. 1989. Function and grammar in the history of English: Periphrastic 'do'. In *Language Variation and Change: Current Issues in Linguistic Theory. vol. 52.*, ed. Fasold and Schiffrin, 133–172. Philadelphia: John Benjamins.

Labov, William. 1994. *Principles of Linguistic Change, Volume 1: Internal Factors*. Blackwell.

Labov, William. 2001. *Principles of Linguistic Change, Volume 2: Social Factors*. Blackwell.

Ohala, J.J. 1983. The origin of sound patterns in vocal tract constraints. In *The Production of Speech*, ed. Peter F. MacNeilage, 189–216. New York: Springer-Verlag.

Paul, Hermann. 1880. *Prinzipien de Sprachgeschichte*. Halle: Niemeyer.

Pierrehumbert, Janet B. 2001. Exemplar dynamics: Word frequency, lenition and contrast. In *Frequency and the Emergence of Linguistic Structure*, ed. Joan Bybee and Paul Hopper. Amsterdam: John Benjamins.

Rogers, Everett M. 1962. *The Diffusion of Innovations*. Free Press.

Weinreich, Uriel, William Labov, and Marvin Herzog. 1968. Empirical foundations for a theory of language change. In *Directions for Historical Linguistics*, ed. W. Lehmann and Y. Malkiel. Austin: U. of Texas Press.

Whitney, William Dwight. 1867/1896. *Language and the Study of Language*. New York: Charles Scribners Sons.

Yaeger-Dror, Malcah. 1994. Sound change in Montreal French. In *Phonological Structure and Phonetic Form. Papers in Laboratory Phonology 3*, ed. P. Keating, 267–292. Cambridge University Press.

Yaeger-Dror, Malcah. 1996. Phonetic evidence for the evolution of lexical classes: The case of a Montreal French vowel shift. In *Towards a Social Science of Language*, ed. G. Guy, C. Feagin, J. Baugh, and D. Schiffrin, 263–287. Benjamins: Philadelphia.

Department of Linguistics
University of Arizona
Tucson, AZ 85721
tabaker@u.arizona.edu