

Part II

Phonetic considerations

OUNCORRECTED PROOF – REVISES, 1/9/2012, SPi

3

Phonetic bias in sound change

ANDREW GARRETT AND KEITH JOHNSON

3.1 Introduction

Interest in the phonetics of sound change is as old as scientific linguistics (Osthoff and Brugman 1878).¹ The prevalent view is that a key component of sound change is what Hyman (1976) dubbed *PHONOLOGIZATION*: the process or processes by which automatic phonetic patterns give rise to a language's phonological patterns. Sound patterns have a variety of other sources, including analogical change, but we focus here on their phonetic grounding.² In the study of phonologization and sound change, the three long-standing questions in (1) are especially important.

- (1) a. Typology: Why are some sound changes common while others are rare or nonexistent?
- b. Conditioning: What role do lexical and morphological factors play in sound change?
- c. Actuation: What triggers a particular sound change at a particular time and place?

In this chapter we will address the typology and actuation questions in some detail; the conditioning question, though significant and controversial, will be discussed only briefly (in section 3.5.3).

¹ For helpful discussion we thank audiences at UC Berkeley and UC Davis, and seminar students in 2008 (Garrett) and 2010 (Johnson). We are also very grateful to Juliette Blevins, Joan Bybee, Larry Hyman, John Ohala, and Alan Yu, whose detailed comments on an earlier version of this chapter have saved us from many errors and injudicious choices, though we know they will not all agree with what they find here.

² Types of analogical change that yield new sound patterns include morphophonemic analogy (Moulton 1960, 1967) and analogical morphophonology (Garrett and Blevins 2009). Of course, the source of a pattern is not always clear. For example, patterns like the linking [ɹ] of many English dialects have been attributed to a type of analogical change called 'rule inversion' (Vennemann 1972a), perhaps not phonetically grounded, but work by Hay and Sudbury (2005) and others calls this into question. Note that some phonological patterns, while phonetically grounded in a broader sense, correspond to no specific phonetic patterns because they arise through the telescoping of multiple phonetically grounded sound changes. Again, it is not always easy to identify such cases confidently.

The typology question concerns patterns like those in (2–3). In each pair of examples in (2), one is a common sound change while the other is nonexistent. The ultimate causes of these patterns are clear enough where there are obvious phonetic correlates, but the mechanisms explaining the relationship—that is, the precise mechanisms of phonologization—are still disputed.

- (2) Typologically common vs. nonexistent sound changes
- a. Common: [k] > [tʃ] before front vowels (Guion 1998)
Nonexistent: [k] > [q] before front vowels
 - b. Common: vowel harmony involving rounding (Kaun 2004)
Nonexistent: vowel harmony involving length
 - c. Common: vowel reduction restricted to unstressed syllables (Barnes 2006)
Nonexistent: vowel reduction restricted to stressed syllables
 - d. Common: consonant metathesis involving sibilants (Blevins and Garrett 2004)
Nonexistent: consonant metathesis involving fricatives generally

Our typological point can be sharpened further. Not only are there generalizations about patterns of sound change, but the typology is overwhelmingly asymmetric. For example, the inverse of each of the common changes in (3) is nonexistent.

- (3) Asymmetries in sound change
- a. Common: [k] > [tʃ] before front vowels
Nonexistent: [tʃ] > [k] before front vowels
 - b. Common: intervocalic stop voicing (Kirchner 2001; Lavoie 2001)
Nonexistent: intervocalic stop devoicing
 - c. Common: [t] > [ʔ] word-finally (Blevins 2004: 120–1)
Nonexistent: [ʔ] > [t]

It is uncontroversial that such asymmetries in sound change must (somehow) reflect asymmetries in phonetic patterns. We will refer to these as *BIASES*.

Our approach to the typology question, then, is grounded in processes of speech production and perception and in the phonetic knowledge of language users. The bulk of our chapter is devoted to an evaluation of various components of speech production and perception, with an eye to identifying asymmetries (biases) that should be associated with each component. Our hypothesis is that various types of sound change can be grounded in the various speech components based on their typological profiles. We hope this approach yields a useful framework for discussing the relation between patterns of sound change and their phonetic correlates.

From a broader perspective the typology question can be seen as a facet of what Weinreich et al. (1968) call the *CONSTRAINTS PROBLEM*: determining ‘the set of

3. *Phonetic bias in sound change* 53

possible changes and possible conditions for change' (p. 183). The second main question we address in this chapter is what they call the **ACTUATION PROBLEM**: Why does a change take place in one language where its preconditions are present, but not in another? Historical linguists sometimes defer this question to sociolinguists by assuming that its answer involves contingencies of social interaction, but a comprehensive model of phonologization should explain how phonetic patterns uniformly characterizing all speakers of a language can give rise to phonological patterns that serve as speech variants or norms for some of them.

Our approach highlights the three elements of phonologization shown in (4).

- (4) a. Structured variation: Speech production and perception generate variants (see sections 3.3–3.4)
- b. Constrained selection: Linguistic factors influence the choice of variants (see section 3.5)
- c. Innovation: Individuals initiate and propagate changes (see section 3.6)

Processes of speech production and perception generate what Ohala (1989) memorably describes as a 'pool of variation' from which new phonological patterns emerge; we emphasize that this variation is structured in ways that help determine phonological typology. Other processes contribute to the phonologized outcome; for example, Kiparsky (1995) and Lindblom et al. (1995) refer to 'selection' from the pool of variants. But our first goal is to understand how the underlying variation itself is structured by bias factors, even if selectional processes also contribute bias (see section 3.5). Finally, actuation begins with innovation; our second goal is to understand why individual innovators would increase their use of certain speech variants from the pool of variation.

This chapter is organized as follows. In sections 3.2–3.5, we address the constraints problem of Weinreich et al. (1968). We begin with a review of sound change typologies in section 3.2; despite differences of detail, many share a taxonomy inherited from the neogrammarians. In section 3.3, we examine elements of speech production and perception and evaluate possible bias factors in each case; we suggest in section 3.4 that certain patterns of sound change may be correlated with certain bias factors based on their phonological typology. We discuss selection in section 3.5, describing facets of phonologization that are system-dependent; they may involve bias factors, but only relative to language-specific or universal systematic constraints.

In section 3.6 we turn to the actuation question, sketching a theory of mechanisms that link bias factors and sound changes. While the former are often universal, the latter are language-specific and at first perhaps even speaker-specific. Successful changes must propagate from innovators before eventually becoming community speech norms; we present the results of simulating aspects of this process. We conclude in section 3.7 with a brief summary and some questions for future research.

3.2 Typologies of sound change

Historical linguistics textbooks (e.g. Hock 1991, Hock and Joseph 1996, Campbell 2004, Crowley and Bowerman 2009) classify sound changes according to a superficial typology, often naming very specific categories: apocope, cluster simplification, metathesis, palatalization, umlaut, etc. Of course it is important for students to learn what these terms mean. But more sophisticated work has always recognized that an explanatory classification of surface patterns should reflect a typology of causes. Two typologies have been especially influential within historical linguistics: a traditional two-way division into articulatorily-grounded and other sound changes, and a newer three-way division into listener-oriented categories; see Tables 3.1–3.2. We will briefly describe each approach, as well as Grammont’s (1939) more elaborated scheme.

The traditional typology is due to the neogrammarians. According to this account, most types of sound change originate through processes of articulatory reduction, simplification, or variability; dissimilation, metathesis, and a few other types comprise a residual type with other origins. Osthoff and Brugman (1878) themselves only briefly comment, indicating that most changes have ‘mechanical’ (i.e. articulatory) causes while dissimilation and metathesis are ‘psychological’ in origin. It was Paul (1880, 1920) who suggested specifically that the first type originates in articulatory reduction, speculating as well that the second may have its basis in speech errors. Crucially, in any case, the neogrammarians and Bloomfield (1933) held that the major type of sound change was phonetically gradual, imperceptible while under way, and

TABLE 3.1 Several influential traditional typologies of sound change

AUTHOR	ORIGIN OF MOST SOUND CHANGES	RESIDUAL TYPE
Osthoff and Brugman (1878)	‘mechanical’ (articulatory)	ORIGIN: ‘psychological’ EXAMPLES: dissimilation; metathesis
Paul (1880, 1920)	articulatory reduction	ORIGIN: speech errors? EXAMPLES: metathesis; non-local assimilation and dissimilation
Bloomfield (1933)	articulatory simplification?	ORIGIN: unclear EXAMPLES: articulatory leaps; dissimilation; haplogogy; metathesis; non-local assimilation
Kiparsky (1995)	variation in production	ORIGIN: ‘perception and acquisition’ EXAMPLES: compensatory lengthening; dissimilation; tonogenesis; context-free reinterpretation, e.g. [k ^w] > [p]

3. *Phonetic bias in sound change* 55**TABLE 3.2** Two recent listener-based typologies of sound change

OHALA (1981; 1993 <i>b</i>)	BLEVINS (2004; 2006 <i>a</i> ; 2008 <i>b</i>)
LABEL: Hypocorrection EXAMPLES: umlaut; many other assimilations	LABEL: ‘CHOICE’ EXAMPLES: vowel reduction and syncope; vowel shifts; stop debuccalization; final devoicing; umlaut; etc.
LABEL: Hypercorrection EXAMPLE: dissimilation	LABEL: ‘CHANCE’ EXAMPLES: dissimilation; metathesis
LABEL: Confusion of acoustically similar sounds EXAMPLES: [θ] > [f]; [gi] > [di]	LABEL: ‘CHANGE’ EXAMPLES: [θ] > [f]; [anpa] > [ampa]; [akta] > [atta]

regular.³ This theory was couched by Paul (1880, 1920) in a surprisingly modern exemplar-based view of phonological knowledge (see section 3.6 below).

More recently, a similar two-way scheme has been defended by Kiparsky (1995). He writes that the first sound change type originates as speech variation with articulatory causes; certain variants are then selected by linguistic systems, subject to further (linguistic) constraints.⁴ The residual type consists of changes that originate as perceptually-based reinterpretations, possibly in the course of language acquisition.

The role of the listener was already crucial for Paul (1880, 1920), according to whom the major type of sound change occurs when articulatory processes create variants that are heard by listeners, stored in exemplar memory, and in turn give rise to new, slightly altered articulatory targets. But in emphasizing the articulatory basis of sound change, neither the neogrammarians nor their successors explored the possible details of listener-based innovation. In recent decades, two influential accounts of sound change have done precisely this. These accounts, due to John Ohala and Juliette Blevins, share comparable three-way typologies. We highlight the similarities between them in Table 3.2, though they also have important differences.

For Ohala, most explicitly in a 1993 paper (Ohala 1993*b*), there are three main mechanisms of sound change.⁵ The one corresponding most closely to the traditional category of articulatorily grounded change is what he calls *HYPOCORRECTION*. This is rooted in correction, the normalization that listeners impose on a signal—for example, factoring out coarticulatory effects to recover a talker’s intention. In *hypocorrection*, a listener undercorrects for some coarticulatory effect, assuming that it is

³ Bloomfield (1933) suggests with some uncertainty that articulatory simplification may underlie the major type of sound change; he expresses no view of the cause(s) of the residual type.

⁴ The role of articulatory reduction in sound change has also been emphasized by other modern linguists (e.g. Mowrey and Pagliuca 1995, Bybee 2001, 2007), but they have not yet presented an overall account of how various types of sound change fit together.

⁵ It is hard to select one or even a few of Ohala’s contributions from within his influential and insightful *oeuvre* in this area; see linguistics.berkeley.edu/phonlab/users/ohala/index3.html for a full list.

phonologically intended; this leads to the phonologization of coarticulatory patterns. (One of the most important features of this account is that it explains why articulatorily driven changes are not even more widespread: through correction, articulatorily motivated variants are usually reinterpreted as intended.) A second mechanism is called *HYPERCORRECTION*: a listener overcorrects, assuming that a phonologically intended effect is coarticulatory; this leads to a dissimilatory change. Ohala's third mechanism of sound change is the confusion of acoustically similar sounds, which he attributes to the listener's failure to recover some feature found crucially in one sound but not the other.⁶

Most recently, Blevins (2004, 2006a, 2008b) uses the terms *CHOICE*, *CHANCE*, and *CHANGE* for what she views as the three basic mechanisms of sound change. In principle they are distinct from Ohala's mechanisms; extensionally they are similar. For example, *CHOICE* refers to innovations grounded in articulatory variation along the hypospeech–hyperspeech continuum, for which (Blevins 2006a:126) assumes 'multiple phonetic variants of a single phonological form'. Mostly these correspond to the major type of change recognized by the neogrammarians, and to cases of what Ohala treats as hypocorrection, though he does not refer to a continuum of phonetic variants from which hypocorrection operates.

Blevins's term *CHANCE* refers to innovations based on intrinsic phonological ambiguity. For example, a phonological sequence /aʔ/ might be realized phonetically as [a̠], permitting listeners to interpret it phonologically either as (intended) /aʔ/ or as /ʔa/; if /ʔa/ is chosen, a metathesis sound change has occurred. Dissimilatory changes described by Ohala as hypercorrection are understood as a special case of *CHANCE*. Finally, the term *CHANGE* refers to innovations in which some perceptual bias leads to misperception. For example, in an /anpa/ > /ampa/ assimilation, it is hypothesized that the speaker crucially did not produce [mp]; rather, a listener perceived [anpa] as [ampa] and interpreted it phonologically as /ampa/. Other examples of this type include context-free place of articulation shifts like [θ] > [f], also mentioned by Ohala as the parade example of confusion of acoustically similar sounds.⁷

⁶ Ohala (1993b: 258) suggests that this can be viewed as a type of hypocorrection; the difference 'is whether the disambiguating cues that could have been used by the listener (but were not) are temporally co-terminous with the ambiguous part [as in the confusion of acoustically similar sounds] or whether they are not', as in hypocorrection.

⁷ On this sound change see section 3.5.1 below. A potential criticism is that of Blevins's three mechanisms, only *CHANGE* is intrinsically asymmetric (assuming that perceptual biases and constraints on misperception are asymmetric). By contrast, nothing about *CHOICE* or *CHANCE* *per se* predicts any directionality; for example, Blevins (2004:35) notes, in *CHANCE* 'there is no language-independent phonetic bias' and 'the signal is inherently ambiguous'. Therefore the explanation for any observed asymmetries must be sought elsewhere. This criticism is not germane to Ohala's system. In that system, however, since hypocorrection and hypercorrection are mirror-image processes, there is no immediate explanation for their many asymmetries (for example, nonlocal laryngeal-feature dissimilation is common but nonlocal laryngeal-feature assimilation is rare).

3. *Phonetic bias in sound change* 57

Perhaps the fullest typology is that of Grammont (1939), the first author to present a theory based on a survey of all known sound change patterns.⁸ For him, sound changes emerge through competition between constraints—he called them ‘laws’ (Grammont 1939:176)—favoring effort reduction and clarity, as well as other factors. Given in (5) is his full scheme; he distinguishes changes where the conditioning environment is adjacent or local (5b) from those where it is nonlocal (5c). Grammont’s typology cannot readily be adapted to the present day, but it is notable that he invoked articulatory reduction, perceptual clarity, and motor planning as key ingredients in sound change. His theory of nonlocal dissimilation is especially interesting (see already Grammont 1895): he argues that the segment undergoing dissimilation is always in a ‘weaker’ position than the trigger; positional strength is defined with reference to accent, syllable position, and, if all else is equal, linear order, in which case the first segment is weaker. He suggests that nonlocal dissimilation occurs when planning for a segment which is in a more prominent position distracts a talker who is producing a similar segment in a weaker position.

- (5) Grammont’s (1939) typology of sound changes
- a. Unconditioned changes: explanation unclear (in some cases language contact?)
 - b. Locally conditioned changes
 - ASSIMILATION: motivated by articulatory ease
 - DISSIMILATION: motivated by perceptual clarity
 - METATHESIS: motivated by perceptual clarity and phonotactic optimization
 - c. Nonlocally conditioned changes
 - ASSIMILATION: explanation unclear, but evidently articulatory in origin
 - DISSIMILATION: originates in motor-planning errors
 - METATHESIS: motivated by perceptual clarity and phonotactic optimization

Our own presentation draws much from the approaches of earlier authors, but it crucially differs from them. With its reference to ‘articulatory’ reduction and variability, the traditional dichotomy inherited from the neogrammarians is too simplistic, even in its modern avatars, and fails to reflect the true complexity of speech production. On the other hand, the listener-oriented typologies of Ohala and Blevins leave

⁸ That is, all patterns known to him over 75 years ago. The only comparable works are by Hock (1991), whose textbook classifies surface patterns without a theory of causes, Blevins (2004), whose broad coverage is exhaustive for certain patterns but is not meant to be complete for all types of sound change, and Kümmel (2007), whose coverage is restricted to a few language families. Today it would be almost impossible to be as thorough as Grammont tried to be; useful modern sources are Blevins’s (2008b) ‘field guide’ and Hansson’s (2008) overview.

essential questions about speech production unanswered; for example, what processes generate and constrain the variable input to Blevins's CHOICE? Finally, while thorough and replete with interesting observations, Grammont's account is too inexplicit and stipulative to be used without change today.

The typology we present is deductive rather than inductive. That is, rather than surveying sound changes, we examine components of speech production and perception, seeking relatively complete coverage, and we ask what biases each component is likely to yield. We propose that biases emerging from the various systems of speech production and perception, respectively, underlie various types of sound change with corresponding phonological profiles. What emerges from this approach has elements of previous typologies, therefore, but cannot be directly mapped onto any of them.

3.3 Biases in speech production and perception

There are several sources of variability in the speech communication process that may lead to sound change. For example, a listener may misperceive what the talker says because the talker is speaking softly or at a distance, or there is some background noise. Similarly, the talker may misspeak, accidentally producing a different sound from that intended or a variant of the sound that is different from usual. Further, children may come to language acquisition with bias to organize linguistic knowledge in ways that turn out to differ from the organization used by their parents.

Variability introduced in these ways by the communication process could be random. For example, misperception of a vowel would result in hearing any other vowel in the language with equal probability. However, most sources of variability are far from random, and instead introduce bias into the process of sound change so that some outcomes are more likely than the others.

For example, when the English lax vowel [ɪ] is misperceived (Peterson and Barney 1952), not all of the other vowels of English are equally likely to be heard. Instead, as Table 3.3 shows, the misperception is likelier to be [ɛ] than any other vowel. This lack of randomness in perceptual variation is one property of bias factors in sound change.

TABLE 3.3 Identification of [ɪ] and [ɛ] in Peterson and Barney (1952). The speaker's intended vowel is shown in the row label, and the listener's perceived vowel in the column label.

	i	ɪ	ɛ	æ	ɑ	er	other
[ɪ]	0.06	92.9	6.75	0.02	0.01	0.25	0.01
[ɛ]	0	2.5	87.71	9.23	0.01	0.5	0.05

TABLE 3.4 Identification of [ʊ] and [ʌ] in Peterson and Barney (1952). The speaker's intended vowel is shown in the row label, and the listener's perceived vowel in the column label.

	u	ʊ	ʌ	ɔ	ɑ	other
[ʊ]	0.93	96.55	1.66	0.5	0.16	0.2
[ʌ]	0	1	92.21	1.24	5.25	0.3

A second (defining) property of bias factors in sound change is that bias is directional. For example, given that [ɪ] is most often misperceived as [ɛ], one might suppose that [ɛ] would reciprocally be misperceived as [ɪ]. As the second line in Table 3.3 shows, this is not the case. Although [ɛ] is misperceived as [ɪ] at a rate that is greater than chance, the most common misperception of [ɛ] was as [æ]. Table 3.4 indicates that the lax back vowels tested by Peterson and Barney showed a similar asymmetric confusion pattern, where [ʊ] was confused with [ʌ] while [ʌ] was more often confused with [ɑ]. Labov (1994) observed that in vowel shifts, lax vowels tend to fall in the vowel space; the perceptual data in Tables 3.3–3.4 suggest that one source of the directionality of the sound change may be a perceptual asymmetry. In any case, our main point is that phonetic bias factors are directional.

Phonetic bias factors thus produce a pool of synchronic phonetic variation (Ohala 1989, Kiparsky 1995, Lindblom et al. 1995) which forms the input to sound change; this is sketched in Figure 3.1. The structure imposed on the phonetic input to sound change, via the directionality of phonetic variation, is a key source of the typological patterns of sound change.

In the following subsections, we will consider potential bias factors arising from the phonetics of speaking and listening, and the extent to which they may provide both non-randomness and directionality in sound change. Speaking and listening as a whole can be said to contain four elements that might provide bias factors in sound change. We will discuss these in turn: motor planning (3.3.1); aerodynamic constraints (3.3.2); gestural mechanics (3.3.3), including gestural overlap and gestural blend; and perceptual parsing (3.3.4). The order roughly mimics the order from thought to speech, and from a talker to a listener. In Section 3.4 we will turn to discuss representative types of sound change that stem from the various bias factors we identify in this section.

3.3.1 *Motor planning*

Motor planning is the process of constructing or retrieving motor plans that will later be executed by speaking. In this process, speech errors may occur as planning elements (syllables, segments, gestures, etc.) influence each other through priming or

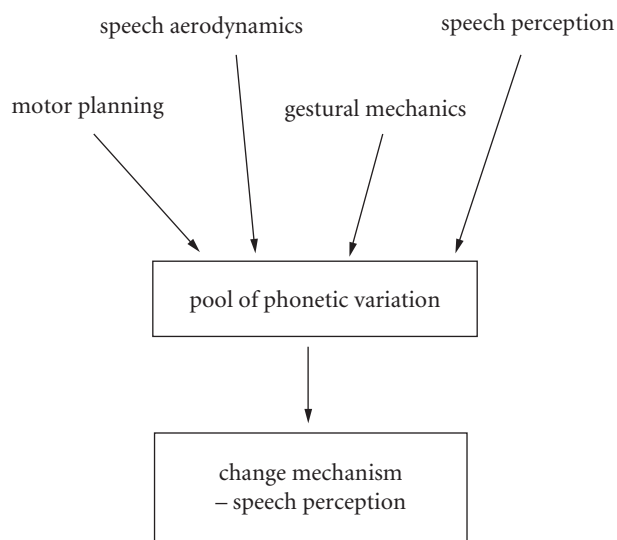
Phonetic bias factors:

FIGURE 3.1 Phonetic bias factors produce a pool of synchronic phonetic variation which can be taken up in sound change.

coactivation, or through the inhibition of one segment by the activation of another. Sound changes may then emerge if such speech errors are incorporated into a language's phonology. Two basic speech error patterns could lead to sound change. The first of these, *BLENDING*, has been extensively studied in the speech error literature. The second, *INHIBITION*, appears to be much less common, though it is a focus of language play.⁹

In motor plan blending, plans for nearby similar segments may influence each other as they are activated; this is motor *PRIMING* (see Tilsen 2009a with further references). The blending or interaction of similar, nearby sounds is exemplified in interchange errors (*snow flurries* → *flow snurries*), anticipations (*reading list* → *leading list*), and

⁹ On motor plan blending see Boomer and Laver (1968), MacKay (1970), Fromkin (1971), Fromkin (1973), Dell (1986), and Shattuck-Hufnagel (1987). Pouplier and Goldstein (2010) have also shown that speech planning and articulatory dynamics interact with each other in complex ways, so that the specific phonetic results of some speech errors may be outside the speaker's ordinary inventory of articulatory routines.

In addition to blending and inhibition, bias may also emerge from what Hume (2004b) calls *ATTES-TATION*, suggesting that some metathesis patterns point to 'a bias towards more practiced articulatory routines' (p. 229). Undoubtedly, there is a tendency for the articulators to be drawn to familiar, routinized patterns. This can be seen in loan word adaptation as words are nativized, and probably also exerts a type of phonotactic leveling as Hume suggests. We consider attestation to be a systematic constraint (section 3.5), different in kind from the phonetic bias factors, though in this case the difference between linguistically universal and language-specific biases is particularly fine.

3. *Phonetic bias in sound change* 61

preservations (*waking rabbits* → *waking wabbits*). Blending of segmental plans due to adjacency of those plans results in bias toward non-randomness in speech production errors. People are more likely to blend plans that are in proximity to each other—in time, phonetic similarity, and articulatory planning structure (that is, onsets interact with onsets, nuclei with nuclei, etc.).

The effects of motor plan inhibition can be seen in tongue twisters where an alternating pattern is interrupted by a repetition (Goldinger 1989). In the sequence *unique New York* we have a sequence of onset consonants [j . . . n . . . n . . . j] and when the phrase is repeated the sequence is thus [. . . j n j j n n j j n n j j . . .], an *aa bb* pattern. Other tongue twisters are like this as well. For example, *she sells sea shells by the sea shore* is [ʃ . . . s . . . s . . . ʃ . . . s . . . ʃ]. Typically, in these sequences the error is toward an alternating pattern [j . . . n . . . j . . . n] instead of the repetition of one of the onsets. It may be worth noting in this context that repeated tongue motion is dispreferred in playing a brass instrument like a trombone or trumpet. With these instruments (and perhaps others) rapid articulation of notes is achieved by ‘double tonguing’—alternating between coronal and dorsal stops, rather than ‘single tonguing’—using a sequence of coronal stops to start notes.

In both motor plan blending and motor plan inhibition, it is likely that rhythm and stress may play a significant role in determining that prominent segments will be preserved while non-prominent segments will be altered, because the prosodic organization of language is extremely important in motor planning (Port 2003, Saltzman et al. 2008).

3.3.2 *Aerodynamic constraints*

Speech production is constrained by aerodynamics even in the absence of interactions among articulators. Aerodynamic bias factors are characterized by a tendency toward phonetic change as a result of changing aerodynamic parameters even when all else (e.g. the position of the articulators) remains constant. Two laws of speech aerodynamics are involved, among others. The first is the aerodynamic voicing constraint: in order to produce vocal fold vibration, air pressure below the glottis must be greater than air pressure above the glottis (Ohala 1983). This physical law governing voicing sets up an ‘ease of voicing’ hierarchy among the phonetic manner classes: stops are the hardest to voice, since air passing through the glottis will raise supraglottal air pressure in the closed vocal tract, and vowels are the easiest to voice. Thus, the aerodynamic voicing constraint introduces phonetic bias into sound change, biasing voiced stops to become voiceless.¹⁰

¹⁰ One might wonder if the aerodynamic voicing constraint biases vowels to be voiced. Despite the symmetry of it, we are reluctant to say so. It seems to us that voiced speech may have some inherent advantages for spoken communication. For example, voicing provides resistance to air flow so voiced breath-groups extend over a longer time than voiceless (e.g. whispered) breath groups. Voiced speech is also louder than voiceless speech, which is a communicative advantage in most situations. We see the

Linguists have noted a number of different linguistic responses to the phonetic bias against voiced stops. Voiced stops have lost their voicing and neutralized with voiceless stops, but in some languages maintenance of a contrast between voiced and voiceless stops is achieved by altering the phonetic properties of the voiced series in some way. Such ‘repair strategies’ include prenasalization, implosion, and spirantization. In our view, the phonetic bias imposed by the aerodynamic voicing constraint should impel voiced stops to become voiceless, all else being equal. The further development of repair strategies is motivated by contrast maintenance; see section 3.5.1 below on perceptual enhancement.¹¹

A second law of speech aerodynamics that provides phonetic bias in sound change is a constraint on frication, which can only be achieved when air pressure behind the fricative constriction is sufficient. This introduces a bias against voiced fricatives because airflow is impounded by the vocal folds, reducing oral pressure (Ohala 1983, Johnson 2003: 124, Ohala and Solé 2010). Without any articulatory adjustments, therefore, voiced fricatives will tend to become glides. As with the aerodynamic voicing constraint, the frication constraint introduces a repelling force—a bias against a particular combination of phonetic features—and a direction of change if no contrast-maintaining repair strategy is applied.

3.3.3 *Gestural mechanics*

The actual movements of articulators introduce variability in speech, and may introduce bias for sound change. Two types of interaction among articulators have been implicated in language sound patterns.

In the first type, GESTURAL OVERLAP, independent articulators like lips and tongue tip are moving at the same time and their movements may obscure each other. For example, in the utterance *hand grenade* the tip and body of the tongue are moving to make stop consonants in rapid succession in the [d] sequence at the word boundary. Cho (2001) found that the relative timing of gestures across a word boundary is more variable than is gestural timing within words, so in some productions of *hand grenade* the tongue body gesture for [g] may precede the tongue tip gesture for [d]. Though the [d] closure is made, it has very little impact on the acoustic output of the vocal tract, so the utterance sounds like [hæŋ grəneɪd]. The coronal gesture of [nd] is hidden by the dorsal gesture which now covers it. A hearer of such an overlapped utterance would think that the alveolar gesture has been deleted (Byrd 1994) and so may not include the hidden gesture in their plan for the word. In gestural overlap, the movement for

aerodynamic voicing constraint as a constraint against voicing—providing a phonetic bias toward the elimination of voicing in segments where voicing is difficult.

¹¹ On prenasalization and voicing see e.g. Iverson and Salmons (1996). It may be helpful to note also that ‘contrast maintenance’—a basic factor that we appeal to in accounting for sound change—is similar to ‘faithfulness’ constraints in Optimality Theory, whose ‘markedness’ constraints likewise correspond almost exactly to our phonetic bias factors.

3. *Phonetic bias in sound change* 63

a construction can be completely obscured by another. This mechanism introduces a directional bias in sound changes involving sequences such that back gestures are more likely to hide front gestures. Debuccalization is an example of this, where a glottalized coda may be replaced by glottal stop, but we very rarely see glottal stops become oral.

In the second type of interaction between articulators, *GESTURAL BLEND*, the phonetic plan for an utterance places competing demands upon a single articulator. For example, in the word *keep* the tongue body is required to move back toward the soft palate for the velar [k], and very soon later move forward for the front vowel [i]. Thus, in this word the location of the tongue during the [k] closure is farther forward in the mouth than it is during the [k] of words with back vowels like *cop*.

Several factors determine the outcome of a gestural blend. One of these comes from the quantal theory of speech production (Stevens 1989). Some gestures are more stable under perturbation, in the sense that the output of an acoustically stable gesture will not be much affected by blending with another gesture. In blending a quantally stable gesture with an unstable gesture, the more stable gesture will tend to determine the acoustics of the output. For instance, the more constricted gesture usually shows a greater acoustic change in gestural blending, while less constricted gestures are less impacted. In this way, even though Stevens and House (1963) and later researchers (Strange et al. 1976, Hillenbrand et al. 2001) found that vowels are significantly influenced by the consonants that surround them, the blending of tongue body gestures when a vowel follows velar /k/ or /g/ consonants results in a more noticeable change of the consonant gesture than of the vowel gesture; this yields fronted /k/ and /g/ adjacent to front vowels.

Because patterns of gestural interaction in blending and overlap are language-specific, languages develop different or even complementary patterns in phonologization.¹² For example, while Japanese has vowel devoicing in words like /kusuri/ → [kɯ̥sɯ̥ri] ‘medicine’ (Hasegawa 1999), other languages instead have intervocalic fricative voicing in similar contexts, as in northern Italian /kasa/ → [kaza] ‘house’ (Krämer 2009:213).

3.3.4 *Perceptual parsing*

The role of listeners and (mis)perception in sound change has been a major research theme in the three decades since Ohala’s ‘The listener as a source of sound change’ (1981). Some changes have been explained as a by-product of perceptual similarity: because two segment types sound similar, they are sometimes confused by listeners. All else being equal, if the likelihood of misperception is symmetrical, the resulting sound changes should be symmetrical; if X and Y are confusable sounds, X > Y should be just as likely as Y > X. If this were true of perceptual confusions generally we would

¹² See Bladon and Al-Bamerni (1976) for a discussion of language-specific coarticulation patterns, and Nolan (1985) on individual differences in coarticulation.

expect perceptual parsing to produce symmetric rather than asymmetric patterns of change. Simple perceptual confusability would then yield no bias factor favoring one direction of change over another.

As noted in section 5.1 above, however, sound change is typically asymmetric. For changes grounded in perceptual parsing, this would mean that listeners sometimes introduce bias and thus asymmetrical patterns of sound change. In principle this could happen in at least two ways, though more research is needed in both cases to determine the nature of the mechanisms. First, in some cases asymmetric misperception may be a bias factor. For instance, in Tables 3.3–3.4 we illustrated perceptual confusions among lax vowels. These reveal a distinct pattern of perceptual asymmetry in vowel perception, suggesting that the tendency for lax vowels to lower in the vowel space (Labov 1994) could have its phonetic roots in asymmetric misperception. Another such case is studied by Chang et al. (2001), who focussed on sounds that differ in the presence or absence of some acoustic element (e.g. a particular band of energy in a stop release burst). They suggest that sounds differing in this way may be asymmetrically misperceived—that listeners are more likely to fail to notice the element than to erroneously imagine it to be present. They relate the asymmetry to patterns of stop palatalization. Asymmetric misperception could also stem from other acoustic properties of segments, like the temporal distribution of retroflexion cues (Steriade 2001), or from properties of the auditory system, like the temporal spread of masking (Wright 1996, Wright and Ladefoged 1997); more research is needed.

A second class of perceptual bias factors, perceptual hypercorrection, was first identified by Ohala (1981). This arises when correction (perceptual compensation for coarticulation) applies to undo coarticulation that is actually absent. For instance, Beddor et al. (2001) found that listeners are relatively insensitive to vowel nasality variation when a nasal segment followed. They attributed this perceptual insensitivity to compensation for coarticulation, and noted that it correlates with a crosslinguistic tendency for vowel nasality contrasts to be suspended before nasal consonants.¹³ We shall note cases in section 3.4.4 where hypercorrection may be a plausible explanation of sound change.

To forestall misunderstanding, we should comment on the relation between hypocorrection (in Ohala's sense) and perceptual parsing bias factors for sound change. As noted above, hypocorrection is Ohala's term for a listener's failure to correct for coarticulation, which may then lead to sound change. A classic example involves interactions between vowels and coronal consonants. In a sequence like /ut/, the coronal tends to front the vowel so that its phonetic realization is closer to [yt].

¹³ Beddor (2009) argues that, in addition to compensation for coarticulation, the sound change $VN > \tilde{V}$ is influenced by natural patterns of gestural mechanics in the coordination of oral and nasal gestures. Interestingly, while there are a number of laboratory demonstrations of correction, there are almost no controlled observations suggesting that listeners hypercorrect in speech perception. The only example known to us is presented by Shriberg (1992); cf. Ohala and Shriberg (1990). This may be a gap in the literature, but it is an important one.

3. *Phonetic bias in sound change* 65

This mechanical effect (the overlap of consonant and vowel tongue gestures) does not ordinarily lead to sound change, due to a perceptual mechanism that compensates for coarticulation (Mann and Repp 1980); coarticulation is perceptually corrected. But if, for some reason, the listener fails to correct for coarticulation, a change may result: /u/ > [y] / __ [cor], with no change before other consonants. Something just like this seems to have happened in Central Tibetan (Dbus), as illustrated in (6). Final consonants were debuccalized or lost; the examples in (6a) show vowels that were unaffected, while those in (6b) show fronting when the final consonant was a coronal.

(6) Central Tibetan precoronal vowel fronting (Tournadre 2005: 28–32)

- a. Written Tibetan (WT) *brag* ‘rock’ > Central Tibetan (CT) *t^haʔ*
 WT *dgu* ‘nine’ > CT *gu*
 WT *p^hjugpo* ‘rich’ > CT *tɕ^hukpo*
- b. WT *bal* ‘wool’ > CT *p^hɛ:*
 WT *bod* ‘Tibet’ > CT *p^høʔ*
 WT *k^hol* ‘to boil’ > CT *k^hø:*
 WT *bdun* ‘seven’ > CT *d̄y*
 WT *sbrul* ‘snake’ > CT *d̄y:*

Hypocorrection is a key ingredient of change, both in Ohala’s and our account, but it is important to add that hypocorrection *per se* does not involve a specific bias factor. The bias factor in cases like (6)—the phonetic force that introduces variability and determines the direction of the change—is gestural. It is coarticulation that determines whether /u/ will drift toward [y] or [a] in coronal contexts. Hypocorrection helps determine whether or not a change will occur on a specific occasion, and as such it is part of a model of actuation; cf. section 3.6.

3.4 Bias factors in sound change

In this section we consider types of sound change that may reflect the bias factors summarized in section 3.3. Sound changes that can be attributed to motor planning, aerodynamic constraints, and gestural mechanics are well documented; perceptual parsing is somewhat harder to substantiate but remains a possible source of sound change.

3.4.1 *Motor planning*

Sound changes that have their origins in motor planning bias factors are, in effect, speech errors that catch on. In recent decades—indeed, since the classic studies of Meringer and Mayer (1895) and Meringer (1908)—it has been *démodé* to suggest that speech errors result in change. But while speech error research shows clearly

that sound change in general cannot be explained as conventionalized speech errors, it does not exclude the possibility that some types of sound change do have precisely that origin. This is our contention here.¹⁴

In section 3.3.1 we discussed two kinds of motor planning errors: blending and inhibition. We surmise that there is one common sound change type whose roots may lie in motor planning inhibition errors: nonlocal dissimilation. Since dissimilation is complex and its analysis is controversial, we discuss it separately in section 3.4.5. As for motor planning blending errors, we expect that sound changes emerging from them should tend to be anticipatory rather than perseverative, and should tend to involve an interaction between relatively similar segments and segments in relatively similar prosodic positions; greater similarity should favor the interaction. At least two types of sound change may conform to our expectations: consonant harmony and long-distance displacement (nonlocal metathesis).

Consonant harmony is illustrated by the Navajo patterns in (7). Note that harmony is symmetric; cf. /j/ → [s] in (7a) and /s/ → [ʃ] in (7b).

- (7) Navajo (Athabaskan) sibilant harmony (McDonough 1991, cited by Hansson 2010:44)
- a. /j-ɨf-mas/ → [jɨsmas] ‘I’m rolling along’
/ʃ-is-ná/ → [sɨsná] ‘he carried me’
 - b. /si-dʒé:ʔ/ → [ʃidʒé:ʔ] ‘they lie (slender stiff objects)’
/dz-ɨf-l-ta:l/ → [dʒɨʃta:l] ‘I kick him [below the belt]’

The data in (7) illustrate one common feature of consonant harmony: it is more typically anticipatory than perseverative.

A second common feature of sibilant harmony in particular is a ‘palatal bias’: many languages have /s/ → [ʃ] assimilation but no /j/ → [s] assimilation, but the reverse asymmetry is rare (Hansson 2010: 352–367). In Aari, for example, as seen in (8), affixal /s/ → [ʃ] when added to a root with /j/; only /s/ is affected.

- (8) Aari (Omoti) sibilant harmony: Causative formation (Hayward 1990)

BASE		CAUSATIVE	
mer-	‘forbid’	mer-sis-	‘cause to forbid’
du:k-	‘bury’	du:k-sis-	‘cause to bury’
di:b-	‘steal’	di:b-zis-	‘cause to steal’
ʃen-	‘buy’	ʃen-ʃɨf-	‘cause to buy’
ʔuf-	‘cook’	ʔuf-ʃɨf-	‘cause to cook’
ʃa:n-	‘urinate’	ʃa:n-ʃɨf-	‘cause to urinate’

¹⁴ In any theory positing occasional events (e.g. misperceptions or failures of perceptual correction) as sources of the variation that becomes conventionalized in change, it is hard to see what would exclude occasional speech errors from contributing to the same variation.

3. *Phonetic bias in sound change* 67

The same asymmetry is found in speech errors (Shattuck-Hufnagel and Klatt 1979). Stemberger (1991) relates this to ‘addition bias’ (Stemberger and Treiman 1986), whereby complex segments are anticipated in planning simple segments; [ʃ] is more complex because it uses the tongue blade and body.

As Hansson (2010) notes, consonant harmony patterns also resemble speech errors in being typically similarity-based: more similar segments interact with each other. In view of this and their other parallels (the nonlocality of consonant harmony, its typically anticipatory nature, and addition bias), Hansson suggests, and we agree, that phonological consonant harmony patterns are likely to have originated diachronically in motor planning errors.

Long-distance displacement (nonlocal metathesis) is a second type of sound change that may have its origin in motor planning. In the typology of metathesis sound changes (Blevins and Garrett 1998, 2004), it is notable that long-distance displacement commonly affects only some segment types. Often, for example, liquids undergo displacement leftward to the word-initial syllable onset. This is especially well documented in Romance varieties and languages influenced by Romance; Old Sardinian examples are shown in (9).¹⁵

- (9) Latin (L) > Old Sardinian (OS) liquid displacement (Geisler 1994:110–11)

L *castrum* ‘fort’ > OS *crástu*
 L *cochlea* ‘snail’ > OS *clocha*
 L *complēre* ‘fill’ > OS *clòmpere*
 L *dextra* ‘right (hand)’ > OS *dresta*
 L *februārium* ‘of February’ > OS *frevariu*
 L *pigrum* ‘slow’ > OS *prigu*
 L *pūblicum* ‘public’ > OS *plubicu*

Such displacements are usually anticipatory and tend to involve comparable syllable positions. For example, as in (9), displacement is often restricted to interchange between obstruent–liquid clusters. We take it that such phonologized patterns are rooted in motor planning. Independent support for this view comes from the fact that such displacements are a well-documented speech error pattern, as in German *Brunsenbenner* for *Bunsenbrenner* ‘Bunsen burner’ (Meringer and Mayer 1895: 91).¹⁶

¹⁵ In Old Sardinian, as Geisler (1994:112) notes, the displacement is restricted to adjacent syllables. In modern dialects, longer-distance displacements are also found: Latin *fenestra* ‘window’ > Old Sardinian *fenestra* > modern dialectal *fronèsta*. This chronological difference between one-syllable and longer displacement patterns undermines an argument by Blevins and Garrett (2004:134–5), based on comparable data in southern Italian dialects of Greek, that certain details of the longer displacement patterns favor the view that such changes originate through misperception.

¹⁶ While displacements of this type are not rare in speech error corpora, we have not studied the data carefully enough to judge whether other displacement patterns that are unattested as sound changes might also correspond to rarer speech error patterns. If they do, as Juliette Blevins points out to us, we would face the problem of explaining why such errors do not sometimes yield sound changes.

3.4.2 Aerodynamic constraints

The aerodynamic constraints on voicing and on frication summarized in section 3.3.2 have consequences for sound change. For example, the familiar change of final obstruent devoicing can be interpreted as an effect of the aerodynamic voicing constraint in a position where voicing is especially vulnerable.¹⁷

The aerodynamic frication constraint is likewise responsible for changes whereby voiced fricatives become glides. An example of the latter is the common pattern of [z] > [ɹ] rhotacism (Solé 1992a, Catford 2001). This change is known from many languages, including Latin and West and North Germanic. Its Old English (OE) effects are seen in words like *xerian* ‘to praise’, *ma:ra* ‘more’, and *xord* ‘treasure’ (cf. Gothic *hazjan*, *maiza*, and *huzd* respectively; OE *r* was probably [ɹ]). In the change in (10), OE [j] and [ɣ] became glides [j] and [w] when surrounded by voiced segments in Middle English (ME). When preceded by a liquid, ME *w* remained intact but in other positions the glides in (10) became diphthong offglides or underwent further changes. (The Middle English forms in (10) are not given in IPA.)

- (10) Middle English (ME) voiced dorsal fricative gliding (Luick 1940: vol. 2, pp. 945–6; the earlier forms shown in each case are from late OE or early ME)
- a. *cæ:ȝ* > ME *keie* ‘key’
e:ȝe > ME *eye* ‘eye’ (cf. German *Auge*)
plæȝian > ME *pleien* ‘play’
 - b. *laȝe* > ME *lawe* ‘law’
jeoȝuȝ > ME *youth* ‘youth’ (cf. German *Jugend*)
 - c. *borȝian* > ME *borwen* ‘borrow’ (cf. German *borgen*)
folȝian > ME *folwen* ‘follow’ (cf. German *folgen*)
morye > ME *morwe* ‘(to)morrow’ (cf. German *Morgen*)
sorye > ME *sorwe* ‘sorrow’ (cf. German *Sorge*)

The precise mechanism by which aerodynamic constraints yield new pronunciations warrants consideration. We prefer to avoid teleological formulations (e.g. [ɣ] > [w] ‘to avoid the combination of frication and voicing’), and we find it more appealing to assume that aerodynamic factors give rise to a biased distribution of variants. In voiced fricatives, for example, the tendency to reduced airflow behind the fricative constriction will automatically yield occasional glide variants. Sound changes like the ones illustrated above then take place when these variants become individual or community speech norms.

¹⁷ Other changes indirectly attributable to this constraint are noted in section 3.5.1.

3.4.3 *Gestural mechanics*

In section 3.3.3 we discussed two types of interaction among articulations: gestural overlap and gestural blending. The latter occurs when segments place competing requirements on a single articulator; gestural overlap involves interaction between independent articulators. Some very common types of sound change are rooted in gestural overlap, including those in (11):

- (11) a. VN > nasalized vowel
 b. Cluster simplifications that originate in gestural masking, e.g. [k^{nm}] > [km]
 c. Stop debuccalizations that originate in glottal coarticulation

For debuccalizations as in (11c), we assume that /k/ > [ʔ] changes may have an intermediate [kʔ] realization. If the glottal closure then masks the oral closure, the audible result is [ʔ].

A less common change originating in gestural overlap is the first stage of the English development in (12):

- (12) English velar fricative labialization: [x] > [f] / round V ___
 Old English **koxxian* > Middle English *kouxe* > *cough*
 Old English *xlæxxan* > Middle English *lauxe* > *laugh*
 Old English *ru:x* > Middle English *rouxe* > *rough*

Note that (as the modern *ou*, *au* spellings indicate) all three English words in (12) had a round vowel [u] before [x]. We follow Luick (1940: vol. 2, pp. 1046–53) and Catford (1977) in assuming that *en route* from [x] to [f] there was a realization like [x^w], resulting from overlap of the round vowel and [x]. Catford notes that a strongly rounded pronunciation can still be heard in southern Scotland: [lɑːx^w] ‘laugh’, [rʌx^w] ‘rough’, etc. The remaining [x^w] > [f] change is not due to gestural mechanics and will be discussed in section 3.5.1 below.

Typical changes due to gestural blend are coronal or velar palatalization (see further 3.4.4 below), the Tibetan precoronal vowel fronting pattern in (6) above, and vowel coalescence. Shown in (13), for example, are Attic Greek coalescence patterns for non-high non-identical short vowels. Here the coalescence of mid vowels preserves height; directionality is relevant in some but not all cases.¹⁸

¹⁸ Omitted in (13) are the coalescence of identical vowels as long vowels and of glide sequences as diphthongs. Note in relation to palatalization that not all ‘palatalization’ is the same: whereas coronal palatalization can be interpreted as an effect of gestural blend, labial palatalization would reflect gestural overlap.

- (13) Selected Attic Greek vowel contraction patterns (Rix 1992:52–3, Smyth 1956:19)

INPUT	CONTRACTION	EXAMPLE
e + o	o:	<i>p^hiléomen</i> > <i>p^hilô:men</i>
o + e	o:	<i>*dɛ:lóeton</i> > <i>dɛ:lô:ton</i>
a + o	ɔ:	<i>*ti:máomen</i> > <i>timḗ:men</i>
o + a	ɔ:	<i>*aidóa</i> > <i>aidḗ:</i>
a + e	a:	<i>*tí:mae</i> > <i>tí:ma:</i>
e + a	ɛ:	<i>génea</i> > <i>géne:</i>

Such examples are categorized as gestural blend because, in terms of vowel height and backness (not rounding), they involve a single articulator, the tongue body, on which successive vowel segments place conflicting demands.

3.4.4 Perceptual parsing

In section 3.3.4 we described three perceptual parsing phenomena that might yield sound change: symmetric misperception; asymmetric misperception; and perceptual hypercorrection. As we noted, symmetric misperception cannot generate asymmetric bias factors as such; in fact, it is rarely correlated with well-established (bidirectional) sound change patterns. Perceptual hypercorrection and dissimilation will be discussed separately in section 3.4.5. In this section, we discuss three types of sound change that have been attributed to asymmetric misperception: velar palatalization; unconditioned [θ] > [f] changes; and obstruent + [w] > labial obstruent shifts. In each case, there is some evidence that perceptual parsing underlies the change but other evidence pointing elsewhere.¹⁹ We regard the matter as unsettled.

Velar palatalization is the best-studied case where there may be a meaningful correlation between a sound change type and asymmetric misperception. One of numerous examples of this type of change is found in English, as shown in (14), where the highlighted examples of *k* and *tʃ* are from original **k*.²⁰

- (14) OE palatalization:
- *k*
- >
- tʃ*
- in syllables with front vowels (Sievers 1898: 101–5)

- a. Word-initial palatalization
tʃeaf ‘chaff’
tʃe:ap ‘cheap’
tʃild ‘child’

¹⁹ For example, Babel and McGuire (2010) report that [θ] perception is more variable than [f] perception in both audio and audio-visual stimuli.

²⁰ Only voiceless **k* palatalization is illustrated, because the interaction of spirantization with the palatalization of **g* would require more detailed exposition.

- b. Internal onset palatalization
**drenki- > drentʃ* ‘a drink’
ortʃeard ‘orchard’
ri:tʃe ‘rich’
- c. Coda palatalization
ditʃ ‘ditch’
pitʃ ‘pitch’
swiltʃ ‘such’
- d. No palatalization in syllables with back vowels
ku:θ ‘known’; cf. (*un*)*couth*
sak ‘sack’

Of course, as noted in section 3.4.3 above, gestural blending is implicated in velar palatalization, which arises from the interaction of articulatory instructions for a front vowel and a velar consonant. But a coarticulatorily palatalized velar is far from being an alveopalatal affricate; it is that distance that perceptual parsing accounts are meant to bridge. For example, Guion (1998) studied the perceptual similarities of velar stops and alveopalatal affricates and found that when stimuli are degraded by gating or noise masking, tokens of [ki] are significantly often misperceived as [tʃi], while tokens of [kɑ], [ku], [tʃi], [tʃɑ], and [tʃu] are more often perceived accurately. In a nutshell, [ki] is misperceived as [tʃi] but [tʃi] is not misperceived as [ki]. Guion suggests that velar palatalization leads to alveopalatal affricates because of this asymmetric misperception.²¹

Our main reservation regarding this argument is that it is not yet supported by phonetic studies of ongoing changes that show a clear articulatory leap from [kʲ] to [tʃ]. We hesitate not only because gestural blending is involved, but because it remains possible that the transition from [kʲ] to [tʃ] is mediated not by perceptual parsing but by processes that include perceptual enhancement (section 3.5.1). In Modern Greek, for example, velar palatalization yields palatals: /k g x ɣ/ → [c ɟ ç ʝ] before front vowels (Arvaniti 2007); some dialects have a further [c ɟ] > [tʃ dʒ] change. If this is a typical pathway for [k] > [tʃ] palatalization, we would want to evaluate the possibility that affrication of [c] reflects perceptual enhancement. But insofar as clear cases of asymmetric misperception are identified, and are correlated with sound changes that do seem to have originated as articulatory leaps between the relevant segment types, it is likely that they are a source of sound change.

We are also uncertain about the asymmetric-misperception account of the [θ] > [f] change found in English and Scots dialects and some other languages.²² A point

²¹ The nature of Guion’s argument is similar to that of Chang et al. (2001), but they discussed asymmetric misperception of [ki] and [ti], which does not correspond to a well-attested sound change pattern.

²² For what it is worth, the change itself is not very common. Though it has occurred in several languages (Blevins 2004:134–5, Kümmel 2007:193), it is less common than the superficially comparable change /s/ > [θ], which evidently targets dental [s] and thus seems to have an articulatory basis.

in favor of this account, to be sure, is that experimental studies (Miller and Nicely 1955, Babel and McGuire 2010) show that [θ] is misperceived as [f] significantly more often than the reverse; this is consistent with the fact that a [f] > [θ] change is unknown.²³ But we suspect that the change may involve first the development of labialization on [θ], i.e. [θ] > [θ^w], with a further [θ^w] > [f] change that is similar to the English [x^w] > [f] change mentioned in section 3.4.3. We have three reasons for our suspicion. First, in Glasgow, to which the English [θ] > [f] change has spread in recent decades, there is a variant that Stuart-Smith et al. (2007) describe as a labialized dental fricative, perceptually intermediate between [θ] and [f].²⁴ Second, in South Saami and Latin there are cases where an interdental > labiodental fricative change is limited to labial contexts (Kümmel 2007:193); we interpret these as shifts targeting phonetically labialized interdentals, equivalent to the [θ^w] > [f] step that we assume for [θ] > [f] shifts generally. Third, within Northern Athabaskan, as analyzed by Howe and Fulop (2005) and Flynn and Fulop (2008), a reconstructible series of interdental fricatives and affricates has the outcomes in (15):

- (15) Selected reflexes of Northern Athabaskan interdental fricatives and affricates
- a. Interdentals: Dene Tha dialect of South Slavey
 - b. Labials ([p], [p^h], [p^ʔ], [f], [v]): Tulita Slavey
 - c. Labial-velars (e.g. [k^w], [k^{wh}], [k^{wʔ}], [ʌ], [w]): Dogrib, Hare, Gwich'in
 - d. Velars: Dene Tha and Gwich'in dialects
 - e. Pharyngealized sibilants: Tsilhqot'in

Howe and Fulop (2005) argue that the Tsilhqot'in development in (15e) was as in (16), and that all the outcomes in (15b–15e) passed through a labialized interdental stage.

- (16) Northern Athabaskan interdental fricatives and affricates in Tsilhqot'in
 [*_ɬθ, *_ɬθ^h, *_ɬθ^ʔ, *θ, *ð] > [*_ɬθ^w, *_ɬθ^{wh}, *_ɬθ^{wʔ}, *θ^w, *ð^w] > [ts^ɬ, ts^{ɬh}, ts^{ɬʔ}, s^ɬ, z^ɬ]

If so, two of the best-documented [θ] > [f] cases (in English and Scots dialects, and in Athabaskan) show evidence for an intermediate [θ^w] stage. Howe and Fulop (2005) and Flynn and Fulop (2008) suggest that the reason labialization emerges is that it enhances the acoustic feature [grave], which, they contend, characterizes interdentals; in their Jakobsonian formulation, [flat] enhances [grave]. In short, on this view of [θ] > [f] shifts, the initial bias factor driving them is not perceptual parsing but perceptual enhancement (section 3.5.1).

²³ As Nielsen (2010:10) points out, however, if it is asymmetric misperception that explains [θ] > [f] shifts, we might expect [θ] > [f] substitutions in English second-language learning; in fact other substitutions appear to be more common.

²⁴ We are not aware of detailed phonetic studies of the ongoing [θ] > [f] change in other dialects. Note that an independent earlier [θ^w] > [f] change is documented in Scots dialects: Old English *θwi:tan* > Buchan Scots *fəjt 'cut'* (Dieth 1932). Of course this does not prove that the same change happened later, but it establishes the change as a natural one within the phonological context of English and Scots.

3. *Phonetic bias in sound change* 73

A final common type of sound change where asymmetric misperception has been assumed is the ‘fusion’ of obstruent + [w] sequences as labial obstruents. In the typical examples in (17), sequences with stops fuse as bilabial stops and those with fricatives fuse as labiodental fricatives.²⁵ Two other examples were mentioned above: the Buchan Scots $\theta w > f$ change in note 24 and the hypothesized Tulita Slavey labiodental > labial shift in (15b).

- (17) a. Stop–glide fusion: Latin $dw > b / \# _$
dwellum > *bellum* ‘war’
dwenos > *bonus* ‘good’
 **dwis* > *bis* ‘twice’
- b. Stop–glide fusion: Ancient Greek $k^w > p$
 **wek^wos* > *epos* ‘word’
 **leik^wo:* > *leipo:* ‘I leave’
 **k^wolos* > *polos* ‘pivot’
- c. Fricative–glide fusion: Old English $xw >$ Buchan Scots f (Dieth 1932)
xwa: > *fa:* ‘who’
xwæt > *fat* ‘what’
xwit > *fajt* ‘white’
xwonne > *fan* ‘when’

Significantly, the fricative changes involve a bilabial > labiodental place of articulation shift. Note also that the Slavey change is non-neutralizing (the phonological inventory previously lacked labials) while the others are neutralizing.

In essence, the perceptual parsing account of changes like these is that [kw] is sufficiently likely to be misheard as [p], and [θw] or [xw] is sufficiently likely to be misheard as [f], for such misperceptions occasionally to give rise to a new phonological representation. Though we do not know of any relevant experimental work, we would not be surprised to learn that asymmetric misperception patterns such as these can be confirmed in the laboratory. Still, one or two points are worth making. First, competing with the perceptual parsing account is one based on articulatory change: an account in which the glide [w] becomes a stop or fricative before the immediately preceding stop or fricative articulation is lost. For example, according to the competing view, [kw] > [p] via intermediate [kp] (or the like) and [xw] > [f] via intermediate [xϕ] (or the like). That such an intermediate stage is possible has support from several sources. For the stop changes in (17), Catford (1977) mentions examples like that of Lak and Abkhaz, where, for example in Lak, /k^w/ is realized as [kp^ʔ]. Catford writes that the ‘the labial element is an endolabial stop: the lips are pushed

²⁵ In some cases the glide is printed as a secondary articulation, in other cases as a distinct segment. This reflects the standard phonological analyses of the languages and probably does not signify any relevant phonetic difference.

forward, but kept flat (not rounded)', and suggests that the Greek change in (17b) may have passed through the same stage. As Larry Hyman reminds us, labialized velar > labial-velar changes are also well documented in Africa, for example in the Eastern Beoid (Niger-Congo) language Noone (Hyman 1981; Richards 1991). To confirm the perceptual parsing account of [k^w] > [p] changes, it would be desirable to identify an ongoing case where such a change involves no intermediate variants.

For fricative changes such as [x^w] > [f], Catford (1977) compares Scots dialects:

The labialisation becomes quite intense, towards the end of the sound, and, intervocalically, almost completely masks the sound of the velar component. Anyone who heard a South Scot saying 'What are you laughing at', [ˈx^wɹt əɾ i ˈlɑːx^wɸən ət] can have no further doubts about how [x] developed to [f].

It is important to note the difference between [ɸ] and [f]. It may be that the shift to a labiodental place of articulation is due to perceptual parsing, but since labiodental fricatives are noisier than bilabial fricatives it may alternatively be possible to assume auditory enhancement (of continuancy). In any case, for the stop changes (e.g. [kw] > [kp]) and the fricative changes (e.g. [xw] > [xɸ]), we are left with the question of whether the emergence of [p] and [ɸ] respectively is due to perceptual parsing (e.g. [kw] misperceived as [kp]), articulatory variability (e.g. [w] occasionally pronounced with lip closure or near-closure), or some other cause.²⁶ The question strikes us as unresolved, and with it the role of perceptual parsing in sound changes of the three broad types examined in this section, which target palatalized and labialized obstruents. We turn in the next section to a final type of sound change that has been attributed to perceptual parsing.

3.4.5 *Nonlocal dissimilation*

Broadly speaking, there are two competing explanations of nonlocal dissimilation.²⁷ As discussed above, the well-known model of Ohala (1981, 1993b) explains dissimilation as an effect of perceptual hypercorrection; cf. Gallagher's (2010) recent study invoking perceptual processing. A traditional competing explanation appeals to motor planning errors (Grammont 1895, Carnoy 1918, Grammont 1939, Frisch 2004, Frisch et al. 2004, Alderete and Frisch 2007). For example, Carnoy (1918: 104) writes that 'when two sounds or two syllables coincide and have to be visualized together and articulated after one another . . . the image of one of them easily crowds out the image of the other'; we take this as a reference to planning. Somewhat less obscurely, Alderete and Frisch (2007: 387, citing Berg 1998 and Frisch et al. 2004) refer to a

²⁶ Dialect variation in the realization of Swedish 'sj' may be fertile ground for studying fricative place of articulation change. This sound, which is described by the IPA as a voiceless simultaneous palatal-velar fricative, has a variety of realizations in dialects of Swedish, including a velarized labiodental variant [ɸx] (Lindblad 1980; Ladefoged and Maddieson 1996).

²⁷ See Alderete and Frisch (2007) and Bye (2011) for overviews and general discussion with reference to further literature.

3. *Phonetic bias in sound change* 75

‘functional motivation . . . in the difficulty of processing words containing repeated segments during speech production.’

We believe that it is worth re-examining the motor planning account of nonlocal dissimilation. As background we begin by presenting four typical dissimilatory sound changes. The first is a less celebrated case of the most famous example of dissimilation, Grassmann’s Law in Indo-European. This term refers to independent changes (in Greek and Sanskrit) whereby the first of two nonadjacent aspirated stops was deaspirated. It has been suggested that the same change may also have happened in the prehistory of Latin; examples are shown in (18).

(18) Grassmann’s Law in Latin (Weiss 2010: 156)

- a. $*b^h ard^h \bar{a} > *bard^h \bar{a}$ (> *barba* ‘beard’; cf. OCS *barda*, English *beard*)
- b. $*g^h lad^h ros > *glad^h ros$ (> *glaber* ‘smooth’; cf. German, Yiddish *glatt*)

The crucial change in (18) was prehistoric: $*b^h > *b$ in (18a), $*g^h > g$ in (18b). The change is shown by the eventual Latin outcomes, with initial *b* and *g* in (18a–b) respectively. Without dissimilation, regular Latin sound changes would have yielded initial $*b^h > f$ in (18a), i.e. †*farba*, and probably initial $*g^h l > l$ in (18b), i.e. †*laber* (just as $*g^h r > r$ in *rāvus* ‘gray’; cf. English *gray*).

Another laryngeal feature is targeted by a Secwepemctsin (Shuswap) change that has been called a Salish Grassmann’s Law (Thompson and Thompson 1985). Dissimilatory deglottalization is shown in (19) with diachronic and synchronic examples.

(19) Secwepemctsin dissimilatory deglottalization

- a. Diachronic examples (Thompson and Thompson 1985)

PROTO-INTERIOR-SALISH		SECWPEMCTSIN
$*k^? ip^?$	‘pinch’	$k^? ip^? -m$
$*q^w at^s^?$	‘full’	$q^w et^s^? -t$
$*ts^? ek^w^?$	‘shine’	$tsək^w - tsək^w^? -t$

- b. Synchronic examples: Reduplication and infixation (Kuipers 1974)

NO DISSIMILATION TRIGGER		DISSIMILATION	
$k^? jəj$	‘be cold, freeze’	$t-kj-k^? ij-t$	‘chilled’
$q^? iχ-t$	‘strong’	$qə-qi-q^? χ-t$	‘stronger’
$q^? iw-t$	‘to break’	$qw-q^? -iw$	‘brittle’
$st^? ek^w$	‘to show off’	$ste-t^? -k^w$	‘smarty’
$k^w inx$	‘how many?’	$k^w i-k^w^? -nx$	‘how many (animals)?’

Finally, in (20–21) we illustrate typical sonorant dissimilations. Liquids are the most common segment type to be affected by nonlocal dissimilation, as in Sundanese, where an infix /-ar-/ surfaces as /-al-/ when it is followed somewhere in the word by

r; examples are in (20).²⁸ Dissimilatory changes involving *l* and *r* in morphology are crosslinguistically common.

- (20) Liquid dissimilation in Sundanese (Western Malayo-Polynesian; Cohn 1992)

	BASE	PLURAL	
A	poho	p<ar>o ho	‘forget’
	gilis	g<ar>i lis	‘beautiful’
	ayim	<ar>a yim	‘patient’
	di-visualisasi-kin	di-v<ar>i sualisasi-kin	‘visualized’
B	dahar	d<al>a har	‘eat’
	pərceka	p<al>ə rceka	‘handsome’
	motret	m<al>o tret	‘take a picture’

In (21), we see cases in Italian where original *n...n* sequences dissimilated to *l...n*. The first of two nasals lost its nasality and became another coronal sonorant.

- (21) Lexically irregular nasal dissimilations in Italian

SOURCE	ITALIAN	
Celtic <i>Bononia</i>	<i>Bologna</i>	‘Bologna’
Greek <i>Panormos</i>	<i>Palermo</i>	‘Palermo’
Latin <i>venenum</i>	<i>veleno</i>	‘poison’
Latin <i>unicornis</i>	<i>licorno</i>	‘unicorn’

The examples in (18–21) are typical of the featural and positional typology of dissimilation. In featural typology, typical dissimilation targets include secondary features such as aspiration as in (18), glottalization as in (19), labialization, and palatalization, as well as some sonorant features, including nasality as in (20) and most especially liquid features as in (21). This profile has been interpreted in two main ways. First, Ohala (1981:193) writes that ‘only those consonantal features should participate in dissimilation which have important perceptual cues spreading onto adjacent segments.’ (This view has the potential problem that in cases like (21), it is necessary to assume that velum lowering in *unicornis* spanned an intervening [k].) Second, Carnoy (1918) suggests that dissimilation typically targets features that are either articulatorily complex (he mentions the trill [r]) or ‘more fugacious and more inconspicuous’ (including aspiration and glottalization).

In any case, we are struck by parallels, having to do with liquids, between the featural profiles of dissimilation and of motor-planning speech errors. The speech errors in (22) are unambiguously dissimilatory in nature; in (22a–c) the output of liquid dissimilation is also a liquid, while the output in (22d) is a nasal.

²⁸ An additional pattern is that with an *l*-initial base, the infix undergoes assimilation and surfaces as /-al-/: *litik* ‘little’ → plural *l<al>itik*.

3. Phonetic bias in sound change 77

(22) Liquid dissimilations in speech errors

- a. *Das ist doch ungrau ... unglaublich* ‘that’s incredible’ (Meringer 1908:93)
- b. *Eine Partei muss auch in den verschiedenen Gle ... Gremien die Fragen der Zeit diskutieren.* (Berg 1998:182–3)
‘A political party also has to discuss the current issues in its various committees.’
- c. *the blide of Frankenstein* (for *the bride of Frankenstein*) (Fromkin 2000: no. 1711)
- d. *zwei Fliegen mit einer Knapp ... Klappe schlagen* (Berg 1998:178)
‘to kill two birds with one stone’

These two speech error outcomes correspond to the two most common diachronic liquid dissimilation patterns.

The examples in (23) are ambiguous because *l* and *r* are both present in the immediate context in each example, so the errors might in principle be assimilatory; but in each case positional parallelism—*la/le* in (23a), *gr/gr* in (23b), *bl/fl* in (23c)—suggests that dissimilation is a likelier interpretation.

(23) Liquid dissimilations in speech errors: Ambiguous examples of planning inhibition

- a. *Kravierlehrer* (for *Klavierlehrer* ‘piano teacher’) (Meringer and Mayer 1895:96)
- b. *ein grosser Gleu ... Greuel* ‘a great abomination’ (Meringer 1908:93)
- c. *übergebri ... gebliebenes Fleisch* ‘left-over meat’ (Meringer 1908:93)

Dissimilatory speech errors are admittedly uncommon; those involving liquids are less than ten per cent as frequent as nonlocal assimilatory errors involving liquids.²⁹ But they are well enough documented, as illustrated in (22–23), that a theory of speech production should take account of them. And if dissimilatory speech errors are a clear pattern, they might in some cases lead to sound change.³⁰

We next consider positional typology: In what positions are segments the targets of dissimilation? A traditional generalization is that nonlocal consonant dissimilation is more often anticipatory, as in (18–21), than perseverative. This view is not supported in recent work (Bye 2011), but it is worth noting that the latter does not count lexically irregular cases or distinguish surface-true patterns from affixal alternations. In any case, based on a range of (mostly Indo-European) examples,

²⁹ The examples in (22–23) include the complete dossier of reasonably persuasive cases in the published corpora of Meringer (Meringer and Mayer 1895; Meringer 1908) and Fromkin (2000).

³⁰ We do not know of speech error studies for languages with phonological glottalization, aspiration, etc. The motor-planning account of dissimilation predicts the existence in such languages of dissimilatory speech errors involving those features.

Grammont (1895, 1939) argues that dissimilation tends to target segments in unaccented positions and in ‘weaker’ syllable positions (e.g. onsets rather than codas). The idea that typical targets of dissimilation are ‘weak’ positions and perhaps ‘weak’ features (secondary features such as aspiration) is consistent with a motor-planning approach. In interactions between nearby segments with identical features, motor plan inhibition (section 3.3.1) eliminates repetition by preserving the more salient (anticipated or positionally ‘stronger’) segment.³¹

3.5 Systemic constraints on phonologization

As discussed in sections 3.3–3.4, biases in speech production and perception provide the starting point in sound change, but they do not exhaust the processes of phonologization. Rather, as noted in section 3.1, they generate a pool of structured variation from which phonological patterns emerge; other processes too contribute to the outcome. In this section we identify some additional elements of phonologization that a full account will need to treat in detail, and we comment on possible associated bias factors.

3.5.1 *Enhancement*

The initial stages of sound changes that emerge from the bias factors discussed in sections 3.3–3.4 are either categorical or incremental. They are categorical if they are already phonetically complete in their initial stage. For example, if motor planning errors are a source of sibilant harmony, the erroneous pronunciation of [s] may already have been a fully changed [ʃ]. Our expectation is that changes rooted in motor planning and perceptual parsing are often categorical.

By contrast, in changes emerging from aerodynamic constraints and gestural mechanics, the structured variation found in the initial stage of phonologization may involve pronunciation variants that differ considerably from the eventual outcome. For example, the first stages of adjacent-vowel coalescence might involve only partial gestural overlap, with complete coalescence resulting only after several generations or longer. Similarly, there is apparently a range of intermediate pronunciations between [Vwx] and [Vx^w], or between the latter and [f]. We use the term *ENHANCEMENT* to refer to processes by which a relatively small initial bias effect is amplified to its eventual categorical result.³² This in turn has two distinct profiles.

³¹ Tilsen (this volume) proposes a connection between motor-planning inhibition and dissimilatory effects, grounded in the following experimental observations (from areas outside language): ‘when movement A to one target location is prepared in the context of planning a distractor movement B to a sufficiently different target location, then the executed trajectory of movement A deviates *away from* the target of movement B... In addition, more salient distractors induce greater deviations...’

³² This use of the term is not what Stevens and Keyser (1989) meant when they wrote about featural enhancement, but there are parallels. Some phonetic property is made more recoverable by changes in

3. *Phonetic bias in sound change* 79

First, in what we call ARTICULATORY ENHANCEMENT, the magnitude of an existing feature is enhanced. For instance, in a typical umlaut change targeting /uCi/, the shift from a partly fronted [ɥ] (the result of gestural blending) to a fully fronted [y] is a shift of gestural magnitude. Numerous changes driven by gestural mechanics can be described in comparable terms. In some such cases a secondary feature may become prominent. For example, the distinction between long and short vowels in English was enhanced by the promotion of redundant vowel quality differences between long and short vowels—yielding the modern tense/lax distinction, cued by both vowel length and vowel quality (e.g. [i:] vs. [ɪ]). In a sense, this is a perceptual phenomenon; a contrast is perceptually strengthened by exaggerating a redundant cue (Stevens and Keyser 1989, Whalen 1990, Kingston and Diehl 1994). But articulatory enhancement has not introduced any new phonetic cues, and thus has no place in a list of phonetic bias factors.

In some cases a feature is temporally realigned, yielding greater perceptual distinctness, rather than having its magnitude as such enhanced. For example, in the development of English [f] from earlier round vowels followed by [x], a crucial step was evidently a shift such as [ʷx] > [xʷ], in which labialization is realigned with the end of the fricative.³³ Presumably this timing change served to enhance the perceptual distinctness of labialization. Similarly, in the development of /ai/ diphthong centralization before voiceless consonants ('Canadian Raising'), Moreton and Thomas (2007) argue from age-graded phonetic data that the effect first emerged in the offglide and subsequently spread to and was enhanced in the nucleus. As they schematize the shift in *tight* vs. *tied*, a [t^haɪt] vs. [t^haed] difference evolved into [t^hΔɪt] vs. [t^haed].

Another more dramatic case of temporal realignment is described in Bessell's (1998) study of anticipatory consonant-vowel harmony in Interior Salish. In Snčicuʔumšcn (Coeur d'Alene Salish), this process targets vowels that are followed in the word by so-called FAUCALS: uvulars, pharyngeals, or /r/. Examples are given in (24).

(24) Snčicuʔumšcn anticipatory harmony (Reichard 1938; Bessell 1998)

NO HARMONY TRIGGER		HARMONY	
[tsɪ́f-t]	'it is long'	[tsɛ́f-ɔlqʷ]	'he is tall'
[sɛ́ttf-nts]	'he twisted it'	[nɛʔ-sáttf-ɛʔqs-n]	'crank (on a car)'

Crucially, as Bessell demonstrates, this process cannot be analyzed as phonetic spreading, because intervening consonants are demonstrably unaffected phonetically. She

pronunciation that highlight the phonetic essence of the sound. Stevens and Keyser noted that featural enhancement may be language-specific; this is consistent with phonetic enhancement in sound change.

³³ Cf. Silverman's (2006b) account of a Trique sound change whereby velars became labialized after [u]: *uk > [ukʷ], *ug > [ugʷ] (e.g. [nukʷah] 'strong', [rugʷi] 'peach'); non-velar consonants were unaffected (e.g. [uta] 'to gather', [duna] 'to leave something'). Silverman suggests that labialization emerged because those velar tokens that happened to be slightly labialized would have been more likely to be categorized correctly by listeners, and in this fashion labialized velars gradually evolved.

suggests that this pattern (which amounts to long-distance agreement) arose directly from the purely local vowel–consonant coarticulation found in closely related Interior Salish languages. She writes that the root cause of the shift is ‘that faucal features are maximally compatible with vocalic rather than consonantal structure . . . [T]he phonologisation of local coarticulation [in related languages] lays the ground for a more general assignment of faucal features to vocalic structure, so that faucal features appear on any preceding vowel’ (Bessell 1998:30). Note that in this as in other cases of articulatory enhancement, the basic direction of change is determined by articulatory factors; the bias emerges from gestural mechanics, not perceptual enhancement.

Second, in what we call AUDITORY ENHANCEMENT, a new articulatory feature is introduced with the effect of enhancing the auditory distinctness of a contrast. A classic example is lip rounding on back vowels, which positions vowels in the acoustic vowel space in a maximally dispersed way (Liljencrants and Lindblom 1972), thus enhancing the overall perceptual contrast in the vowel system. Other redundant secondary features that can be analyzed in a similar way include the labialization of [ʃ]. In our discussions of individual sound changes above, we have also identified several developments, listed in (25), that may be attributable to auditory enhancement.

- (25) Possible examples of sound change due to auditory enhancement
- a. Prenasalization in voiced stops enhances voicing (section 3.3.2)
 - b. [θ] > [θ^w] enhances [flat] (section 3.4.4)
 - c. [xϕ] > [f] enhances continuancy (section 3.4.4)

The emergence of auditory enhancement could be envisioned in at least two ways. One possibility is that talkers possess linguistic knowledge of acoustic targets, and that new articulatory features are sometimes introduced in speech when a contrast is insufficiently salient. Such new features then spread like any other linguistic innovations. Another possibility is that features that emerge through auditory enhancement are occasionally present in natural speech, simply by chance along with other phonetic variants, but that because they enhance a contrast they have a privileged status in listeners’ exemplar memories, and are then more frequently propagated. We cannot judge which account is likelier. But whether the speaker-oriented or the listener-oriented approach ultimately proves more satisfactory, it is worth noting that auditory enhancement, unlike articulatory enhancement, does define a set of bias factors for linguistic change: new features may arise that auditorily enhance existing contrasts. This is a bias factor, but unlike those described in sections 3.3–3.4, it is system-dependent.³⁴

³⁴ Note that enhancement need not be regarded as teleological. For example, Blevins and Wedel’s (2009) account of anti-homophony effects may generate articulatory (and perhaps even auditory) enhancement effects as an automatic by-product of phonetic categorization.

3. *Phonetic bias in sound change* 813.5.2 *Selectional bias*

Since phonologization involves the transformation of a phonetic pattern into a categorical speech norm, part of a language's phonological system, it is possible that some selectional constraints intervene in this transformation. This could happen in several ways. For example, given a language whose pool of variation includes two equally robust patterns, both corresponding to known sound changes, the phonological structure of the language might make one pattern likelier to be selected. So, in a language with intervocalic lenition of some segment types, perhaps it is likelier that lenition of other segment types will be phonologized. Arguments along these general lines have been made by Martinet (1955): the structure of a language favors certain selections.

Another possible profile for selectional bias is that the phonologization of a phonetic pattern may be disfavored by the structure of a language or by universal properties of language, even if the precursor pattern is phonetically robust. This position is defended by Kiparsky (2006), who argues that final obstruent voicing never emerges as a sound change despite what he contends is the possibility that natural changes could conspire to yield a suitable phonetic precursor; his explanation is that there is a universal constraint against final obstruent voicing. Wilson (2006) also suggests that learning biases favor phonetically natural patterns. Similarly, Moreton (2008a, 2010) argues that comparably robust types of phonetic pattern are phonologized at different rates. For instance, phonologized dependencies between adjacent-syllable vowel heights are common, while interactions between vowel height and consonant voicing are rarely phonologized despite being phonetically robust. Moreton (2008a) attributes this to a learning constraint: single-feature dependencies are easier to learn.

Concerning these possibilities, we should emphasize two points. The first is that if selectional constraints exist, they constitute a second-order bias type, operating on patterns that are already structured along the lines we have discussed above. We have focused here on first-order bias types because we think it is helpful to sort these out first. Our approach thus differs from that of Kiparsky (2006), who acknowledges that selection (constrained by universal properties of language) operates on a pool of phonetic variation, but does not emphasize that phonetic variation is already structured. One of the key questions in phonological theory concerns the relative burden of selectional bias, as opposed to production and perception biases, in determining patterns of phonological typology.

The second point is that many aspects of selectional bias remain unclear. For example, it seems plausible that a language's phonological system could make some patterns likelier to be selected in phonologization, and it is easy to point to examples that can be interpreted in such terms after the fact. It is harder to show that this is what happened, and we think it is fair to say that the jury is still out. It is likewise obvious in principle that any universal constraints on grammar in general must also constrain

selection in particular, and that the discovery of selectional bias patterns with no other explanation may be evidence for universal constraints. But the details are debated; on final voicing compare Yu (2004), Blevins (2006a, b), and Kiparsky (2006). Finally, Moreton's suggestion of general learning constraints on learning seems reasonable *a priori*, but requires more investigation to be securely established as a source of linguistic asymmetries (cf. Yu 2011).

3.5.3 *Lexical and morphological effects*

A final system-dependent aspect of phonologization is worthy of brief discussion (we have little to add to existing literature) because it concerns the question of conditioning in (1b) on p. 51 above. This question has been a source of controversy since the neogrammarian era: what role do a language's lexical and morphological patterns play in sound change?

Concerning morphology, the question is whether the neogrammarians and many later historical linguists are right to claim that when morphological patterns seem to have played a role in sound change, what actually happened is that a later (independent) analogical change has interfered with its effects. It is often possible to reanalyze supposed cases of morphologically conditioned sound change along these lines. Nonetheless, the fact remains that apparently 'analogical' effects can be discerned before a phonological innovation has become categorical. First described by Bloomfield (1933: 364–6), who called it *SUBPHONEMIC ANALOGY*, this phenomenon has been studied by Trager (1940), Steriade (2000), and others in the recent laboratory phonology literature. At this point, we do not know in general how early in their life-cycle, and under what circumstances, morphological patterning plays a role in strictly phonological changes.

Concerning a language's lexical patterns, the main question has to do with the role of word frequency in sound change.³⁵ In the experimental literature, lexical effects on pronunciation variation are well established. For example, less frequent words tend to be pronounced with greater duration or greater articulatory effort than their more frequent homophones (Guion 1995); see Gahl (2008) and Bell et al. (2009) with references to other earlier work describing a range of leniting effects. This leads to an expectation that leniting sound changes should show frequency conditioning across a range of languages and historical contexts, but this expectation is not yet well supported in the literature. To be sure, cases of the expected type have been

³⁵ Another question concerns homophony avoidance; it has been suggested that a sound change is less likely if it neutralizes a contrast that distinguishes relatively many words (cf. Jakobson 1931; Martinet 1955 vs. King 1967), or that a sound change can be blocked in words where it would yield homophony (Gessner and Hansson 2004; Blevins 2005; Blevins and Wedel 2009). Research in this area is intriguing but not yet definitive. Hume's (2004b: 229) idea that 'more practiced articulatory routines' may influence sound change raises yet another possibility; she suggests that language-specific phonotactic frequencies may influence the direction of changes such as metathesis. This idea is attractive, though its overall role in the typology of sound changes remains to be assessed.

3. *Phonetic bias in sound change* 83

described in changes such as English vowel reduction (Fidelholtz 1975) and flapping (Rhodes 1992), among others summarized by Bybee (2001, 2002) and Phillips (2006), but three problems remain. First, many well-studied leniting changes show no frequency effects; examples include Latin rhotacism, Verner's Law, and the degemination of Latin geminate stops in languages like Spanish.³⁶ If word frequency effects are implicated in sound changes from their earliest stages, the difference between changes where these effects vanish and changes where they persist is unexplained. Second, the nature of the effects identified experimentally (a gradient relationship between frequency and duration) and in studies of phonological patterns (where words may fall into two frequency-determined groups, only one of which shows a change) are not precisely the same, and the relationship between them is not clear. And third, more than one sociolinguistic study has found, echoing the classical view of Bloomfield (1933: 352–362), that ongoing changes tend to exhibit lexical irregularities only late in their development, after they have become sociolinguistically salient, whereas 'the initial stages of a change' are regular (Labov 1994:542–3; cf. Labov 1981; Harris 1985). In our judgment not enough is understood yet about the emergence of frequency effects in sound change to build a coherent picture out of the contradictory facts. In any case, the role played by lexical and morphological patterns in grammar and usage is independent of the role played by bias factors for asymmetric sound change. Important as the question is, it falls outside the scope of this chapter.

3.6 A model of actuation

Weinreich et al's (1968) actuation question (1c), on p. 51, concerns the historically contingent appearance of a sound change in a particular place at a particular time. The phonetic and systemic bias factors identified above represent preconditions for change, and determine the direction of change if it does occur, but they do not explain why a change emerges in one community rather than another, or in one decade rather than another. What causes actuation?

Among the elements of actuation it seems necessary to distinguish two phenomena. First, given that bias factors are in principle present throughout a language community, in the speech of one or more individuals there must be a deviation from the norm for some reason. Whatever the phonetic precursor(s) of a change, someone must first use it (or them) more often or to a greater degree than is the community norm. Second, based on this, some other individuals must then modify their speech, or the nascent change will not endure. Milroy and Milroy (1985) refer to the two types of individuals as INNOVATORS and EARLY ADOPTERS, identifying social differences between them.

³⁶ Latin rhotacism comprised an intervocalic *s > *z change followed by a *z > r change. Verner's Law was a Germanic process of intervocalic fricative voicing (also conditioned by accent); notably, Verner (1877: 102–3) himself evaluated and rejected a frequency-based explanation of the exceptions to Grimm's Law that motivated his discovery.

Of course it is hard to observe innovators in the wild, but we can still ask the crucial question: What causes them to deviate from the norm? Why do some individuals speak differently from all the people around them?

To this first part of the actuation question there are several possible answers.³⁷ One answer, following Yu (2010a, this volume), appeals to individual differences in perceptual compensation. As discussed in section 3.3.4, perceptual compensation ordinarily leads listeners to ignore coarticulation effects. In an exemplar model of linguistic knowledge, this would have the effect of focusing an exemplar cloud more closely on its phonological target. Individuals with systematically attenuated perceptual compensation would therefore have more divergent exemplars in memory, mirroring the bias patterns discussed in section 3.4, and might then produce such variants more often.

A second possible answer would appeal to individual differences in linguistic development and experience. For example, language learners may develop different articulatory strategies for realizing the ‘same’ acoustic target. It may be that two such strategies yield perceptibly different outcomes in some contexts, such as coarticulation; this could be the point of entry of a sound change.³⁸ Or perhaps small random differences in experience—differences in what are sometimes called ‘primary linguistic data’—yield differences in the phonetic systems that learners develop.³⁹

A third possible answer, which we explore here, appeals to differences in sociolinguistic awareness. The basic idea is that individuals (or groups) may differ in how they assign social meaning to linguistic differences. We speculate that some individuals in a language community, but crucially not others, may attend to linguistic variation within their own subgroup but not to variation in other subgroups. If such individuals become aware of a particular phonetic variant in their subgroup, but are unaware that it is also present in other subgroups, they may interpret the variant as a group identity marker, and they may then use it more often. One social parameter that may give rise to such a dynamic is power; Galinsky et al. (2006:1071) suggest that power may ‘inhibit the ability to pay attention to and comprehend others’ emotional states.’ To this we might add a converse linguistic principle: lack of power sharpens one’s attention to linguistic variation (Dimov et al. 2012). What follows is meant as a proof of concept.

³⁷ The truth may involve a combination of answers. Or perhaps there is no answer—Labov (2010: 90–91) compares mass extinctions caused by a meteor: there is nothing biologically interesting about the causes of a meteor collision. But for linguistic innovation, we can at least hope to find some underlying linguistic or psychological causes.

³⁸ Individual phonetic differences without sociolinguistic salience have been identified in English vowel production (Johnson et al. 1993b), rhotic production (Westbury et al. 1998), and coarticulation patterns (Mielke et al. 2010); other such differences undoubtedly exist.

³⁹ This view of how change is triggered is common in the historical syntax literature (Lightfoot 1999); cf. Blevins’s (2006a:126) comment that sound change of the type she calls *CHOICE* ‘can depend on simple frequency changes of variants across generations, as well as differential weightings of variants based on social factors . . .’.

3. *Phonetic bias in sound change* 85

We are aware that it makes sociolinguistic assumptions that remain to be tested; we hope that this will stimulate future discussion of the details of linguistic innovation.

The approach we take, simulating the behavior of a collection of autonomous agents, has been used by previous researchers studying language change (Klein 1966, Pierrehumbert 2001a, Culicover and Nowak 2003, Galantucci 2005, Wedel 2006). Common to these and other models of phonological systems is the assumption that speakers are generally faithful in their reproductions of the phonetic forms of language, perhaps with the involvement of a phonetic retrenchment mechanism (Pierrehumbert 2001a); most also assume phonetic bias factors like those discussed above.⁴⁰ In addition to these model parameters, the simulations presented below add social variation so that social identity is a filter on variation.⁴¹

The bias factors discussed in sections 3.3–3.4 are sources of variance in linguistic performance. Ordinarily, in the course of speaking and hearing, the phonetic distortions introduced by these factors (whether in speech production or perception) do not result in sound change. This is because listeners usually disregard the phonetic variants introduced by bias factors. For example, as a result of categorical perception, listeners are less likely to notice small phonetic variations within phonetic regions associated with a phonetic category, while the same amount of variation is much more noticeable for sounds near a category boundary (Liberman et al. 1957; Kuhl 1991). Perceptual compensation for coarticulation is also known to ‘remove’ phonetic variation due to coarticulation; for example, nasalized vowels sound more nasal in isolation than when immediately followed by a nasal segment (Beddor et al. 2001). Similarly, listeners are able to detect mispronunciations and other speech errors and may disregard them. Even simple misperceptions may be disregarded by listeners when the speaker’s intent is discernible from context, as in the similarity of *can* or *can’t* in normal conversational English. ‘Corrected’ misperceptions, like speech errors, may be disregarded by listeners.

Given all this, if the usual pattern is for the variants introduced by bias factors to be filtered out by perceptual processing, how can bias factors play a role in initiating sound change? We suggest that at one level of representation bias variants are not filtered out, and that they are available for reanalysis in sound change. We will further suggest that social factors interact with bias variation in ways that lead to sound change. Our theory linking bias factors to sound change is based on the assumption that linguistic categories are represented by clouds of exemplars, and that speech production is based on such constellations of remembered instances.

⁴⁰ Within the framework of Optimality Theory the two assumptions correspond generally to faithfulness and markedness constraints (Prince and Smolensky 2004).

⁴¹ Another mechanism that has been utilized recently in multi-agent modeling of sound change is the ‘probabilistic enhancement’ proposed by Kirby (this volume).

The rest of section 3.6 has three parts, first establishing some parameters for the multi-agent modeling of sound change and then presenting a set of simulations. In subsection 3.6.1, we review exemplar models of linguistic memory and relate them to the study of sound change. In subsection 3.6.2, we review research on imitation and a variety of factors that influence it. Finally, subsection 3.6.3 presents the simulations.

3.6.1 *Exemplar memory*

Exemplar-based models of phonology (Johnson 1997a, 2006; Pierrehumbert 2001a) are based on the idea that the cognitive representation of a phonological object consists of all experienced instances of that object. This view of phonology is compatible with traditional theories of sound change that have referred to similar notions in explaining articulatory drift. Thus, already Paul (1880 [1920:49]) wrote that sound change is mediated by a set of ‘representations in memory’:

Even after the physical excitement [the direct experience of articulation and perception] has disappeared, an enduring psychological effect remains, representations in memory, which are of the greatest importance for sound change. For it is these alone that connect the intrinsically separate physiological processes and bring about a causal relation between earlier and later production of the same utterance.

In his view, random variation in the cloud of representations yields gradual articulatory drift. Similarly, Hockett (1965: 201) wrote about a density distribution in acoustic space measured over years:

In the long run (measured in years), the time-dependent vector that is the speech signal for a given speaker—both for what he himself says and for what he hears from others—spends more time in some regions of acoustic space than in others. This yields a density distribution defined for all points of the space. The density distribution is also time-dependent, since the speech signal keeps moving; we may also imagine a decay effect whereby the importance for the density distribution of the position of the speech signal at a given time decreases slowly as that time recedes further and further into the past.

The key aspect of exemplar memory models for sound change is that, in such models, the representation of a category includes variants. This is important because the cloud of exemplars may gradually shift as new variants are introduced. Exemplar theory provides an explicit model of how variability maps to linguistic categorization, and for sound change this model is important because it permits the accumulation of phonetically biased clouds of exemplars that serve as a basis for sound change. Exemplars retain fine phonetic details of particular instances of speech, so phonetic drift or sudden phonological reanalysis are both possible (as will be discussed in more detail below). Other models of the mapping between phonetic detail and linguistic categorization assume that phonetic detail is discarded during language use, and

3. *Phonetic bias in sound change* 87

therefore these theories offer no explanation of how phonetic detail comes to play a role in sound change.

There is a central tension in exemplar theory, however, which relates directly to sound change. We mentioned above several mechanisms (categorical perception, compensation for coarticulation, and mispronunciation detection) that lead listeners to disregard exemplars. More generally, it has become evident that not all exemplars have the same impact on speech perception or production. One particularly obvious point concerns differences between the phonetic space for listening and the phonetic space for speaking. Listeners may be perfectly competent in understanding speech produced in accents or dialects that they cannot themselves produce. For example, we are able to perfectly well understand our young California students at Berkeley, but neither of us can produce a plausible imitation of this variety of American English. The space of familiar exemplars utilized for speech perception is thus, evidently, larger and more diverse than the space of exemplars utilized for speech production. When we say, as above, that specific exemplars may be disregarded by listeners, this can be interpreted to mean that the variants introduced by bias factors are not added to the set of variants used in speech production.

Building on this idea that speech production and perception are based on different sets of phonetic exemplars, following Johnson (1997a) we posit that the perceptual phonetic space is populated with word-size exemplars for auditory word recognition. We follow Wheeldon and Levelt (1995) and Browman and Goldstein (1990a) in assuming that the speech production phonetic space is populated with smaller (segmental or syllabic) exemplars used in calculating speech motor plans. These articulatory exemplars are also recruited in certain speech perception tasks, and in imitation.

Evidence for this dual-representation model comes from a number of different areas of research. For example, in neurophonetics Blumstein et al. (1977) noted the dissociation of segment perception from word recognition in certain forms of aphasia. Hickok and Poeppel (2004) fleshed out a theory of speech reception in which two streams of processing may be active. A DORSAL stream involves the speech motor system in perception (Lieberman et al. 1967; Liberman and Mattingly 1985), and is engaged in certain segment-focussed listening tasks. More commonly in speech communication, speech reception is accomplished by a VENTRAL stream of processing that involves more direct links between auditory and semantic areas of representation.

Speech errors and perceptual errors differ qualitatively as a dual representation model would predict. In the most common type of (sound-based) slips of the tongue, segments in the speech plan interact with each other, to transpose or blend with the main factors being the articulatory similarity and structural position similarity of the interacting segments. For example, the [f] and [t] in the speech error *delayed auditory feedback* → ... *audif—auditory* ... share voicelessness and are in the onsets of adjacent

stressed syllables. Slips of the ear, on the other hand do not usually involve interaction of segments in an utterance, but are much more sensitive to whole-word similarity and availability of an alternative lexical parse (Bond 1999). For example, *He works in an herb and spice shop* was misheard as *He works at an urban spice shop* and *at the parasession* was misheard as *at the Paris session*.

Another source of support for a dual-representation model comes from the study of phonetic variation in conversational speech (Pitt and Johnson 2003). Johnson (2004) studied phonetic variation in conversational speech and found that segment and syllable deletion is extremely common. He concluded that auditory word recognition models that rely on a prelexical segment processing stage would not actually be able to perform accurate (human-like) word recognition and that whole-word matching is a better approach to deal with the massive phonetic variation present in conversational speech.

Proponents of the Motor theory of speech perception (Lieberman et al. 1967) argued for a special *SPEECH MODE* of segment perception. We can now hypothesize that in experiments that require listeners to pay careful attention to phonetic segments, this mode will dominate (Burton et al. 2000). But when listeners are mainly attuned to the meaning of utterances, the speech mode of listening will not be engaged (as much) and a *LANGUAGE MODE* of word perception will dominate. Lindblom et al. (1995) refer to the contrast as the ‘how’-mode vs. the ‘what’-mode of perception.

A dual-representation model of phonology is also consistent with several strands of thinking in psycholinguistics. For example, Cutler and Norris’s (1979) dual-route model of phoneme monitoring (as implemented in Norris 1994) holds that phonemes may be detected by a phonetic route, in a speech mode of listening, or via a lexical route where the presence of the phoneme is deduced from the fact that a word containing the phoneme has just been detected. They identified a number of factors that influence which of these two routes will be fastest. Two modes of perception were also implemented in Klatt’s (1979) model of speech perception. Ordinary word recognition in his approach was done using a whole-word matching system that he called *LAFS* (lexical access from spectra), and new words were incorporated into the lexicon using a segmental spell-out system that he called *SCRIBER*. This approach recognizes that reception of speech may call on either of these systems (or perhaps both of them in a race).

Dual representation is important in our model of sound change because articulatory targets tend to be resistant to change, and in particular sound change is not dominated by pronunciations found in conversational speech, as a naive exemplar model might predict given the predominance of ‘massive reduction’ (Johnson 2004) in conversational speech. This resistance to change is consistent with the idea that the speech mode of perception (and the consequent activation of articulatory representations) is somewhat rare in most speech communication.

3.6.2 *Imitation*

Laboratory studies of phonetic accommodation have shown that speakers adjust their speech on the basis of recent phonetic experience, i.e., that phonetic targets are sensitive to variation. In phonetic accommodation studies, subjects simply repeat words that they hear and are seen to adopt phonetic characteristics of words presented to them (Babel 2009 on vowel formant changes; Nielsen 2008 on consonant aspiration changes). Speech motor plans are maintained by feedback, comparing expected production with actual production, and evidently in phonetic accommodation the expected production (the target) is computed on the basis of one's prior speech exemplars, together with phonetic representations derived from hearing other speakers.

The feedback tuning of speech motor control can also be seen in the laboratory in studies of altered auditory feedback (Katseff et al. 2012). In altered feedback experiments, the talker hears (in real time) re-synthesized copies of his/her speech with the pitch (Jones and Munhall 2000), formants (Purcell and Munhall 2006; Houde and Jordan 1998; Katseff et al. 2012), or fricative spectra (Shiller et al. 2009) altered. Talkers respond by reversing the alterations introduced by the experimenter, even though they don't notice that a change was introduced. In both phonetic accommodation and altered auditory feedback studies, we see the operation of a phonetic mechanism that may be responsible for sound change: a feedback control mechanism that incorporates phonetic exemplars that the speaker hears others produce, or in other words a subconscious phonetic imitation mechanism.

Studies of phonetic accommodation and altered auditory feedback have found a number of parameters that are relevant for a theory of imitation in sound change. First, imitation is constrained by prior speaking experience. People do not imitate perfectly and do not completely approximate their productions to those of others (Pardo 2006; Babel 2009). Some of the inhibition is due to the speaker's own personal phonetic range; Babel (2009) found that vowels with the most variation in a subject's own speech showed the greatest accommodation. We speculate, though this has not been tested, that the degree of match between voices may influence imitation.

Second, imitation is socially constrained. People do not automatically or uncontrollably imitate others, but are more likely to imitate someone they identify with at some level (Bourhis and Giles 1977, Babel 2009). This has implications for sound change because it indicates that the use of bias variants in speech production is socially conditioned.

Third, imitation generalizes. Thus instances of long VOT influence speech in words or segments not heard; for example, /p/ with long VOT produces long (imitative) VOT in /k/ (Nielsen 2008). This finding has important implications for the regularity of sound change. The 'speech mode' system that we propose, by virtue of using segment-sized chunks provides an account of the regularity of sound change (where the receptive whole-word exemplar space would not). Interestingly, Nielsen's results

suggest that phonetic features, or gestural timing relations, may be represented in a way that they can be imitated in different segmental contexts.

Fourth, imitation is constrained by feedback in both auditory and proprioceptive sensory domains (Katseff et al. 2012). This finding is important because it helps define the range of phonetic imitation that is possible with ‘self-exemplars’—namely that proprioceptive feedback is involved. The implication of this is that imitation may be limited by sensory factors that are not immediately apparent to the linguist.

In addition to these observations drawn from prior research on imitation in phonetic accommodation, there are two general properties of imitation that we assume in our model of sound change. First, the only exemplars produced by others that have an impact on imitation are those that are processed in the speech mode of perception. Our dual-representation model entails that articulatory phonetic analysis of items does not always take place, thus not all instances of heard speech contribute to the pool of exemplars used in computing a motor plan.⁴²

Finally, speech production targets are calculated from a population of phonetic exemplars as a sort of weighted average where the ‘activation’ of each exemplar determines its weight in the calculation. Among the many factors that determine exemplar activation the intended linguistic category obviously matters a great deal, and there will also be residual activation from exemplars that have just been said (priming) and exemplars activated by what you have just heard.

It may be objected that imitation does not provide a link between bias factors and sound change, because the phonetic accommodation mechanism must presume that some speaker in the community has already undergone a sound change toward which other speakers are ‘drifting’. According to this objection, imitation is a mechanism for the spread but not the actuation of sound change. This fails to take account of two facts. First, listeners do not know whether the speech they are hearing is what the speaker intended to say, or if it has been altered by a bias factor. The listener’s inclination to imitate applies regardless of whether other speakers intend to produce changed variants or not. Second, listeners do not know whether they are hearing what the speaker actually produced or a perceptually distorted variant of the speaker’s pronunciation. In this case, the listener may imitate a figment of her own imagination. In either case, phonetic accommodation yields sound change, whether the target of accommodation is the result of a production or perception ‘error’ or not.

3.6.3 *Simulating sound change*

We implemented the assumptions discussed above in three simulations. They are in the spirit of Labov’s (1994:586–7) suggestion that ‘misunderstood tokens may never

⁴² Several researchers studying exemplar phonology have noted that word frequency effects are not as strong as a single-representation exemplar model would predict: Morgan et al. (under review); Pierrehumbert (2001a).

3. Phonetic bias in sound change 91

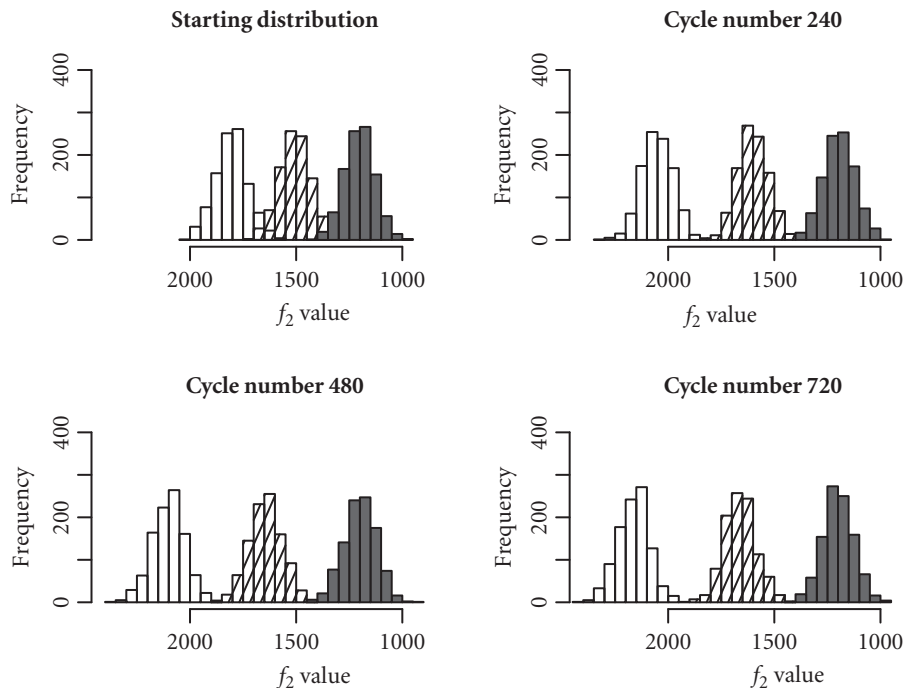


FIGURE 3.2 Simulating Labov's (1994) conception of how 'misunderstanding' is involved in sound change. The starting distribution graph shows histograms of vowel second formant (f_2) values of three vowels in a crowded low vowel space. The vowels overlap slightly because of articulatory phonetic variability. The remaining panels show how the vowels shift in acoustic space as we add heard exemplars to each vowel space. Each cycle involves sampling the space 1000 times, and then recalculating the mean vowel target for each vowel category. The model has two assumptions: (1) f_2 below 1000 Hz is unlikely, and (2) perceptually misidentified tokens are not added to a category's exemplar cloud.

form part of the pool of tokens that are used', so that if a listener 'fail[s] to comprehend [a] word and the sentence it contains...this token will not contribute to the mean value' of the target segment.⁴³ According to this view, perceptual confusion may result in conservation of a boundary between confusable phonemes, by limiting the exemplars of adjacent categories to only those that are correctly identified. The results of the simulation, shown in Figure 3.2, illustrate this. We created hypothetical vowel formant distributions that overlapped slightly and took a random sample of one thousand tokens from each distribution. Each vowel token was classified as an

⁴³ Simulations by Pierrehumbert (2001a) and Wedel (2006) echo in various ways the simulations presented here; see also Kirby (this volume). Like many authors, Labov assumes that the mean value of a cloud of exemplars is a rough indicator of a vowel target. This view may not be accurate (Pierrehumbert 2001a), but serves as a viable simplifying assumption in our model.

example of one of the three vowel categories based on its distance to the category centers. The category centers were then recomputed, with the misrecognized vowel tokens removed, and a new random sample of one thousand tokens was then drawn from each vowel category. In order to make the simulation more realistic we limited the possible vowel space and started the simulation with the back vowel (lowest f_2 value) located at the back edge of the space. This essentially fixed it in place with a mean of about 1200 Hz. As the figure indicates, after several cycles of selective exclusion of exemplars in the vowel categories, the category centers of the front vowel and the mid vowel shift so that they no longer overlap. This simulation illustrates a mechanism in speech perception that results in vowel dispersion (Liljencrants and Lindblom 1972).

In extending this style of simulation to study how bias factors result in sound change we included a social component. This was because we wanted to study not only how sound change might emerge from simple assumptions about exemplar-based phonological categories, but we also wanted a better understanding of the normal case where bias does not result in sound change. Therefore, the remaining simulations in this section track the development of phonetic categories in adjacent speech communities, where a sound change occurs in the system for one group while the other group does not experience the change. For both groups of speakers, we constructed phonetic categories that were represented by clouds of exemplars which include both normal variants and, crucially in both communities, a few exemplars (ten per cent) that have been altered by a bias factor. The key difference between the groups is whether or not the bias variants are disregarded. It seems reasonable to assume that variants produced by phonetic bias factors are usually ‘corrected’, either by perceptual processes like compensation or by rejection of speech errors. Stability of phonetic categories is thus the norm. As we shall discuss, we assumed that these ‘correction’ processes were not implemented to the same degree by all speakers; one group of speakers more actively applied perceptual compensation mechanisms than the other. Thus, the difference between groups is modeled as a difference in the exemplars selected by group members to define the phonetic category.

The top row of Figure 3.3 shows the starting phonetic and social distributions of our first simulation of social stratification and sound change. The simulation tracks the pronunciation of /z/ in two social groups. As discussed above, voiced fricatives like /z/ are biased by aerodynamic constraints, and sometimes are realized with reduced frication (more like an approximant). This simulation of a gradient phonetic effect is appropriate for modeling many types of sound change including context free vowel shifts, the despirantization of voiced fricatives, vowel fronting near coronal consonants, vowel nasalization, and vowel coalescence, among other changes. In this simulation, a bias factor produced a slightly skewed phonetic distribution. Most productions (ninety per cent) clustered around the phonetic target value, which was arbitrarily set to zero. A few productions (ten per cent), however, were a little biased

3. Phonetic bias in sound change 93

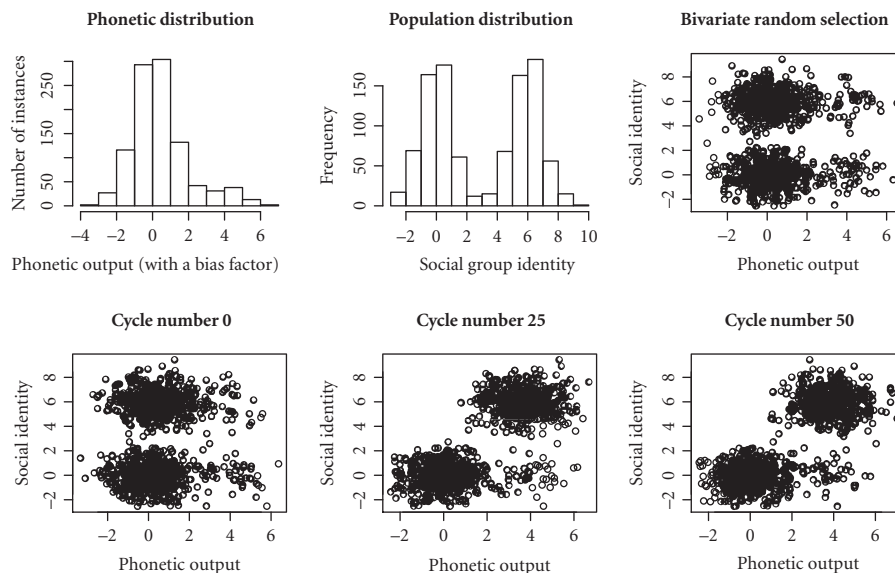


FIGURE 3.3 Simulation of a gradient phonetic bias. The starting phonetic and social identity distributions are shown in the histograms. The results of a bivariate random selection from these distributions is shown in the top right panel. Social group differences are indicated on the vertical axis, which measures an arbitrary ‘social identity’ parameter. Phonetic output is shown on the horizontal axis, where a value of zero indicates a voiced fricative production, and a value of four indicates a voiced approximant production. The bottom panels show the gradual phonetic drift, from iteration 0 to iteration 50 of the simulation, as the phonetic target includes approximated variants for one social group, and persistent phonetic instability for the other group who do not allow the inclusion of approximated variants to influence the target.

so that the phonetic distribution has a longer tail in one direction than it does in the other. The speech community in this simulation was also characterized by a bimodal social stratification with fifty per cent of exemplars produced by one social group and fifty per cent by another group of talkers. Each dot in the top right graph represents an exemplar in the sociophonetic space defined by phonetic output and social identity. At the start of the simulation there is no correlation between the phonetic and social values; the bias factor is equally likely to affect the speech of each population group. The bottom row of graphs shows how this phonetic system evolved over the course of fifty iterations of simulated imitation.

As seen in Figure 3.3, the phonetic output of the two simulated groups of speakers diverges. One group (centered around social identity index value 0) maintained the starting phonetic realization—a situation of persistent phonetic instability, where an aerodynamic bias factor influences about ten per cent of all /z/ productions, but this bias factor does not induce phonetic drift. The other group (centered around social identity index value 6) shows gradual phonetic drift, so that by the end of the

simulation the original /z/ is now /r/. Speakers in both groups are assumed to base their productions on a cloud of exemplars (using the mean value of a set of exemplars as a target). The difference is in the selection of exemplars to include in the cloud. The ‘o’ group, who did not experience a sound change, disregarded the phonetic bias variants—they successfully compensated for the bias and removed it from their exemplar-based phonetic definition of /z/. The ‘6’ group, who did experience the sound change, INCLUDED the bias variants in /z/, and thus the phonetic target was changed by the bias.

Why would different groups of speakers treat bias variants in different ways? Although bias variants occur with equal frequency for both groups of speakers, we assume that phonetically unusual productions may take on indexical meaning for ‘6’ group. Speakers who seek to identify with the group may be more likely to notice phonetic variation among group members and thus include it in as a group indexical property, even though that same variation exists in the population as a whole. Prospective group members may thus notice variants when they are produced by the target group even though they disregard those same variants when produced by other speakers. Considered from another point of view, a group that is aware of some social distance from another group may attend to phonetic deviations from the norm as marks of social differentiation.

It has to be admitted, though, that change caused by gradient bias may also be more inevitable than change induced by more discontinuous bias factors, in that listeners may be less likely to disregard bias variants that are only very minimally different from unbiased variants. Thus, variation introduced by a gradient phonetic bias may be less useful for social differentiation than a more discontinuous bias factor because it may fuel sound change regardless of social identity factors.⁴⁴ It is important, therefore, to study the link between discontinuous bias factors (such as those introduced by speech production or perception errors) and sound change.

To model more discontinuous phonetic bias factors such as the motor planning errors that we posited for cases of consonant harmony, the same basic model can be used. However, discontinuous bias is often structure preserving in the sense that speech errors often result in sounds already present in the language, so we assume that the basic mechanism is one of probability or frequency matching (Vulkan 2000; Gaissmaier 2008; Koehler 2009; Otto et al. 2011). For example, we can model the harmony process that results in a change from [s] to [ʃ] by assuming that one group includes harmonized instances of [ʃ] in the exemplar cloud for /s/ while the other group does not. Then, following Pierrehumbert (2001a), we assume that speech production involves a process that results in frequency matching so that the likelihood of drawing from one or the other mode in the phonetic distribution (that is [s] or [ʃ]) matches the frequency of exemplars in those regions of phonetic space.

⁴⁴ But note that this is definitely not Labov’s (1994) view.

3. Phonetic bias in sound change 95

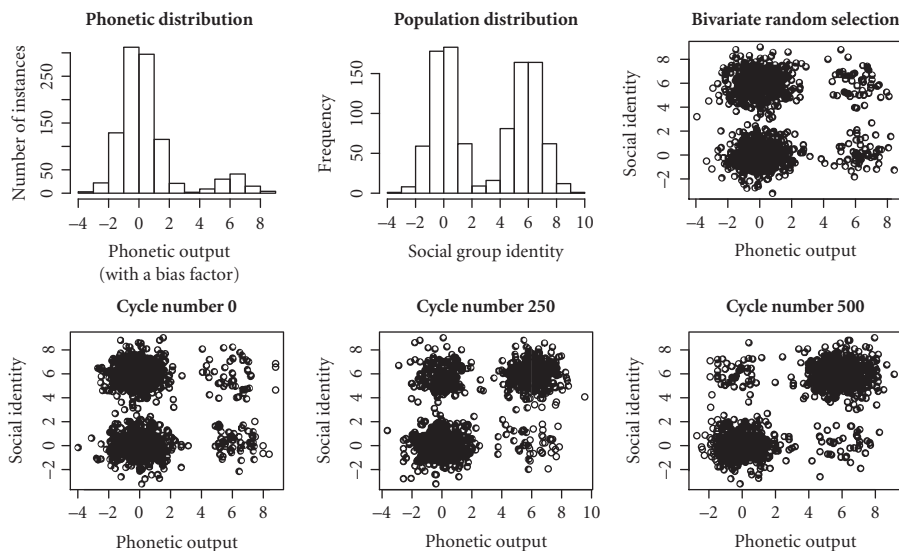


FIGURE 3.4 Simulation of a sound change caused by a discontinuous phonetic bias (such as a motor planning error that results in a consonant harmony)

The simulation (Figure 3.4) was structured in much the same way as the previous one. We have a population of individuals who are evenly divided into two social groups. We also have a phonetic distribution in which ten per cent of the output tokens are mutated by a phonetic bias factor. In this case, though, the bias factor produces a discontinuous jump in phonetic space. However, we cannot suppose that acceptance of the bias variants into a phonological category would result in gradual phonetic drift because the intermediate phonetic space may be unpronounceable, or the bias variants are good instances of an existing phonetic category. So the average phonetic target centered around /s/ (phonetic output equal to zero in the model) stays as it was, as does the average phonetic target centered around /ʃ/ (the bias variant, modeled with phonetic output equal to 6). However, speakers in one group are willing to accept bias variants as acceptable ways to say forms with an /s...ʃ/ sequence, while speakers in the other group do not accept bias variants. Thus with a frequency matching production model, where the speaker's produced distribution of variants matches the distribution of the exemplar cloud, the bias factor may lead to wholesale change.⁴⁵

These simulations of the link between phonetic bias factors and sound change have shown that exemplar-based models provide a useful, explicit method for studying the

⁴⁵ This simulation provides a useful reminder of the importance of compensation mechanisms, for phonetic stability. If the simulation is allowed to run over thousands of epochs the frequency matching mechanism, plus the phonetic bias factor, leads to oscillation between [s] and [ʃ]. The model does not stabilize unless the group who shifted from [s] to [ʃ] begin to treat instances of [s] as errors which should be corrected and thus removed from the exemplar cloud.

role of bias factors in sound change. We have also shown, with citations from Paul, Hockett, and Labov, that an exemplar-based conception of human phonetic memory is the mainstream view.⁴⁶

The simulations also identified a crucial role for exemplar selection in sound change, and in particular concluded that socially motivated exemplar selection rules make it possible to model both sound change and phonetic stability. Building on this finding, we speculate that a group who tend to accept bias variants (phonetic variants caused by bias factors) is likely to be engaged in a project of social differentiation, and are looking for cultural material that could be of value in this project. Thus, bias variants, though phonetically confusing, may be socially useful. Although this is stated as if it is a phonetically conscious activity, it need not be. To the extent that changes are ‘involuntary’ and ‘unconscious’ (Paul 1880; Paul 1920, chapter 2; Strong et al. 1891, chapter 1), we can speculate that a low status group who seek social identity with each other, against some other group, may be more attentive to phonetic detail than a group who feel secure in their social standing.

Finally, although we used an exemplar memory in all of the simulations, we used two kinds of mechanism to model sound change—phonetic target recalculation for gradient bias factors (Figure 3.3) and frequency matching for discontinuous bias factors (Figure 3.4). This difference relies on what Hockett (1965) called the ‘Quantization hypothesis’—the idea that the continuous range of phonetics is, for speakers, divided into discontinuous quanta of phonetic intentions. In the exemplar model, the difference boils down to whether the bias factor should be interpreted as changing the articulatory plan for a specific gesture, or changing the production rule used to select gestures in word production. One is tempted to associate this difference also with neogrammarian sound change, as against lexical diffusion (as Labov 1981 did). But there is no reason to believe that frequency matching is any less regular than target changing—that is to say, there is no reason to think that the shifting frequency distributions of [s] and [ʃ] would not affect all tokens of [s].

3.7 Conclusion

In this chapter we have outlined a framework for categorizing and understanding some key features of sound change. Much remains to be examined from this point of view, of course, including questions only touched on above. For example, how do processes of enhancement (section 3.5.1) work? How do we interpret lexical and morphological effects in sound change (section 3.5.3)? And what actual sociolinguistic and psychological evidence bears on the specific theories of actuation discussed in section 3.6?

⁴⁶ That is to say, the exemplar approach is mainstream in that part of linguistic research that Strong et al. (1891:1) called the ‘science of language’, as opposed to ‘descriptive grammar’.

3. *Phonetic bias in sound change* 97**TABLE 3.5** Well-established bias factors and representative changes

BIAS FACTORS	REPRESENTATIVE SOUND CHANGES
PRODUCTION AND PERCEPTION BIAS	
Motor planning (3.3.1)	Consonant harmony; anticipatory displacement (3.4.1)
Aerodynamic constraints (3.3.2)	Rhotacism, other fricative-to-glide shifts; final devoicing (3.4.2)
Gestural mechanics (3.3.3)	Palatalization; umlaut; VN > \tilde{V} ; vowel coalescence (3.4.3)
SYSTEMIC BIAS	
Auditory enhancement (3.5.1)	Interdental fricative labialization; back vowel rounding

We have described two broad classes of bias factors that may help explain asymmetries in sound change. The first, our main focus (sections 3.3–3.4), consists of bias factors emerging in speech production and perception through motor planning, aerodynamic constraints, gestural mechanics, and perceptual parsing. Despite its familiarity, we suggested that perceptual parsing is the least securely established factor; its prototypical examples may have other interpretations. More research is in order on this and all the other production and perception bias factors we discussed.

Systemic constraints (section 3.5) are a second broad class of bias factors, arising from language-specific or universal features of a phonological system. This class includes perceptual enhancement and in particular auditory enhancement, which can yield asymmetries in sound change; selectional bias (favoring certain variants, universally or in certain phonological systems); and perhaps lexical effects. Since some of the bias factors in this broad class are less well established at this point, the eventual dossier may be smaller than what we have identified. In Table 3.5 we summarize some of the best established bias factor types in both broad classes, with a few representative sound changes that we have mentioned.

Finally, since any full account of phonologization must address the emergence of speech norms (in an individual or community) from occasional phonetic variants, we have sketched the outline of a linking theory that relates them (section 3.6). Whether this sketch and our discussion of bias factors are on the right track or in need of substantial revision, we hope in any case to stimulate further discussion of the phonetic bases of phonologization.