10

Interpreting Misperception

Beauty is in the Ear of the Beholder

Juliette Blevins

10.1 INNOCENT MISPERCEPTION

In 1910, forty years after receiving his doctorate in Leipzig, Baudouin de Courtenay (1845–1929) published *O prawach grosowych* ("Phonetic Laws"). In this work Baudouin summarizes much of his life's work in modeling sound patterns and sound change, dissecting phonetic laws into psychological laws (linguistic representations), and laws of the manifestation of linguistic representations by means of organic physical media (implementation of linguistic representations as linguistic forms). It is in this work that Baudouin outlines one of his most significant contributions to the theory and method of phonological analysis: he characterizes misperception as a significant source of sound change, and suggests investigations of the nature of such misperceptions by experimental methods.

> I must emphasize the importance of errors in hearing (lapsus auris), when one word is mistaken for another, as a factor of change at any given moment of linguistic intercourse and in the history of language as a social phenomenon. Experimental methods can help to define the types and directions of these errors which depend on the physical conditions, on the sense of hearing of individuals, and on the degree of morphologization and semasiologization of the mobilized articulatory and auditory representations. (Baudouin 1972*b*: 267–8)

In the modern era, there is perhaps no single person who has given more substance to Baudouin's insight than John J. Ohala. Ohala has identified recurrent common sound changes with sources in misperception, and has used experimental methods to simulate misperception in the laboratory (e.g. J. Ohala 1981*b*, 1990*a*; see also Jonasson 1971). The influence of Ohala's work in this area has been dramatic. In addition to his own numerous publications, there is now a growing literature on sound patterns whose origins are primarily perceptual (e.g. Janson 1983; Beddor *et al*. 1986; Blevins and Garrett 1998, 2004; Foulkes 1997; Guion 1998; Majors 1998; Hume and Johnson

2001b; Plauché 2001). And this work has, in turn, had a dramatic influence on modern phonological theory: within the last decade there has been widespread recognition that certain recurrent sound patterns have phonetic explanations grounded in aspects of speech perception and production. These range from Optimality-Theoretic models (e.g., Steriade 1999, 2001) and generative rule-based accounts (e.g. Hale and Reiss 2000), to functional adaptive approaches (e.g. Lindblom *et al.* 1995), and Evolutionary Phonology (Blevins 2004, forthcoming), where markedness constraints are eliminated from synchronic grammars.

Nevertheless, there has been some resistance among phonologists to accept two fundamental implications of Ohala's research: (i) that innocent misperception can lead directly to attested recurrent sound patterns; and (ii) that sound change is non-teleological. In this chapter I look at sources of this resistance. In some cases, experimental results are simply ignored. In others, interpretations of perception experiments are not empirically motivated, and fail to recognize lexical effects. A final source of resistance to non-teleological models of sound change involves simplification of the model: "innocent misperception" is argued to be incapable of accounting for a particular phenomenon alone, motivating the implication of markedness effects. This simplification is a serious misinterpretation of Ohala's position and that of many others, where speech perception, speech production, and language use all play important roles in the evolution of sound patterns.

This chapter, then, concerns itself generally with scientific method. More specifically, I review the interpretation of experimental results in the domain of speech perception and their relevance for phonological theory. In phonology, as in many other fields, interpretation of results is often biased by theoretical orientation. Continued constructive dialogue between those carrying out experiments and others who might interpret them can help to identify biases and idealizations within the phonological literature, and can contribute to the common goal of developing plausible and accurate grammatical models.

10.2 Velar palatalization: not hearing results

While coarticulation is the source of many local assimilations, perceptual factors have also been shown to play an important role in sound changes sometimes viewed as assimilatory. One of the most influential papers in this area is J. Ohala (1990*a*), "The phonetics and phonology of aspects of assimilation." In this article, Ohala presents experimental results supporting recurrent asymmetries in the phonology of major place assimilation. In VC_1C_2V sequences, where the medial sequence is heterorganic, listeners misinterpret the sequence as having the place features of C_2 with greater-than-chance frequency when the cluster closure duration is relatively short. This is attributed to the more reliable and robust place cues of the C_2 burst than formant transitions into C_1. Ohala's results are important because they account for two features of sound patterns involving VC_1C_2V sequences: (i) there is a tendency for the consonant sequence to become homorganic; (ii) where major place features are concerned, there is a tendency for assimilation to be regressive. As Ohala (1990*a*) points out, the standard generative account makes neither prediction. There is no explanation for why place, as opposed to, say, manner features assimilate. And there is no account of the directional asymmetry found for major place features.[1]

Another sound pattern often attributed to coarticulation is velar palatalization. Many languages show sound changes or synchronic alternations involving velar palatalization before palatal segments. The most common change of this kind is *k > t͡ʃ / _{i,j}. Palatalizations of voiced velar stops are less frequent than voiceless ones, and velar palatalization is more common before front high vowels/glides than non-high front vowels. Within the Indo-European family, Old English palatalization of /k,g/ before front vowels and Slavic palatalization of /k,g/ when preceded by front non-low vowels are well documented. Other similar cases are found in the history of Indo-Iranian, Bantu, Chinese, Salish, and Mam (Guion 1998: 20, and references therein). What is remarkable about these velar palatalizations is not only their frequency in the world's languages, but also that the output of the rule is not velar, but coronal. This shift in articulation is problematic for a purely coarticulatory account. Coarticulation predicts fronting of the tongue body, producing [kʲ] a palatalized velar, or [c], a pure palatal. But articulatory factors are unable to explain the shift from velar to coronal, which involves a change of articulator: from the tongue body for [k], to the tongue blade for [t͡ʃ].

However, in this case, too, there is evidence that speakers misperceive palatalized velars as coronals with greater-than-chance frequency (Guion 1998). A number of experiments carried out by Guion demonstrate that velar stops before front high vowels are acoustically similar to palatoalveolar affricates and that velar stops in the same contexts are easily confused by listeners with palatoalveolar affricates. An additional finding is that the acoustic and perceptual similarity of voiceless velar stops to palatoalveolars is greater than that for voiced velar stops. In sum, by making reference to perceptual properties of palatalized velars, Guion (1998) is able to explain the high frequency of velar palatalization in the world's languages (in contrast, for example, to labial palatalization), the higher frequency of this change with [k] than [g], the higher frequency of this change with high front vowels than with other vowels, and the shift of articulator from the tongue dorsum to the tongue blade.

In this context, the content of Morris Halle's plenary address ''Moving On'', given in honor of the 80th anniversary of the founding of the Linguistic Society of America in January 2004, is somewhat surprising. In this lecture, Halle, a founder of modern generative phonology, suggests that synchronic velar palatalizations in a range of

[1] Subsequent to this study, Ohala's insights were integrated into phonetically based Optimality accounts, for example, Steriade (2001).

languages are best accounted for in terms of feature geometry (Halle *et al.* 2000). Front vowels, the triggers of velar palatalization, have both a primary designated articulator, the tongue body (dorsal), and a secondary designated articulator, the tongue blade (coronal). Assimilation expressed as feature spreading accounts for velar palatalization by spreading of the coronal node from a high segment to an adjacent velar.

The question which arises is why Guion's (1998) work is not noted by Halle. Given acoustic similarities between velar stops before front high vowels and palatoalveolar affricates and the confusability of these two consonant types by listeners, should one not rethink a purely articulatory and representational account of the same facts? In this case, as with the regressive major place assimilation, there are fine details which the articulator-feature account does not predict: velar palatalization occurs adjacent to other front vowels, though it is most common with high front vowels; and velar palatalization occurs more often with [k] than [g]. The perceptual account finds further support in context-free shifts from pure palatals to palatoalveolar affricates, like that which occurred in the history of Athabaskan (Krauss 1982). So why are these results ignored?

Interested in the answer to this question, I asked Morris Halle myself, after his plenary address, whether he was aware of Guion's work. He said he had looked at it, but that really, this was all about the results of the articulatory model: sound patterns are the way they are because articulation informs phonological universal feature representations. In general, it seems that results of this sort encroach on territory which generativists are accustomed to viewing as part of the synchronic system. When this happens, there is resistance to the idea of incorporating experimental results. One consequence of this policy is an unrealistic conception of synchronic systems which, in many cases, duplicate explanations in other domains.

10.3 Regressive place assimilation in VNCV: mishearing results

In contrast to Halle's seeming disinterest in the misperception literature, there are recent phonological studies which look in detail at similar experimental results. A case in point is the experimental literature on place assimilation in VCCV sequences referred to earlier. Place assimilation of a nasal stop to a following oral stop is a common sound change, and also reflected by alternations in many of the world's languages. Perception studies, including Fujimura *et al.* (1978) and J. Ohala (1990*a*) show a match between misapprehension and sound change. The CV transition dominates the percept, giving rise to a single homorganic interpretation of place for a medial heterorganic sequence. In Fujimura *et al.*'s experiment, homorganicity correlated with duration of consonantal interlude, while in Ohala's Experiment 1, non-homorganic sequences like [VŋpV], [VnpV] were judged as homorganic 93 percent of the time. Nevertheless, Steriade (2001: 232–3) and Hayes and Steriade

(2004: 26–7) doubt that innocent misapprehension alone is capable of driving systematic phonological change in nasal place assimilation. What is the source of this doubt?

Steriade (2001) cites the results of Hura *et al*. (1992). In this experiment, where heterorganic VNCV sequences were presented to English speakers (N = [m, n, ŋ] and C = [p, t, k]), nasals showed an error rate of 6.9 percent, significantly higher than the 3 percent error rate for fricatives in the same position. The result of interest to Steriade, however, was that most errors were non-assimilatory. For example, 76.1 percent of all incorrect nasal responses were /n/. If listeners had been responding based on the place of articulation of the following stop, responses would be balanced among the three distinct places of articulation. Steriade takes these results as rejection of the general hypothesis that misperception explains regressive nasal-place assimilation. She argues instead that misperception of nasal place in VNCV may be common, but that assimilation (or neutralization to [n]) in this context is the consequence of optimization which characterizes synchronic grammars. She attributes to the speaker the knowledge that the nasal in VNCV is likely to be misperceived. Given this, an ''unmarked'' nasal is produced instead.

Steriade's interpretation of this particular experiment is surprising, since the authors themselves advise against this conclusion:

> it would be a mistake to reject Ohala's hypothesis [of perceptually based assimilation; JB] on the basis of our results, because the VC_1C_2V intervals used in our experiment appear to have been longer than the duration at which perceptual assimilation errors typically occur (Repp 1978; J. Ohala 1990*a*). In other words, our experiment was not designed to provide a clear test of Ohala's hypothesis. (Hura *et al*. 1992: 69)

Not only were intervocalic consonantal intervals long in this study, but the stimulae were made from a set of nonsense names, with N##C sequences spanning the end of the first name, and the beginning of the last name. In the case of final nasals, the first names were *Shanim*, *Shanin*, and *Shaning*, while the last names were *Perry, Terry*, and *Kerry.* The fact that 76.1 percent of errors in perception of nasals involved hearing *Shanim* or *Shaning* as *Shanin* [ʃanɪn] may be due, not to the default status of /n/ generally, but to very specific facts about the English lexicon, such as the existence of names like *Shannon* and *Sharon*, in the same general phonological neighborhood as the ambiguous tokens, and the fact that disyllabic names with initial stress ending in lax-vowel + nasal sequences are much more likely to end in /n/ than in a non-coronal nasal. Compare *Aaron*, *Alan*, *Brian*, *Brendon*, *Dustin*, *Dylan*, *Evan*, *Ivan*, *Jasmine*, *Karen*, *Kevin*, *Kristen*, *Lauren*, *Logan*, *Martin*, *Megan*, *Morgan*, *Stephen*, etc. vs. much less frequent /m/- and /ŋ/-final names like *Adam* or *Henning*.[2]

[2] In 1992, the year the Hura *et al*. study was published, the top 100 baby names in the USA. included 31 distinct non-homophonous disyllabic names ending in unstressed /Vn/ (boys: Brandon, Ryan, Justin, Kevin, Steven/Stephen, Dylan, Aaron, Brian/Bryan, Jordan, Christian, Austin, Nathan, Jason, Cameron, Dustin, Evan, Dillon, Devin, Ethan, Logan. Girls': Megan, Lauren, Jasmine, Morgan, Kaitlyn/Caitlin, Erin, Kristen, Kathryn, Jordan, Shannon, Meghan, Kristin), and only two (William, Adam), ending in unstressed vowel + non-coronal nasal (from www.popular-baby-names.org). The bigger the English

In this particular case, reference to synchronic markedness constraints reflects failure to identify and isolate different sources of phonological knowledge: the knowledge of general English phonotactics, and the knowledge of purely contingent patterns involving proper names in the English lexicon.

10.4 VOWEL REDUCTION: AN IMAGINARY PARADOX?

Steriade (2001) is not alone in attributing recurrent sound patterns to teleological constraints invoking enhancement of contrast or minimization of effort. Crosswhite (2001, 2004) looks at vowel reduction in a wide range of languages and argues that there are two distinct types of vowel reduction, with distinct phonetic teleologies. Contrast-enhancing reduction is vowel reduction where non-peripheral vowels within the psycho-acoustic vowel space are avoided in unstressed positions. Since unstressed vowels are more likely to be misperceived than stressed vowels, Crosswhite argues, this type of reduction serves to minimize perceptual confusion by enhancing contrast. A second type of vowel reduction is classified as "prominence reduction". Under this type of reduction, long or otherwise salient vowel qualities are avoided in unstressed syllables.

Crosswhite (2004: 191) claims that this bipartite typology is key to explaining the empirical facts of "reduction paradoxes"—her term for the phenomenon that, cross-linguistically, a particular vowel quality is frequently subject to reduction while at the same time the same vowel quality is frequently the output of reduction.[3] A pair of examples will serve to illustrate the paradox. In Belarusian, unstressed /e, o/ reduce to [a], while unstressed /i, u, a/ undergo no phonemic shifts. In Standard Bulgarian, however, with underlying vowels /i, u, e, o, ə, a/, unstressed /e, o/ raise to [i,u], while unstressed /a/ raises to [ə]. The paradox, Crosswhite suggests, is that [a] is the output of reduction in one language, but the target of reduction in another.

However, nowhere in her work does Crosswhite ask what to many might be an obvious question, namely, Is there a paradox here? Beckman, de Jong, Jun, and Lee (1992) sketch what appears to be a reasonable and well-supported phonetic explanation for prosodically conditioned vowel reduction. On their account:

given-name database, the more skewed ratios are. For example, the top 500 baby names of 2002 show 150 disyllabic names ending in unstressed /Vn/, with only five ending in unstressed V + non-coronal (William, Adam, Liam, Tatum, Malcolm).

Pam Beddor (pers. comm., 2006) suggests another factor which should be taken into account in comparing Ohala's (1990*a*) results with those of Hura *et al*. (1992): range of listener choice. Ohala's study offered three choices, two homorganic and one "other", which could have encouraged homorganic responses. This contrasts with the Hura *et al*. study, where listeners were simply asked to identify the first name in each two word sequence. For a different experimental design with assimilation rates intermediate between these two studies, see Beddor and Evans-Romaine (1995).

[3] Of course, it could be that the same phonetic vowel quality is not involved, which would also render the paradox non-existent.

any prosodic effect that increases the gestural overlap, or that decreases the acoustic salience of an overlapped gesture, would increase the likelihood of a listener reinterpreting the coarticulation as an intentional feature of the affected phoneme segment. For example, an unstressed vowel might be very short, so that a greater proportion of its dorsovelar gesture overlaps with the preceding consonant. A listener might misinterpret this resulting coarticulation as an intentional vagueness about the vowel's quality—that is, an underlying full vowel might be replaced with /ə/ (Beckman, de Jorg, Jun, and Lee 1992: 48)

Experimental support for this account is found in a study of jaw kinematics, and more general studies of undershoot. As noted by Beckman *et al.* (1992), study of jaw movements in accented vs. unaccented vowels in closed syllables is consistent with an interpretation of "truncation" effects, where a consonantal gesture cuts off the oral gesture for the vowel before the jaw reaches its target. A study of /pɑp/ sequences by Edwards *et al.* (1991) shows that the jaw moves a shorter distance when lowering for shorter unaccented vowels than for longer accented ones, but that the joint kinematic movements are not those associated with a higher target for the unaccented vowels. Beckman *et al.* (1992: 48) also note that Lindblom's (1963) early model of target undershoot as the source of short vowel reduction had good empirical coverage, despite its basis in temporally invariant movements. Taking other facts about gestural organization and timing into account, truncation of closely phased gestures "seems to account for a great deal of vowel reduction synchronically."

Further support for non-teleological models of vowel reduction can be found in the effect of linguistic experience on perceptual similarity and notions of prominence. The paradox suggested by Crosswhite is non-existent if it can be shown that reduced mid vowels /e, o/ can be heard as [a] in one language, while reduced /a/ can be heard as [ə] in another. While I know of no study testing precisely this combination of perception facts, a growing number of studies show that perceptual similarity judgments can reflect language-specific phonetics of one's native language (Bladon and Lindblom 1981; Flege *et al.* 1991; Mielke 2003). More importantly, perhaps, there are significant differences across languages and dialects in the extent to which temporal information is used to distinguish vowel quality. Studies demonstrating that changes in vowel duration give rise to percepts of different qualities include Ainsworth (1972), Mermelstein (1978), Strange (1989), Whalen (1989), and Miller and Grosjean (1997). Finally, there is evidence of cross-speaker differences in the categorization of reduced vowels highlighting difficulties for speakers in distinguishing between acoustic and lexical reduction (van Bergem 1993).

10.5 Constraints on strong positions: misconceptions of the model

Another area where diachronic non-teleological explanations have been claimed to be inadequate is in the phonology of prosodic prominence. J. Smith (2004*a*, *b*) argues that the synchronic phonology of prosodically prominent positions is generally incompatible with diachronic phonetic explanations. The two arguments presented in J. Smith (2004*b*) are that (i) the addition of perceptual salience to a strong position is not the kind of sound change that is expected to result from misperception, and (ii) that the relationship between strong positions and salient properties found in augmentation is too abstract to have reliable origins in the acoustic signal.

However, this work simplifies models like Evolutionary Phonology (Blevins 2004), in which phonologization of misperception is just one of multiple mechanisms leading to regular sound change and regular sound patterns. Within Evolutionary Phonology, as in traditional Neogrammarian models, variability along the hyper-to-hypoarticulation continuum provides the exemplar space from which new phonological representations can emerge. The rigid associations in some languages between, for example, a main stress and a heavy syllable (Aguacatec, Yurok) or between a main stress and high tone (Serbo-Croatian, Slave) are expected, given the range of variability in the realization of stress before such patterns are phonologized. For example, variation in syllable duration, pitch contour and amplitude have been observed under emphatic stress (e.g. Dahan and Bernard 1996), and can lead to phonological associations between word stress and distinct types of prominence. As Gordon (2002) demonstrates, a general phonetic feature characterizing most types of "heavy syllable" is greater acoustic energy. Smith's argument, then, is essentially a straw man. No one has claimed that misperception characterizes fortition, or that more specifically, that properties associated with phonologically prominent syllables must or typically originate in innocent misperception. Fortition and lenition are processes that occur at opposite ends of the hyper-hypo-articulation continuum, associated with durational expansion and contraction of the utterance as well as more and less forceful articulations respectively.

J. Smith's (2004*b*) second argument for synchronic markedness constraints over historical phonetic modeling of sound patterns involves languages in which roots are stressed. The claim is that since roots have no intrinsic relation to stress, this relationship cannot have evolved through misperception. Again, it is unclear whose claims Smith is contesting. No one, to my knowledge, has ever claimed that root stress is a direct consequence of misperception. Historical studies of the evolution and movement of stress are not numerous, and there seems to be just as many cases of languages where stress shifts between roots and suffixes (e.g. Lithuanian, Japanese), as those mentioned by J. Smith (2004*a*), in which roots are always stressed (e.g. Mbabaram, Tahltan). With respect to this last class, it seems significant that proto-Pama-Nyungan, the ancestral language of Mbabaram, is reconstructed with root-stress. If this is a directly inherited trait, then it is unclear why a synchronic account is necessary for this particular case. Finally, consider the simplest models of grammaticization. A free morpheme becomes a clitic which subsequently becomes an affix: if affixes are fossil clitics, and stress is originally a property of lexemes, then, all else being equal, a pattern of root/stem stress will emerge.

The reductionism found in Smith's arguments is typical of modern generative and post-generative traditions. These approaches show a strong bias to treat recurrent sound patterns as reflections of synchronic markedness or naturalness constraints. But there is continued recognition in the history of phonology, from the Neogrammarians to the school of experimental phonology practised by Ohala and others, that there are multiple sources of recurrent sound patterns. Recurrent sound patterns can be the result of direct inheritance among genetically related languages; they can be a consequence of recurrent phonetically motivated sound change, or they can be the result of contact-induced change. Since explanations for the majority of recurrent sound patterns in the world's languages do not require reference to synchronic markedness constraints (Blevins 2004), the burden of proof lies squarely with those invoking such constraints.[4]

10.6 PHONETIC KNOWLEDGE, PHONOLOGICAL KNOWLEDGE, AND REALISTIC GRAMMARS

Having come to some understanding of sources of resistance in the literature to innocent misperception as the source of regular sound patterns, and to its implication that sound change is non-teleological, we can turn to more constructive questions. Once diachronic phonetic explanation is excised from synchronic phonology, what will grammars look like? What architecture characterizes the description of sounds and sound patterns, and what types of experimental evidence and methods are most likely to shed light on the content of phonological knowledge? Before answering these questions, it will be useful to dispel three common misconceptions of grammatical models in which phonetic explanation is confined primarily to the diachronic component.

One common misconception is that the existence of explanation in the diachronic dimension is only illusory, since language change itself reflects constraints on synchronic grammars. This view is most succinctly stated by Joseph and Janda (1988) and taken by others (e.g. Hume and Johnson 2001*b*) as a cogent argument for importing phonetic explanations into synchronic grammars. However, when we look closely at the structure of the argument, we can see where it goes wrong:

> Diachrony is best viewed as the set of transitions between successive synchronic states, so that language change is necessarily something that always takes place in the present and is therefore governed in every instance by constraints on synchronic grammars. (Joseph and Janda 1988: 194)

The set of transitions between successive synchronic states is discontinuous, involving an initial state, where the newborn does not have an identifiable grammar of a language, and a final state where the young child or adult does. The claim of models like Evolutionary Phonology is that the majority of regular sound changes have seeds

[4] One recent study attempting to provide positive proof of synchronic markedness constraints is Wilson (forthcoming).

in misperception, resolution of ambiguity, and frequency-based choice of ''best exemplar'' and that these transforms take place in the course of language acquisition (Blevins 2004: 32–44; Blevins forthcoming).[5] The fact that this acquisition takes place in the present does not mean that change must be governed by constraints on synchronic grammars. On the contrary, if the content of synchronic grammar is what is being discovered in the course of acquisition, then it cannot play a primary role in acquisition. This is the view taken by Lindblom (2000) and Wedel (2004*a*, *b*), for instance, where formal properties of sound patterns are modeled as emergent structures formed in self-organizing ways through the feedback of the perception–production loop in the course of language acquisition.

Two other misconceptions are common regarding localization of phonetic explanation in the diachronic dimension. One is that synchronic systems no longer characterize or incorporate phonetic knowledge.[6] Another is that, with phonetic explanations excised from synchronic grammars, there is nothing left for synchronic systems to characterize. These two misconceptions may stem from the failure properly to distinguish between phonetic and phonological knowledge, explanation and description, or between innate versus learned knowledge of language. Once phonetic explanations are excised from synchronic phonological systems, these systems can describe purely and systematically all and only the phonological knowledge for which speakers show positive evidence, via natural linguistic behavior or experimental data. This phonological knowledge may have a very different character from structuralist, generativist, and optimality conceptions, and explanations for the structure of this knowledge may require deeper understanding of association and categorization strategies of humans more generally (Johnson, this volume). For example, Ernestus and Baayen (2003) show that Dutch speakers interpret neutralized final devoiced segments in Dutch as voiced or voiceless by making use of phonological/phonetic similarity patterns in the lexicon, with new words interpreted in such a way as to conform to these learned patterns A synchronic grammar of Dutch should be able to characterize knowledge of phonological/phonetic similarity in the sense that it is used by speakers in this particular experiment. In this case, the resulting grammar will be highly descriptive, since word-forms are the basis of analogical generalizations. At the same time, a model of this sort highlights the extent to which knowledge of phonetics and phonology is learned knowledge of language, since the basis of emergent analogical generalizations are learned sound patterns of individual words.

What models of synchronic phonology have this degree of descriptive detail, adopt learning as a primary mechanism, and treat phonetic explanations as primarily

[5] This model accepts that ''mature'' phonologies change over time, as shown by among others Labov (1994, 2001) and Harrington *et al*. (2000). Indeed, exemplar models (e.g. Pierrehumbert 2001*a*; Johnson, this volume) predict continued change, though at slower rates than initial acquisition transforms, as entrenched system-internal phonological categories have stronger and stronger feedback effects.

[6] I refer here to language-specific phonetic knowledge in the sense of Kingston and Diehl (1994), as well as universal aspects of phonetics, for example, the more sonorous percept of [a] in contrast to [i] implicated in universal sonority scales (cf. Blevins 1995).

historical? One which immediately comes to mind is Stochastic Phonology or Probabilistic Phonology (Pierrehumbert 2001*b*, 2003*b*; Johnson, this volume). In Stochastic Phonology, phonetic knowledge is represented in detail. Frequencies of sound patterns play a crucial role in the acquisition of phonological and phonetic competence, and it is precisely this competence which one attempts to model. At the phonetic level, exemplar theory provides one model of how probability distributions over cognitive maps may be used in speech perception and production (Johnson 1997*b*, this volume; Pierrehumbert 2001*a*, 2003*b*). Modeling of the lexicon is most accurately viewed in terms of stronger and weaker connections between words with more and less shared properties. Finally, in Stochastic Phonology, the actual grammar provides a very concrete tracking of generalizations over the lexicon. The form and content of these generalizations are addressed in work on analogical modeling, from the formal work of Skousen (1989, 1992), and the experimental studies like Ernestus and Baayen (2003), to the computational modeling of Wedel (2004*a*), for instance. Since, as summarized by Pierrehumbert (2001*b*), the lexical network and the cognitive map each explain "a large and diverse battery of findings about implicit knowledge of speech, and no viable alternative has been proposed for either concept," it is surprising that not more theories of grammar take lexical networks and cognitive maps as architectural starting points for the characterization of phonetic and phonological knowledge.

The growing experimental evidence, then, suggests that synchronic phonological grammars are not the domain of phonetic explanation. This is not surprising. Large synchronic systems were proposed as part of an explanation for regular patterns for which no other explanation appeared to be available. But expansive synchronic grammars were a means to an end, not an end in themselves. As compelling phonetic historical explanations for sound patterns have become available, the need to invoke synchronic phonological explanations has receded.

In this chapter I have focused on two of the most important implications of Ohala's research for phonological modeling: that innocent misperception can lead directly to attested recurrent sound patterns; and that sound change is non-teleological. However, the greatest achievement of experimental phonology is not simply being able to account for patterns of sound change, but to provide accounts that make real predictions outside the domain of direct investigation. Validating these predictions reinforces the value of the experimental method, and confirms the discovery of genuine explanations rather than *post facto* descriptions.