

Phonetically Driven Phonology:
The Role of Optimality
Theory and Inductive Grounding¹

Bruce P. Hayes

UCLA

October 1996

Written for the proceedings volume of the 1996 Milwaukee Conference
on Formalism and Functionalism in Linguistics

Abstract

Functionalist phonetic literature has shown how the phonologies of human languages are arranged to facilitate ease of articulation and perception. The explanatory force of phonological theory is greatly increased if it can directly access these research results. There are two formal mechanisms that together can facilitate the link-up of formal to functional work. As others have noted, Optimality Theory, with its emphasis on directly incorporating principles of markedness, can serve as part of the bridge. Another mechanism is proposed here: an algorithm for *inductive grounding* permits the language learner to access the knowledge gained from experience in articulation and perception, and form from it the appropriate set of formal phonological constraints.

¹ I would like to thank Marco Baroni, Patricia Keating, Donca Steriade, and members of talk audiences at MIT, Arizona, Tilburg, Utrecht, Cornell, UC Santa Cruz, UCLA, and the Formalism/Functionalism conference for helpful input in the preparation of this paper.

1. Phonological Functionalism

The difference between formalist and functionalist approaches in linguistics has taken different forms in different subfields. For phonology, and particularly for the study of fully-productive sound patterns, the functionalist approach has traditionally been phonetic in character. For some time, work in the phonetic literature, such as Ohala (1974, 1978, 1981, 1983), Ohala and Ohala (1993), Liljencrants and Lindblom (1972), Lindblom (1983), and Westbury and Keating (1986), has argued that the sound patterns of languages are effectively arranged to facilitate ease of articulation and distinctness of contrasting forms in perception. In this view, much of the patterning of phonology reflects principles of good design.

In contemporary phonological theorizing, such a view has not been widely adopted. Phonology has been modeled as a formal system, set up to mirror the characteristic phonological behavior of languages. Occasionally, scholars have made a nod towards the phonetic sensibleness of a particular proposal. But on the whole, the divide between formal and functionalist approaches in phonology has been as deep as anywhere else in the study of language.

It would be pointless (albeit fun) to discuss reasons for this based on the sociology of the fields of phonetics and phonology. More pertinently, I will claim that part of the problem has been that phonological theory has not until recently advanced to the point where a serious coming to grips with phonetic functionalism would be workable.

2. Optimality Theory

The novel approach to linguistic theorizing known as Optimality Theory (Prince and Smolensky 1993) appears to offer the prospect of a major change in this situation. Here are some of the basic premises of the theory as I construe it.

First, phonological grammar is not arranged in the manner of Chomsky and Halle (1968), in essence as an assembly line converting underlying to surface representations in a series of steps. Instead, the phonology *selects* an output form from the set of logical possibilities. It makes its selection using a large set of constraints, which specify what is “good” about an output, in the following two ways:

- (1)a. Phonotactics: “The output should have phonological property X.”
- b. Faithfulness: “The output should resemble the input for property Y.”

Phonotactic constraints express properties of phonological markedness, which are typically uncontroversial. For example, they require that syllables be open, or that front vowels be unrounded, and so on. The Faithfulness constraints embody a detailed factorization of what it means for the output to resemble the input; they are fully satisfied when the output is identical to the input.

Constraints can conflict with each other. Often, it is impossible for the output to have the desired phonotactic properties and also be faithful to the input; or for two different phonotactic constraints to be satisfied simultaneously. Therefore, all constraints are prioritized; that is, ranked. Prioritization drives a specific winnowing process (not described here) that ultimately selects the output of the grammar from the set of logical possibilities by ruling out all but a single winner.²

I will take the general line that Optimality Theory is a good thing. First, it shares the virtues of other formal theories: when well implemented, such theories provide falsifiability, so that the errors in an analysis can lead to improvement or replacement. Further, formal theories characteristically increase the *pattern recognition capacity* of the analyst. For example, it was only when the formal theory of moras was introduced (Hyman 1985) that it became clear that compensatory phonological processes always conserve mora count (see Hyman, and for elaboration Hayes 1989).³

Second, Optimality Theory has permitted solutions to problems that simply were not treatable in earlier theories. Examples are the metrical phonology of Guugu Yimidhirr (Kager, to appear), or the long-standing ordering paradoxes involving phonology and reduplication (McCarthy and Prince 1995).

Most crucially, Optimality Theory has the advantage of allowing us to incorporate general principles of markedness into language-specific analyses. Previously, a formal phonology consisted of a set of somewhat arbitrary-looking rules. The analyst could only look at the rules “from the outside” and determine how they reflect general principles of markedness (or at best, supplement the rules with additional markedness principles, as in Chomsky and Halle (1968, Ch. 9), Schachter (1969), or Chen (1973)). Under Optimality Theory, the principles of markedness (stated explicitly and ranked) form the *sole ingredients* of the language-specific analysis. The mechanism of selection by ranked constraints turns out to be such an amazingly powerful device that it can do all the rest. Since rankings are the only arbitrary element in the system, the principled character of language-specific analyses is greatly increased. This is necessarily an argument by assertion, but I believe a fair comparison of the many phonological analyses of the same material in both frameworks would support it.⁴

² For the original presentation of Optimality Theory, the reader is referred to Prince and Smolensky (1993) or to McCarthy and Prince (1994) for a recent summary.

³ Here are two other examples. (a) The formal theory of intonational representation developed by Pierrehumbert (1980) led her to discover English intonational contours not previously noticed. (b) The theory of prosodic domains of Selkirk (1980) led Hayes and Lahiri (1991) to find close links between intonation and segmental phonology in Bengali that would not have otherwise been observed.

⁴ Much of the Optimality Theory analytical literature is currently posted on the World Wide Web at <http://ruccs.rutgers.edu/roa.html>, and may be downloaded.

3. What is a Principled Constraint?

The question of what makes a constraint “principled” is one that may be debated. The currently most popular answer, I think, relies on typological evidence: a principled constraint is one that “does work” in many languages, and does it in different ways.

But there is another answer to the question of what makes a constraint principled: a constraint can be justified on *functional* grounds. In the case of phonetic functionalism, a well-motivated phonological constraint would be one that either renders speech easier to articulate or renders contrasting forms easier to distinguish perceptually. From the functionalist point of view, such constraints are *a priori* plausible, under the reasonable hypothesis that language is a biological system that is designed to perform its job well and efficiently.

Optimality Theory thus presents a new and important opportunity to phonological theorists. Given that the theory thrives on principled constraints, and given that functionally motivated phonetic constraints are inherently principled, the clear route to take is to explore how much of phonology can be constructed on this basis. One might call such an approach “Phonetically-Driven Optimality-theoretic Phonology.” A theory of this kind would help close the long-standing and regrettable gap between phonology and phonetics.

4. Sample Result: Segment Licensing

The position just taken regarding phonetics and Optimality Theory is not original with me, but is inspired by ongoing research of Steriade (1993, 1995b, in progress), as well as work such as Archangeli and Pulleyblank (1994), Jun (1995a,b), Flemming (1995), Kaun (1995a,b), Silverman (1995), Myers (1996), Gordon (forthcoming), and Kirchner (in progress). Although I obviously cannot review all of this work, I will mention one particularly vivid example.

Steriade (1995b, in progress) considers the very basic question of segmental phonotactics in phonology: what segments are allowed to occur where? Her perspective is a novel one, taking the line that *perception* is the dominant factor. Roughly speaking, Steriade suggests that segments preferentially occur where they can best be heard. The crucial part is that many segments (for example, voiceless stops) are rendered audible largely or entirely by the contextual acoustic cues that they engender on neighboring segments through coarticulation. In such a situation, it is clearly to the advantage of particular languages to place strong restrictions on the phonological locations of such segments.

Following this approach, and incorporating a number of results from research in speech perception, Steriade is able to reconstruct the traditional typology of “segment licensing,” including what was previously imagined to be an across-the-board preference for consonants to occur in syllable onset position. She goes on to show that there in fact

are areas where this putative preference fails as an account of segmental phonotactics: one example is the preference for retroflexes to occur postvocalically (in either onset or coda); preglottalized sonorants work similarly. As Steriade shows, these otherwise-baffling cases have specific explanations, based on the peculiar acoustics of the segments involved. She then makes use of Optimality Theory to develop explicit formal analyses of the relevant cases.

5. The Hardest Part

What is crucial here (and recognized in Steriade's work) is that a research result in phonetics is not same as a phonological constraint. To go from one to the other is to bridge a large gap. Indeed, the situation facing Phonetically-Driven Optimality-theoretic Phonology is a rather odd one. In many cases, the phonetic research that explains the phonological pattern has been done very well and is quite convincing; it is only the question of how to incorporate it into a formal phonology that is difficult. An appropriate model for the research program described here is: *we seek to go beyond mere explanation to achieve actual description.*

In what follows, I will propose a particular way to attain phonetically-driven phonological description.⁵ Since I presuppose Optimality Theory, what is crucially needed is a means to obtain phonetically-motivated constraints.

In any functionalist approach to linguistics, an important question to consider is: *who is in charge?* That is, short of divine intervention, languages cannot become functionally well designed by themselves; there has to be some mechanism responsible. In the view I will adopt, phonology is claimed to be phonetically natural because the constraints it includes are (at least partially) the product of *grammar design*, carried out intelligently (that is, unconsciously, but with an intelligent algorithm) by language learners.

Before turning to this design process, I will first emphasize its most important aspect: there is a considerable gap between the raw patterns of phonetics and phonological constraints. Once the character of this divergence is clear, then the proposed nature of the design process will make more sense.

6. Why Constraints Do Not “Emerge” From The Phonetics

There are a number of reasons that suggest that phonetic patterns cannot serve as a direct, unmediated basis for phonology. For more discussion of this issue, see Anderson (1981) and Keating (1985).

⁵ Since this is only one approach among many, the reader is urged to compare it with Steriade's work, as well as Flemming (1995) and Kirchner (in progress).

6.1 Variation and Gradience

First, phonetics involves gradient and variable phenomena, whereas phonology is characteristically categorial and far less variable. Here is an example: Hayes and Stivers (1996) set out to explain phonetically a widespread pattern whereby languages require postnasal obstruents to be voiced. The particular mechanism we propose is reviewed below; for now it suffices that it appears to be verified by quantitative aerodynamic modeling and should be applicable in any language in which obstruents may follow nasals.

Since the mechanism posited is automatic, we might expect to find it operating even in languages like English that do not have postnasal voicing as a phonological process. Testing this prediction, Hayes and Stivers examined the amount of closure voicing (in milliseconds) of English /p/ in the environments / m ___ versus / r ___. Sure enough, for all five subjects in the experiment, there was significantly more /p/ voicing after /m/ than after /r/, as our mechanism predicted. But the effect was *purely quantitative*: except in the most rapid and casual speech styles, our speakers fully succeeded in maintaining the phonemic contrast of /p/ with /b/ (which we also examined) in postnasal position. The phonetic mechanism simply produces a quantitative distribution of voicing that is skewed toward voicedness after nasals. Moreover, the distribution of values we observed varied greatly: the amount of voicing we found in /mp/ ranged from 13% up to (in a few cases) over 60% of the closure duration of the /p/.

In contrast, there are other languages in which the postnasal voicing effect is truly phonological. For example, in Ecuadorian Quechua (Orr 1962), at suffix boundaries, it is phonologically illegal for a voiceless stop to follow a nasal, and voiced stops are substituted for voiceless; thus *sača-pi* ‘jungle-loc.’ but *atam-bi* ‘frog-loc.’ For suffixes, there is no contrast of voiced versus voiceless in postnasal position. Clearly, English differs from Quechua in having “merely phonetic” postnasal voicing, as opposed to true phonological postnasal voicing.⁶ We might say that Ecuadorian Quechua follows a categorial strategy: in the suffix context it simply doesn’t even try to produce the (phonetically difficult) sequence *nasal + voiceless obstruent*. English follows a “bend but don’t break” strategy, allowing a highly variable increase in degree of voicing after nasals, but nevertheless maintaining a contrast.

I would claim then, that in English we see postnasal voicing “in the raw,” as a true phonetic effect, whereas in Ecuadorian Quechua the phonology treats it as a categorial phenomenon. The Quechua case is what needs additional treatment: it is a kind of leap from simply allowing a phonetic effect to influence the quantitative outcomes (in a variable

⁶ The difference is clearly reminiscent of the notion of “phonologization” discussed in Hyman (1976) and earlier, though Hyman’s main focus is on historical contrast redistributions such as those found in tonogenesis.

way) to arranging the phonology so that an entire contrast is wiped out.⁷

6.2 Symmetry

Let us consider a second argument. I claim that phonetics is asymmetrical, whereas phonology is usually symmetrical. Since the phonetic difficulty of articulation and perception follows from the interaction of complex physical and perceptual systems, we cannot in the general case expect the regions of phonetic space characterized by a particular difficulty level to correspond to phonological categories.

To make this clear, consider a particular case, involving the difficulty of producing voiced and voiceless stops. The basic phonetics (here, aerodynamics) has been studied by Ohala (1983) and by Westbury and Keating (1986). Roughly, voicing is possible whenever a sufficient drop in air pressure occurs across the glottis. In a stop, this is a delicate matter for the speaker to arrange, since free escape of the oral air is impeded. Stop voicing is influenced by quite a few different factors, of which just a few are reviewed here.

(a) *Place of articulation.* In a “fronter” place like labial, a large, soft vocal tract wall surface surrounds the trapped air in the mouth. During closure, this surface retracts under increasing air pressure, so that more incoming air is accommodated. This helps maintain the transglottal pressure drop. Since there is more yielding wall surface in labials (and more generally, at fronter places of articulation), we predict that the voiced state should be relatively easier for fronter places. Further, since the yielding-wall effect actually makes it harder to turn off voicing, we predict that voicelessness should be harder for fronter places.

(b) *Closure duration.* The longer a stop is held, the harder it will be to accommodate the continuing transglottal flow, and thus maintain voicing. Thus, voicelessness should be favored for geminates and for stops in post-obstruent position. (The latter case assumes that, as is usual, the articulation of the stop and the preceding obstruent are temporally overlapped, so no air escape can occur between them.)

(c) *Postnasal position.* As just noted, there are phonetic reasons why voicing of stops should be considerably favored when a nasal consonant immediately precedes the stop.

⁷ Actually, this paragraph slights the complexity of phonetic implementation. Following Pierrehumbert (1980) and Keating (1985), I assume that there is also a phonetic component in the grammar, which computes physical outcomes from surface phonological representations. It, too, I think, is Optimality-theoretic and makes use of inductive grounding (below). I cannot address these issues here for lack of space.

(d) *Phrasal position*. Characteristically, voicing is harder to maintain in utterance-initial and utterance-final position, since the subglottal pressure that drives voicing tends to be lower in these positions.

As Ohala (1983) and others have made clear, these phonetic factors are abundantly reflected in phonological patterning. (a) Gaps in stop inventories that have both voiced and voiceless series typically occur at locations where the size of the oral chamber makes voicing or voicelessness difficult; thus at *[p] or *[g], as documented by Ferguson (1975), Locke (1983), and several sources cited by Ohala (p. 195). (b) Clusters in which a voiced obstruent follows another obstruent are also avoided, for instance in Latin stems (Devine and Stephens 1977), or in German colloquial speech (Mangold 1962, 45). Geminate obstruents are a similar case: they likewise are often required to be voiceless, as in Japanese (Vance 1987, 42), West Greenlandic (Rischel 1974), or !Xõõ (Traill 1981, 165). (c) Languages very frequently ban voiceless stops after nasals, with varying overt phonological effects depending on how the constraints are ranked (Pater 1995, 1996; Hayes and Stivers 1996). (d) Voicing is favored in medial position, and disfavored in initial and final position, following the subglottal pressure contour (Westbury and Keating 1986).⁸

Plainly, the phonetics can serve here as a rich source of phonological explanation, since the typology matches the phonetic mechanisms so well. However, if we try to do this in a naive, direct way, difficulties immediately set in.

Suppose that we concoct a *landscape of stop voicing difficulty* (2) which encodes difficulty values for a set of phonological configurations. For simplicity, we will consider only a subset of the effects mentioned above.

⁸ Interestingly, Westbury and Keating's (1986) modeling work found no articulatory support for the large typological difference between final devoicing (ubiquitous) and initial devoicing (somewhat unusual; see Westbury and Keating for cases). Recent work by Steriade (in progress) that relates the phonology of voicing to its *perceptual* cues at consonant releases would appear to fill this explanatory gap.

(2) Landscape of Difficulty for Voiced Stops: Three Places, Four Environments

	b	d	g	
[-son] ____	43	50	52	
# ____	23	27	35	
[+son, -nas] ____	10	20	30	
[+nas] ____	0	0	0	contour line: 25

The chart in (2) was constructed using a software aerodynamic vocal tract model implemented at UCLA (Keating 1984). The basis of the chart is explained below in section 10; for now, it may be considered simply a listing of “difficulty units” for voicing in various phonological configurations. It can be seen that the model has generated patterns that are qualitatively correct: the further back in the mouth a place of articulation is, the harder it is to maintain voicing. Moreover, the rows of the chart reflect the greater difficulty of maintaining voicing after obstruents and initially, as well as the greater ease after nasals.

What is crucial about the chart is that it reflects the *trading relationships* that are always found in the physical system for voicing. One cannot say, for example, that velars are always harder to voice, because velars in certain positions are easier to voice than labials in others. Similarly, the environments / # ____ versus / [+son, -nas] ____ do not define a consistent cutoff in voicing difficulty, since [g] in the environment / [+son, -nas] ____ is harder than [b,d] in the environment / # ____.

The dotted line on the chart represents a particular “contour line” for phonetic difficulty,” analogous to a contour line for altitude on a physical map. A language that truly “wanted” to behave in a phonetically rational way might ban all phonological configurations that exceeded the contour line, as in (3a). Translating this particular contour line into descriptive phonological language, we have the formulation of (3b):

(3) A Hypothetical Phonological Constraint

- a. *any voiced stop that characteristically requires more than 25 units of effort
- b. *post-obstruent voiced stops,
 - *[d,g] in initial position,
 - *[g] after oral sonorants

Note that [g] is permitted by (3), but only postnasally.

I would contend that a constraint like (3) (however formulated) is relatively unlikely to occur in a real phonology. What occurs instead are constraints that are likewise phonetically sensible, but which possess formal symmetry. Here are some real-world examples, with the languages they are taken from:

- (4)a. *Voiced obstruent word-finally (Polish)
- b. *Voiced obstruent after another obstruent (Latin)
- c. *Voiced obstruent geminate (Japanese)
- d. *Voiced velar obstruents (Dutch)

These constraints ban symmetrical regions of phonological space, not regions bounded by contour lines of phonetic difficulty. Nevertheless, they are phonetically sensible in a certain way: in the aggregate, the configurations that they forbid are more difficult aerodynamically than the configurations that they allow. Thus constraints like (5) would be quite unexpected:

- (5)a. *Voiceless obstruent word-finally (compare (4a))
- b. *Voiceless obstruent after another obstruent (compare (4b))
- c. *Voiceless obstruent geminate (compare (4c))
- d. *Voiceless velar obstruents (compare (4d))

To generalize: I believe that constraints are typically natural, in that the set of cases that they ban is phonetically harder than the complement set. But the “boundary lines” that divide the prohibited cases from the legal ones are characteristically storable in rather simple terms, with a small logical conjunction of feature predicates. In other words, phonological constraints tend to ban phonetic difficulty in simple, formally symmetrical ways (cf. Kiparsky 1995, 659). The constraint (3) is very sensible phonetically, but apparently too logically complex to appear in natural languages (or, at least, in more than a very few of them).

A further demonstration makes this point in a different way. Consider first that Egyptian Arabic (Harrell et al. 1963) bans the voiceless bilabial stop [p]. This is both phonetically sensible and empirically ordinary, as noted above. What is very striking about the ban, however, is that it extends even to geminates: Cairene has words like [yikubb] ‘he spills’, but no analogous words like *[yikupp].⁹ As noted earlier, voiced obstruent geminates are cross-linguistically rare, for good phonetic reasons.

A near-minimal comparison with Arabic is Japanese, which (some unassimilated borrowings aside) is one of the languages that bans voiced obstruent geminates. Since Japanese has [pp] but not [bb], there is an interesting contradiction: in Arabic [bb] is well formed and [pp] is ill formed, whereas in Japanese it is just the opposite.

⁹ It is fairly safe to infer that this is not just idealized phonemic transcription on Harrell et al.’s part, since elsewhere they do record allophonic [p] resulting from a process of regressive voicing assimilation in obstruents.

The contradiction is resolved in the context of the formal phonological constraints that are responsible. Japanese allows [pp], and forbids [bb], as part of a general ban on voiced obstruent geminates. Such a ban is phonetically sensible, because obstruent voicing is hard to maintain over long closures. Arabic allows [bb], and bans [pp], as part of a phonetically sensible ban on voiceless labial stops. The latter ban is phonetically sensible because of the large expanding oral chamber wall surface in labials. The opposite effects thus result from formally general phonological constraints, each with a phonetically natural core.

The tentative conclusion here is that the influence of phonetics in phonology is not direct, but is mediated by structural constraints that are under some pressure toward formal symmetry. A phonology that was directly driven by phonetic naturalness would, I think, be likely to miss this point.

The gap between the phonetic difficulty patterns and the phonology is thus still there, waiting to be bridged. Clearly, languages are well designed from a phonetic point of view. What is needed, I believe, is a way of accounting for this design that also allows principles of structural symmetry to play a role.

7. A Scheme for Phonological Grammar Design

Grammars could in principle be designed at two levels. Within the species as a whole, it is often held there is a Universal Grammar, invariant among non-pathological individuals, which determines much of the form of possible languages. Another sense in which grammar could be designed, outlined by Kiparsky and Menn (1977, 58), is at the level of the individual, who is engaged from infancy on in the process of constructing a grammar, one that will ultimately generate the ambient language or something close to it. Could the language learner be a Designer of Grammars? If so, how might she go about it?

From the discussion above, it would seem plausible that grammatical design within phonology aims at a compromise between formal symmetry and accurate reflection of phonetic difficulty. What follows is a tentative attempt to specify what phonological design could be like. It is very far from being confirmed, but I think it important at least to get started by laying out a concrete proposal.

The task of phonological grammar design, under Optimality Theory, has two parts: gaining access to constraints (here, by inventing them), and forming a grammar by ranking the constraints. The strategy taken is to suppose that constraints are invented in great profusion, but trimmed back by the constraint ranking algorithm.

The particular process whereby constraints are invented I will call *inductive grounding*. The term “grounded,” which describes constraints that have a phonetic basis, was introduced by Archangeli and Pulleyblank (1994). “Inductive” means that the constraints are learned by processing input data.

8. Inductive Grounding I: Evaluating Constraint Effectiveness

The language learner has, in principle, an excellent vantage point for learning phonetically grounded constraints. Unlike any experimenter observing her externally, the child is actually operating her own production and perception apparatus, and plausibly would have direct access to the degree of difficulty of articulations and to the perceptual confusability of different acoustic signals.

Beyond the capacity to judge phonetic difficulty from experience, a language learner would also require the ability to generalize across tokens, creating a *phonetic map* of the range of possible articulations and acoustic forms.

Considering for the moment only articulation, I will suppose that the language learner is able to assess the difficulty of particular phonological configurations, using measures such as the maximum articulatory force needed to execute the configuration, or perhaps simple energy expenditure.¹⁰ Further, we must suppose that the learner is able to generalize from experience, arriving at a measure of the *characteristic* difficulty of particular phonological configurations, which would abstract away from the variation found at various speaking rates and degrees of casualness, as well as the variable perceptual clarity that different degrees of articulatory precision will produce. Pursuing such a course, the learner could in principle arrive at a phonetic map of the space of articulatory difficulty.¹¹ A tentative example of a phonetic map is given below in section 10.

Given a phonetic map drawn from experience, a language learner could in principle use it to construct phonetically grounded constraints; hence the term “inductive grounding.” The inductive grounding algorithm I will suggest here supposes the following.

First, I assume constraints are constructed profusely, as arbitrary well-formed combinations of the primitive elements of phonological theory. In principle, this involves some risk, since the number of constraints to be considered grows exponentially with the number of formal elements included in their structural descriptions. However, if as suggested above, constraints are under some pressure toward formal simplicity, it is likely that the size of the search space can be kept under control.

¹⁰ As Katherine Demuth has pointed out to me, one should probably also consider motor-planning difficulty; for example, the difficulty very young children have in employing more than one place of articulation per word. Since such difficulty is at present impossible to estimate, I must stick to physical difficulty for now.

¹¹ Obviously, this task itself involves quite non-trivial learning. An encouraging reference from this viewpoint is Kelly and Martin (1994), who provide a fascinating survey of the ability of humans and other species to form statistical generalizations and to estimate relative magnitudes from experience.

Second, candidate constraints are assessed for their degree of grounding, accessing the phonetic map with a procedure I will now describe.

A grounded constraint is one that is phonetically sensible; that is, it bans things that are phonetically hard, and allows things that are phonetically easy. Taking a given candidate phonological constraint C, and any two entries E₁ and E₂ in the phonetic map, there are four logical possibilities:

- (6)a. Both E₁ and E₂ violate C.
- b. Both E₁ and E₂ obey C.
- c. E₁ violates C and E₂ obeys C.
- d. E₁ obeys C and E₂ violates C.

We will ignore all pairs of types (a) and (b) (same-outcome) as irrelevant to the assessment of C. Among the remaining possibilities, we can distinguish cases where the constraint makes an error from those in which it makes a correct prediction.

- (7)a. Correct predictions

E₁ violates C and E₂ obeys C; E₁ is harder than E₂.
E₁ obeys C and E₂ violates C; E₁ is easier than E₂.

- b. Errors

E₁ obeys C and E₂ violates C; E₁ is harder than E₂.
E₁ violates C and E₂ obeys C; E₁ is easier than E₂.

Since the goal of a constraint is to exclude hard things and include easy things, we can establish a simple metric of constraint effectiveness simply by examining all possible pairs {E₁, E₂} drawn from the phonetic map. The definition below presumes a particular phonological structural description defining a constraint, and a phonetic map against which the constraint may be tested:

- (8) Constraint effectiveness

$$\text{Effectiveness} = \text{Correct predictions} / (\text{Correct predictions} + \text{Errors})$$

On this scale, “perfect” constraints receive a value of 1, since they always ban things that are relatively harder, and never things that are relatively easier. Useless constraints, which ban things in an arbitrary way with no connection to their phonetic difficulty, receive a value of 0.5; and utterly perverse constraints, which ban only relatively easy things, get a value of 0. Clearly, the language learner should seek constraints with high effectiveness values.

It is more complicated to define constraint effectiveness for perceptual distinctness. Flemming (1995) has argued persuasively that perceptual distinctness can only be defined *syntagmatically* in perceptual space: for instance, [ɪ] is a fine vowel, indeed the preferred high vowel in a vertical vowel system such as Marshallese, where it is the only high vowel (Choi 1992). But where [i] and [u] occur as phonemes, as in most languages, [ɪ] is a poor vowel, due to its acoustic proximity to (thus, confusability with) [i] and [u]. Assuming the correctness of Flemming's position, we must evaluate not individual entries in the phonetic map, but *pairs* of entries. And since constraint effectiveness is determined by comparing cases that a constraint treats differently, we must deal with pairs of pairs. In various cases I have explored, this procedure leads to coherent results, but as there are further complications, I will consider only articulation here, with the intent of dealing with perception elsewhere.

9. Inductive Grounding II: Selecting the Grounded Constraints

Merely defining constraint effectiveness does not provide an explicit definition of a grounded constraint. If we only allowed constraints that showed a maximally good fit to the phonetic map (effectiveness value 1), then only a few simple constraints would be possible, and most of the permitted constraints would be very complex, like the “contour line constraint” in (3) above. This would be wrong on both counts. First, my judgment, based on experience in phonological typology, is that there are *many* constraints, in fact, dismayingly many, unless we come up with a reasonable source for them. Thus, we want the inductive grounding algorithm to generate a very rich (but thoroughly principled) constraint set. Second, as already argued, we want to keep constraints from being heavily “tailored” to fit the phonetic pattern. Real constraints seldom achieve such a perfect fit; rather, they deviate in the direction of structural simplicity.

A simple way to accomplish this deviation, as well as to provide a rich constraint set, is to rely on the notion of *local maximum*; in particular, local maxima of constraint effectiveness. Typically, local maxima are recognized as difficult problems for language learners, preventing the learner from arriving at the correct final state. A complex, multiply dimensioned pattern typically has many local maxima, but (definitionally) only one global one. But for our purposes, a local maximum is an excellent thing, because it permits a large number of constraints to emerge from a given phonetic map.

To make the idea explicit, here are some definitions:

- (9) *Constraint space* is the complete (infinite) set of possible constraints. It is generated by locating all legal combinations of the primitive formal elements of a particular phonological theory.
- (10) Two constraints are *neighbors* in constraint space if the structural description of one may be obtained from that of the other by a single primitive formal substitution (switching a feature value; addition or loss

of a feature or association line, etc.; the exact set of substitutions will depend on the phonological theory employed).

- (11) Constraint C_1 is said to be *less complex* than constraint C_2 iff the structural description of C_1 is properly included in the structural description of C_2 (cf. Koutsoudas et al. 1974, 8-9).

Using these definitions, we can now state an explicit characterization of phonetic grounding:

- (12) Defn.: *grounded*

Given a phonological constraint C and a phonetic map M , C is said to be *grounded* with respect to M if the phonetic effectiveness of C is greater than that of all neighbors of C of equal or lesser complexity.

Definition (12) uses the notion of local maximum, by requiring that C only exceed its neighbors in effectiveness. But (12) also goes beyond local maxima in a crucial sense: the neighbors that one must consider are only neighbors of equal or lesser complexity. It is this bias that permits the system to output relatively simple constraints even when their match to the phonetic map is imperfect.

The definition of phonetic grounding in (12) is obviously quite speculative, but I would claim the following virtues for it: (a) Assuming that a reasonably accurate phonetic map can be constructed, it specifies precisely which constraints are grounded with respect to that map, thus satisfying the requirement of explicitness. (b) The formally simple constraints that a given map yields are not just a few phonetically-perfect ones, but a large number, each a local effectiveness maximum within the domain of equally or less-complex constraints. (c) Constraints are able to sacrifice perfect phonetic accuracy for formal symmetry, since the competitors with which they are compared are only those of equal or lesser complexity.

10. An Application of Inductive Grounding

Here is a worked out example. To begin, we need a plausible phonetic map, for which I propose (13):

(13) A Phonetic Difficulty Map for Six Stops in Four Environments

	p	t	k	b	d	g
[-son] ____	7	0	0	43	50	52
# ____	10	0	0	23	27	35
[+son,-nas] ____	45	28	15	10	20	30
[+nas] ____	155	135	107	0	0	0

I obtained this map by using a software aerodynamic vocal tract model. This model was developed originally by Rothenberg (1968) as an electrical circuit model, and is currently implemented in a software version in the UCLA Phonetics Laboratory. This version (or its close ancestors) are described in Westbury (1983), Keating (1984), and Westbury and Keating (1986). Roughly, the model takes as input specific quantitative values for a large set of articulations, and outputs the consequences of these articulations for voicing, that is, the particular ranges of milliseconds during which the vocal folds are vibrating. The units in chart (13) represent articulatory deviations from a posited maximally-easy average vocal fold opening of 175 microns; these deviations are in the positive direction for voiceless segments (since glottal abduction inhibits voicing) and negative for voiced (since glottal adduction encourages it).

I used the model in an effort to give plausible quantitative support to the scheme to be followed here. However, it should be emphasized that obtaining reasonable estimates of articulatory difficulty from the model requires one to make a large number of relatively arbitrary assumptions, reviewed in the footnote below.¹² What makes the procedure defensible is that the outcomes that it produces are qualitatively reasonable: examining the map, the reader will find that all the relevant phonetic tendencies described above in

¹² (a) In real life, numerous articulations other than glottal adduction influence voicing (Westbury 1979, 1983); I have used glottal adduction alone, despite the lack of realism, to reduce phonetic difficulty to a single physical scale. To permit a uniform criterion of perceptual adequacy, the right-side environment for all stops was assumed to be prevocalic, which of course adds another caveat to the results.

(b) Inputs to the aerodynamic model were as in Keating (1984), modified for the postnasal environment as in Hayes and Stivers (1996).

(c) The criterion for adequate perceptual voicelessness was that the release of the stop should be voiceless and there should be at least a 50 msec. voiceless interval (half of the stop's 100 assumed msec. closure duration). The criterion for perceptual voicing was that the release of the stop should be voiced, and at least half of the stop closure should be voiced. Preceding obstruents and nasals were assumed to overlap with the target stop, so they added only 50 msec. to the total consonant closure.

(d) Since I had no basis for assessing what the true maximally-easy vocal fold opening is, I was forced (for this one parameter) to "let the theory decide"; picking the value of 175 as the one that best matched observed phonological typology.

section 6.2 are reflected quantitatively in the map. Thus, voiced stops are most difficult after an obstruent, somewhat easier in initial position, easier still after sonorants, and easiest postnasally. The reverse pattern holds for voiceless stops. Further, for any given environment, stops are easier to produce as voiced (and harder as voiceless) when they are in frontier places of articulation.

I will now derive a number of phonological constraints from the phonetic map of (13) by means of inductive grounding. The chart in (14) lists some of the work that must be done. The first column gives what I take to be a fairly substantial list of the most plausible constraints (given what the chart is suitable for testing), along with all of their simpler neighbors. I have imposed a relatively arbitrary limit of formal complexity on this candidate set, under the assumption that language learners either cannot or will not posit extremely complex constraints. The second column gives the phonetic effectiveness value for the candidate constraints, calculated by the method laid out in (9)-(12) and exemplified below.¹³ Finally, the third column lists all the neighbor constraints for each main entry that are equally or more simple, taking the assumption that these neighbors are obtained by either a feature value switch or by deletion of single elements from the structural description.

(14) Constraint	Effective-ness	Neighbors
a. *[+nasal][+voice]	0.000	*[+nasal][-voice], *[-nasal][+voice], *[+voice], *[+nasal]
b. *[+nasal][-voice]	1.000	*[+nasal][+voice], *[-nasal][-voice], *[-voice], *[+nasal]
c. *[-nasal][+voice]	0.701	*[-nasal][-voice], *[+nasal][+voice], *[+voice], *[-nasal]
d. *[-nasal][-voice]	0.357	*[-nasal][+voice], *[+nasal][-voice], *[-voice], *[-nasal]
e. *[+son][+voice]	0.500	*[+son][-voice], *[-son][+voice], *[+voice], *[+son]
f. *[+son][-voice]	0.861	*[+son][+voice], *[-son][-voice], *[-voice], *[+son]
g. *[-son][+voice]	0.841	*[-son][-voice], *[+son][+voice], *[+voice], *[-son]
h. *[-son][-voice]	0.094	*[-son][+voice], *[+son][-voice], *[-voice], *[-son]
i. *[LAB, +voice]	0.425	*[LAB, -voice], *[COR, +voice], *[DORS, +voice], *[LAB], *[+voice]
j. *[LAB, -voice]	0.633	*[LAB, +voice], *[COR, -voice], *[DORS, -voice], *[LAB], *[-voice]
k. *[COR, +voice]	0.500	*[COR, -voice], *[LAB, +voice], *[DORS, +voice], *[COR], *[+voice]
l. *[COR, -voice]	0.443	*[COR, +voice], *[LAB, -voice], *[DORS, -voice], *[COR], *[-voice]
m. *[DORS, +voice]	0.608	*[DORS, -voice], *[LAB, +voice], *[COR, +voice], *[DORS], *[+voice]
n. *[DORS, -voice]	0.371	*[DORS, +voice], *[LAB, -voice], *[COR, -voice], *[DORS], *[-voice]
o. *[+voice] unless LAB	0.568	*[-voice] unless LAB, *[+voice] unless COR, *[+voice] unless DORS, *[] unless LAB, *[+voice]
p. *[-voice] unless LAB	0.388	*[+voice] unless LAB, *[-voice] unless COR, *[-voice] unless DORS, *[] unless LAB, *[-voice]

¹³ Note that for some constraints, the effectiveness value cannot be calculated. When a constraint excludes or permits every entry in the map, then the formula for effectiveness in (8) will have a zero denominator. The only constraints for which this arose here were constraints included just because they were neighbors of other constraints.

q. *[+voice] unless COR	0.521	*[-voice] unless COR, *[+voice] unless LAB, *[+voice] unless DORS, *[] unless COR, *[+voice]
r. *[-voice] unless COR	0.513	*[+voice] unless COR, *[-voice] unless LAB, *[-voice] unless DORS, *[] unless COR, *[-voice]
s. *[+voice] unless DORS	0.453	*[-voice] unless DORS, *[+voice] unless LAB, *[+voice] unless COR, *[] unless DORS, *[+voice]
t. *[-voice] unless DORS	0.556	*[+voice] unless DORS, *[-voice] unless LAB, *[-voice] unless COR, *[] unless DORS, *[-voice]
u. *[LAB]	0.541	*[COR], *[DORS]
v. *[COR]	0.466	*[LAB], *[DORS]
w. *[DORS]	0.491	*[LAB], *[COR]
x. *[] unless LAB	0.459	*[] unless COR, *[] unless DORS
y. *[] unless COR	0.534	*[] unless LAB, *[] unless DORS
z. *[] unless DORS	0.509	*[] unless LAB, *[] unless COR
aa. *[+voice]	0.519	*[-voice]
bb. *[-voice]	0.481	*[+voice]
cc. *[+nasal]	(undetermined)	*[-nasal]
dd. *[-nasal]	(undetermined)	*[+nasal]
ee. *[+son]	(undetermined)	*[-son]
ff. *[-son]	(undetermined)	*[+son]

Here is an example of how effectiveness was computed for individual constraints. The constraint *[LAB, -voice] bans [p]; this ban is phonetically natural (for reasons already given) and would thus be expected to have a reasonably high effectiveness value. I repeat the phonetic map below, this time with letters a-x, permitting reference to the entries:

(15)

	p	t	k	b	d	g
[-son] ____	a: 7	b: 0	c: 0	d: 43	e: 50	f: 52
# ____	g: 10	h: 0	i: 0	j: 23	k: 27	l: 35
[+son, -nas] ____	m: 45	n: 28	o: 15	p: 10	q: 20	r: 30
[+nas] ____	s: 155	t: 135	u: 107	v: 0	w: 0	x: 0

*[LAB, -voice] bans the shaded region of the map. If it is to be effective, then pairwise comparisons between banned cells and unbanned ones should predominantly come out with the banned cells being more difficult. Here is the outcome; “>” means “is harder than”:

(16) a. Correct Predictions: 50

- a** > b, c, h, i, v-x
- g** > b, c, h, i, v-x
- m** > b, c, d, h-l, n-r, v-x
- s** > b-f, h-l, n-r, t-x

b. Incorrect Predictions: 29

- a** < d-f, j-l, n-r, t, u
- g** < d-f, j-l, n, o, q, r, t, u
- m** < e, f, t, u
- s** < (none)

The computed effectiveness value is $50/(50 + 29)$, or .633, which is what was listed in (14j).¹⁴

The neighbors of ***[LAB, –voice]** that have equal or lesser complexity are listed below with their effectiveness values:

(17)

Constraint	Effectiveness	Justification for neighbor status
*[LAB, +voice]	0.425	switch value of [voice]
*[COR, –voice]	0.443	switch value of PLACE
*[DORS, –voice]	0.371	switch value of PLACE
*[LAB]	0.541	delete [+voice]
*[–voice]	0.481	delete [LAB]

Since ***[LAB, –voice]** at .633 exceeds all of its neighbors in effectiveness, the definition (12) designates it as phonetically grounded with respect to the phonetic map (13).

Repeating this procedure, we find that the constraints listed in (18) emerge as phonetically grounded. In the chart below, I give some mnemonic labels, often embodying a particular effect that a constraint might have. However, the reader should bear in mind that in Optimality Theory the empirical effects of a constraint can range much more widely than the label indicates; see for example Pater (1995, 1996).

(18)

Constraint	Effectiveness	Characteristic Effect
a. *[+nasal][–voice]	1.000	postnasal voicing
b. *[+son][–voice]	0.861	postsonorant voicing
c. *[–son][+voice]	0.841	postobstruent devoicing
d. *[–nasal][+voice]	0.701	postoral devoicing
e. *[LAB, –voice]	0.633	*p
f. *[DORS, +voice]	0.608	*g
g. *[+voice] unless LAB	0.568	/b/ is the only voiced stop
h. *[–voice] unless DORS	0.556	/k/ is the only voiceless stop
i. *[LAB]	0.541	*labials
j. *[] unless COR	0.534	COR is the only place
k. *[+voice]	0.519	voicing prohibited

The other constraints are designated by the algorithm as *not* grounded, because they are not local effectiveness maxima:

¹⁴ A simple BASIC program to carry out this tedious procedure may be downloaded from the author's web page: <http://www.humnet.ucla.edu/humnet/linguistics/people/hayes/hayes.htm>.

(19)	Constraint	Effectiveness	Characteristic Effect
a.	*[+voice] unless COR	0.521	/d/ is the only voiced stop
b.	*[-voice] unless COR	0.513	/t/ is the only voiceless stop
c.	*[] unless DORS	0.509	DORS is the only place
d.	*[COR, +voice]	0.500	*d
e.	*[+son][+voice]	0.500	postsonorant devoicing
f.	*[DORS]	0.491	*dorsals
g.	*[-voice]	0.481	voicing obligatory
h.	*[COR]	0.466	*coronals
i.	*[] unless LAB	0.459	LAB is the only place
j.	*[+voice] unless DORS	0.453	/g/ is the only voiced stop
k.	*[COR, -voice]	0.443	*t
l.	*[LAB, +voice]	0.425	*b
m.	*[-voice] unless LAB	0.388	/p/ is the only voiceless stop
n.	*[DORS, -voice]	0.371	*k
o.	*[-nasal][-voice]	0.357	postoral voicing
p.	*[-son][-voice]	0.094	postobstruent voicing
q.	*[+nasal][+voice]	0.000	postnasal devoicing

The neighbor constraint that “defeats” each of (19) may be determined by consulting chart (14).

Lastly, there are four constraints (*[+nasal], *[-nasal], *[+son], and *[-son]) for which the algorithm makes no decision, since the map of (13) does not bear on their status. These constraints were included simply to provide neighbors for the truly relevant constraints. I assume they could be evaluated by a more comprehensive map.

Did the simulation work? If the map in (13) is valid, and if languages adopt only grounded constraints, then the constraints of (18) should be empirically attested, and those of (19) not.

(a) The “finest” grounded constraint, with effectiveness value 1, is (18a), **[+nasal][-voice]**. This constraint is indeed widely attested, with noticeable empirical effects in perhaps 7.6% of the world’s languages (estimate from Hayes and Stivers 1996). Voicing in sonorant-adjacent positions ((18b), ***[+son][-voice]**) and devoicing in obstruent clusters ((18c), ***[-son][+voice]**) is also quite common.

(b) The chart also includes all the characteristic place-related voicing patterns: the bans on fronter voiceless stops and on backer voiceless ones (18e-h).

(c) Two of the simpler constraints, (18i) ***[LAB]** and (18j) ***[] unless COR**, do play a role in phonologies (see Rood 1975 and Smolensky 1993), but their appearance in the chart is probably accidental. The phonetic map used here is suitable only for testing

constraints on obstruent voicing, not place inventories. A legitimate test of the constraints that target place would require a much larger phonetic map.

(d) Likewise, the blanket ban on voicing ((18k)*[+voice]) makes sense only if one remembers that the map (18) only compares obstruents. Since voicing in sonorants is very easy, it is likely that in a fuller simulation, in which the map included sonorants, the constraint that would actually emerge is *[-sonorant, +voice]. This is well attested: for example, 45 of the 317 languages in Maddieson's (1984) survey lack voiced obstruents.

(e) The only non-artifactual constraint designated as grounded that probably is not legitimate is (18d), [-nasal][+voice], which would impose devoicing after oral segments. It has been suggested by Steriade (1995a) and others that [nasal] is a privative feature, being employed in phonological representations only to designate overt nasality. If this is so, then [-nasal][+voice] would not even appear in the candidate set.

(f) We can also consider the constraints of (19), which emerge from the simulation designated as not grounded. My impression, based on my own typological experience, is that these constraints are indeed rare or unattested in actual languages. Obviously, careful typological work would be needed to affirm this conclusion.

I would conclude that the inductive grounding procedure, applied in this narrow domain, does indeed single out the phonologically-stated constraints that match typology. It is interesting that some of the constraints (for example (18e) *[LAB, -voice]) do not record extremely high effectiveness scores, but are nevertheless fairly well attested (19 languages of the 317 in Maddieson 1984 shows a stop gap at [p]). This suggests, as before, that formal symmetry, and not just phonetic naturalness, plays a role in constraint creation.

11. The Remainder of the Task of Phonological Acquisition

Above I have outlined a procedure that, equipped with full-scale phonetic maps, could generate large numbers of grounded constraints. What are we to do with them, in order to obtain actual grammars?

In Optimality Theory, the answer is simply: rank them. Tesar and Smolensky (1993, 1995) have demonstrated an algorithm that ranks constraints using input data, with high computational efficiency. I suggest that the promiscuously-generated constraints from inductive grounding could simply be fed into the Tesar-Smolensky algorithm. The algorithm will rank a few of them high in the grammar, the great majority very low. In Optimality Theory, a constraint that is ranked low enough will typically have no empirical effects at all. Thus, the Tesar-Smolensky algorithm can in principle weed out the constraints that, while grounded, are inappropriate for the language being learned.

The combined effect of inductive grounding and the Tesar-Smolensky algorithm is in principle the construction of a large chunk of the phonology. The further ingredients

needed would be constraints that have non-phonetic origins. These include: (a) the Faithfulness constraints; these perhaps result from their own inductive procedure, applied to the input vocabulary; (b) functionally-based constraints that are not of phonetic origin: for example, rhythmically-based constraints (Hayes 1995), or constraints on paradigm uniformity. Moreover, the child must also learn the phonological representations of the lexicon, a task that becomes quite non-trivial when these diverge from surface forms. Even so, I believe that getting the phonetic constraints right would be a large step towards phonology.

12. Acquisition Evidence

The above discussion was entirely formal in character, attempting to develop an abstract scheme that was at least explicit enough to be confronted with actual data. But what of real children? Is there any evidence that they can generate formally symmetrical constraints from their own phonetic experience?

In considering this question, I will refer to a very substantial research tradition in phonological acquisition. To summarize the results quickly and in inadequate detail, it appears that the following hold:

(a) Children's perceptions are well ahead of their productions (Smith 1973; Braine 1974, 284; Ingram 1989, 162-8; Eimas 1996, 32). Although in certain cases (Macken 1980a, 1995) a child's errors can be shown to be the result of misperception, there is strong evidence that children can internalize many adult-like lexical forms that are neutralized only in their own productions.

(b) Children naturally develop procedures to reduce the complexity of adult forms to something they can handle with their limited articulatory abilities. These procedures frequently develop sufficient regularity that it is reasonable to refer to them as the child's own phonology; that is, a phonology that serves to map adult surface forms (or perhaps something deeper) into child surface forms.

(c) The phonology of children is elaborate beyond what is required to reduce the child's speech to something easily pronounceable. For example, Amahl, the subject of Smith (1973), developed a remarkable form of "labiality flopping," whereby the labiality of the /w/ in (for example) /kwi:n/ 'queen' migrated rightward, surfacing on the final consonant and converting it to /m/: [gi:m]. Another extraordinary migration (*string*: [ˈtriŋs]) is documented by Hamp (1974).

(d) Lastly, children's phonologies are to a fair degree specific to the individual child (indeed, to a particular child at a particular phase of acquisition). There is no such thing as "English infantile phonology"; only the phonologies created by particular children.

These results, which I take to be relatively uncontroversial, lead to the conclusion (Kiparsky and Menn 1977; Macken 1995) that to some degree, phonology is not merely *learned* by children, but is to some extent also *created* by them.

Let us consider some of the constraints that have been created by children.

(a) Amahl Smith, at age 2 years, 60 days, rendered all stops (irrespective of underlying form) as voiceless unaspirated lenis initially, voiced in medial position, and voiceless finally; thus ['t̥ebu] 'table', [a:t] 'hard', ['wə:ɡin] 'working'. Plainly, such realizations cannot have been an imitation of adult speech; they were Amahl's own invention. Equally plainly, the constraints Amahl adopted have a real role in the phonology of languages other than English; consider for instance the voicing of intervocalic stops in Korean, or the devoicing of final stops in German. Finally, as noted above, Amahl's constraints render articulation easier, by imposing the default values predicted on aerodynamic grounds.

(b) Amahl also required every consonant to be either prevocalic or final, so he produced no consonant clusters. The phonetic naturalness of such a pattern has been argued for by Steriade (1995b); and it has been observed in adult language in the phonology of Gokana (Hyman 1982, 1985).

(c) Children who impose gaps in their stop inventories at [p] or at [g], contrary to adult input, are described by Ferguson (1975) and Macken (1980b). These gaps are analogous to the gaps of adult languages noted in section 6.2. They are phonetically natural, and indeed are predicted by the phonetically grounded constraints (18e,f) derived in the simulation above.

(d) Both Ferguson (1975, 11) and Locke (1983, 120) report cases of children who (against input evidence) require all postnasal obstruents to be voiced. Again, this is phonetically natural, derived under my simulation (18a), and typologically commonplace.

In all of these cases, the point to observe is that children have the capacity to obtain constraints that are phonetically grounded, formally simple, and not available from the ambient language data. I conclude that a good case can be made that children really do have a means of fabricating constraints that reflect phonetic naturalness, perhaps by

something like the method of inductive grounding laid out above.¹⁵

To conclude this section, we must complete the explanatory chain by establishing appropriate links between child-innovated constraints and adult phonology. There are two possibilities.

First, there is the link of *learnability*: it is possible that the child's search for the adult grammar is aided by the child's hypothesis that the adult grammar will contain grounded constraints. Thus, in principle, the ability to access the set of grounded constraints could speed acquisition, though I think it would be hard at present to obtain serious evidence on this point.

Second, there is the *diachronic* link. Suppose that certain constraints fabricated by individual children manage to survive into the adult speech community, perhaps by being adopted in a peer group of young children.¹⁶ This would account for the characteristic naturalness and formal symmetry of adult constraints, without positing that naturalness is a criterion for learnability.

¹⁵ This is not to say that all children's constraints are the same as adults'. For example, the slower tempo of child speech (Smith 1978) means that children escape the phonetic difficulties of "antigemination," which have been explained phonetically by Locke (1983, 174) and Odden (1988, 470). For this reason, children can indulge in widespread consonant harmony, which the antigemination effect rules out for adults. The lesser articulatory skill of children is probably the cause of frequent stop-for-fricative substitutions; adults, who are more skillful but in a bigger hurry, tend instead toward lenition, with intervocalic spirantization. I assume that as children come to speak faster and with greater articulatory control, their phonetic maps change, with an accompanying shift towards adult-like constraints. For further discussion of this issue, see Locke (1983) and Macken (1995).

¹⁶ The reader, who almost certainly speaks a normatively-imposed standard language, might find this counterintuitive, unless (s)he remembers that most languages are colloquial, non-standard varieties. As Hock (1986, 466-7) remarks, nonstandard languages change quite a bit more rapidly than standard ones. I would conjecture that this is because they suppress the innovations of children with considerably less force. The abundance of non-standard English dialects that replace [θ, ð] with [f, v] or [t, d] (both normal childhood substitutions) is a good illustration. Given that such dialects are geographically remote from each other, it seems very likely that these substitutions are childhood inheritances (Wells 1982, 96-7).

Either of these hypotheses would account for the characteristic appearance of grounded constraints in adult grammars.¹⁷

13. Innate Knowledge

Lurking in the background of much of this discussion is a belief widely held by formal linguists: that much (most?) of linguistic structure is specified innately, and does not have to be learned by any procedure at all. For Optimality Theory, it is suggested by (for example) Tesar and Smolensky (1993, 1) that all the constraints might be innate, so that the creation of grammar in the child would largely reduce to the task of ranking these already-known constraints. To the contrary, I have been assuming that constraints need not necessarily be innate, but only *accessible in some way* to the language learner, perhaps by inductive grounding.

On the whole, it is very hard to make this issue an empirical one. I know of two sources of facts that might bear on the question.

First, there are phonetically grounded constraints that govern uncommon sounds. Among these are the constraints discovered by Steriade requiring postvocalic position for retroflexes and for preglottalized sonorants. In Maddieson's (1984) survey, only 66 of the 317 languages sampled had retroflexes, and only 20 had laryngealized sonorants of any sort. Similarly, implosives and ejectives display place asymmetries much like the place asymmetries for voiced and voiceless stops, respectively (though more robust), and have similar aerodynamic explanations (see Maddieson, Chap. 7, and references cited there). Implosives occur in only 32 languages of the Maddieson sample, ejectives in 52.

If the proto-stages of human language likewise seldom deployed retroflexes, preglottalized sonorants, ejectives, and implosives, then during most of the period for the evolution of language, there can have been little selectional pressure to deploy these sounds in any particular way. There is no selective advantage to possessing an innate constraint on the distribution of retroflexes if the language you are learning doesn't have any. From the viewpoint of inductive grounding, in contrast, such constraints are unproblematic: children can obtain the phonetic maps necessary for acquiring them from

¹⁷ The indebtedness of this whole section to the work of Stampe (1972) and Donegan and Stampe (1979) should be clear. The approach I have taken could be viewed as an attempt to extend Stampe and Donegan's work, making use of Optimality Theory to establish a more direct connection between phonetics and child phonology. In Optimality Theory, one need merely specify in a constraint what is phonetically hard, with the Faithfulness constraints determining what particular "fix" is adopted to avoid phonotactic violations. In contrast, Natural Phonology requires a massive proliferation of processes, each needed to characterize one particular strategy for avoiding phonetic difficulty.

the practice they obtain in imitating an ambient language that happens to have the relevant sounds.¹⁸

A very different source of evidence on the innateness question comes from Locke and Pearson (1990). These authors studied a child who was deprived of articulatory practice for part of her infancy because of a temporarily-installed tracheal tube. What they found suggests that learning through phonetic self-exploration may indeed be important to acquisition, as the child they studied was delayed considerably in phonological acquisition once the tube was removed. Locke and Pearson are cautious in interpreting this result, but in principle such research could provide serious empirical data on the question of the innateness of phonetic constraints.

14. Ungrounded Constraints

It has often been emphasized that a language's phonological structure is not always sensible. A language may have a system that is synchronically odd, as a result of a conspiracy of historical circumstances such as borrowing, or a peculiar sequence of changes, each one natural (Bach and Harms 1972; Hayes 1995, 219-21).

One possible example comes from Northern Italian, which shows the rather odd pattern of voicing /s/ intervocalically but not postnasally. The pattern is productive, as Baroni's (1996) recent testing indicates. The sequence of events that gave rise to this pattern historically was (a) loss of all nasals before /s/ (early Romance); (b) intervocalic /s/ voicing; (c) reintroduction of postnasal /s/ in learned borrowings from Latin, pronounced faithfully to the Latin original (Maiden 1995; 14, 63, 76, 84). While it is not clear whether purely-intervocalic voicing is grounded (in my simulation, it depends on the feature system used), nevertheless the Northern Italian phenomenon does seem somewhat peculiar in light of the pattern of phonetic difficulty involved.

A perusal of Maddieson (1984) will show a number of stop systems that have gaps at places other than the expected *[p] and *[g]. Although Ohala (1983) suggests additional factors that may influence voicing-gap patterns, it appears likely that many of these systems are also accidents of history, and must be attributed to ungrounded constraints.

Two points seem worth making about ungrounded constraints. First, if grammars really do permit them, then they must have some source. I would conjecture that the source is induction, in this case not over the learner's phonetic experience but over the input data: eventually, the child figures out such constraints from (systematic, consistent, long-term) negative evidence. Such constraints would be the rough analogues in the present theory of Stampe's (1973) 'rules', as opposed to the grounded constraints, which correspond roughly to Stampe's 'processes'.

¹⁸ This paragraph presupposes that any innate principles of language did arise by natural selection. For defense of this view, and criticism of the alternative possibilities, see Pinker and Bloom (1990) and Dennett (1995).

Second, if the distinction between inductively grounded constraints (learned from internal experience) and learned constraints (learned from gaps in input data) is true, then it should be detectable. Here are some possible ways to detect it:

(a) Children who innovate constraints in their own speech should never innovate an ungrounded constraint.

(b) In principle, grounding could influence intuitive judgments. For instance, Donca Steriade has pointed out in lectures that in English, hypothetical forms like [rtap], with a gross sonority violation, sound much worse than forms like [ktap], with a lesser violation. This is despite the fact that neither form occurs in the English input data. I would conjecture that the difference in judgment has its origins in the phonetic naturalness of the two configurations. By way of contrast, we might expect purely learned, ungrounded constraints to provide judgments related to the lexicon; that is, to the degree to which the child's input justifies the inductive conclusion that a particular segment or sequence is absent.

(c) Borrowed words might also provide evidence: new borrowed phonemes and sequences should be more easily pronounced if they merely violate arbitrary learned constraints than if they violate phonetically grounded ones.¹⁹

What emerges here is that, while the existence of ungrounded constraints makes it harder to test a theory of phonetic grounding, it does not make it impossible.

15. Consequences of Inductive Grounding for Feature Theory

A major line of evolution in feature theories (traceable, for example, through Jakobson, Fant and Halle 1951, Chomsky and Halle 1968, and Sagey 1986) has been one of increasing *phonetic literalism*: the features have gradually come closer to depicting what is going on in the mouth during speech. Autosegmental representations, which permit an idealized depiction of the timing of individual articulators, increase the degree of literalism.

In one sense, this has been a positive development: since phonology is mostly phonetically grounded, formal representations that include a more precise depiction of the phonetics will do better in many cases than those that do not. However, I believe that detailed consideration of various cases indicates that the “phonetic literalist” research program for feature theory has not really achieved its goals. Inductive grounding suggests what may be a better direction for feature theory to follow.

¹⁹ A frustrating interfering factor here is that the adult speakers have had massive practice, for years, in pronouncing precisely the sounds of their language. Presumably, this has substantial effects on their phonetic maps.

The problem is that phonetics is very complicated, and involves physical and perceptual systems that interact in many ways. Ordinary phonological representations, even those designed with an eye on phonetic form, are simply not rich enough to characterize all the things that can happen (Flemming 1995).

Perhaps the plainest example of this is the mechanism of postnasal voicing, investigated by Hayes and Stivers (1996). Hayes and Stivers suggest that the widespread preference for postnasal voicing follows from a quite elaborate set of contraposed phonetic tendencies. First, obstruents tend to voice in nasal-*adjacent* position because nasals induce on them a slight, coarticulatory nasal “leak,” at a level that does not render obstruents perceptually nasal, but does encourage voicing by venting oral pressure. Second, a peculiar tendency of the velum to *rise while the velar port is closed* during a nasal-to-obstruent transition (and correspondingly, fall while closed in obstruent-to-nasal transitions) produces a kind of “velar pumping,” which yields an anti-voicing effect in obstruent + nasal sequences (thus negating the voicing effect of nasal leak) but a pro-voicing effect in nasal + obstruent sequences (reinforcing the nasal leak effect). Putting these effects together (and modeling them quantitatively), Hayes and Stivers predict specifically post-nasal voicing, which is what agrees with typology.

In principle, the highly detailed and specific phonetic effects studied by Hayes and Stivers could be encoded in the phonology: spreading principles would depict the coarticulation, and special new features would depict the resulting aerodynamic effects. With such features, the constraint against postnasal voiceless obstruents would come out as something like (20):

$$(20) * \left[\begin{array}{l} \text{–sonorant} \\ \text{–voice} \\ \text{+minor nasal leak} \\ \text{+rarefactive velar pumping} \end{array} \right]$$

But the last two features in (20) are hopeless as members of a *phonological* feature inventory, as they play no role at all elsewhere in phonology: they define no natural classes, do not spread, and are completely redundant.

Inductive grounding covers the ban on postnasal voicelessness by addressing a phonetic map, as shown above. The features it uses in formulating and scanning the map ([voice], [nasal], and [sonorant]) are almost totally uncontroversial, being pervasively relevant to many aspects of phonology. Moreover, inductive grounding accounts for why nasal-adjacent voicing of obstruents is always in *post*-nasal position, never prenasal. As Pater (1996) notes, this is a conspicuous gap in the recent analysis by Ito, Mester, and Padgett (1995), which treats the phenomenon as voicing spread.

Consider another area in which phonetically literalist feature theory fails: phonological assimilations that have more than one trigger. For example, Hyman (1973) notes various tonal rules in which a H(igh) tone becomes L(ow) or rising (LH) after the

sequence *L-toned vowel + voiced obstruent*. In such a process, both the L tone and the voiced obstruent must be considered as factors contributing to the change, as each one often triggers the lowering effect by itself in other languages. To my knowledge, there is no featural account that covers “two-trigger” phenomena, because the autosegmental theory of assimilation only allows a single trigger to spread its feature value onto the target.

Inductive grounding appears to be a more promising approach here, because phonetic effects can be additive. A H tone is harder to produce after a voiced obstruent (for the phonetics, see Ohala 1978 and Hombert 1978); it is also harder when the preceding vowel is L; and it is hardest of all when both environments are present. Thus inductive grounding would plausibly single out the crucial constraint in (21) as grounded:

$$(21) \begin{array}{ccc} * & V & \left[\begin{array}{c} C \\ -\text{son} \\ +\text{voice} \end{array} \right] & V \\ & | & & | \\ & L & & H \end{array}$$

Two-trigger processes in phonology are quite common: typical examples are intersonorant voicing and intervocalic lenition (Kirchner, in progress).

The upshot of this discussion is this: it would be a mistake for phonologists to continue to formulate feature theory by attempting to construct simple schematic representations capable of mirroring the extremely complex behaviors of phonetics. This is not really a feasible task, and inductive grounding provides a more realistic alternative.

What *is* the right direction for feature theory, then? A better approach, I think, is to construe the feature system as describing how the language learner/user *categorizes* phonological form. The phonetic experience that must be entered into phonetic maps is extremely variegated; so for a map to be at all coherent or useful, the experience must be sorted into salient phonological categories. I believe that features form the basis of these categories. The categories of feature theory are also what serve as the raw material for the constraints that are tested against the phonetic maps. In principle, a feature inventory that is not especially literalist would, because it is small, reduce the hypothesis space that must be considered in the fabrication of constraints by inductive grounding, and thus render the search for effective constraints more feasible.

As for what research strategy would confirm particular constraints: the crucial diagnostic would be based on a property of constraints covered above in section 6.2, namely their tendency to be formally symmetrical at the expense of close fit to the phonetics. It is precisely when constraints deviate in minor ways from perfect grounding that we can infer that formal symmetry is playing a role. The features can be justified by whether they capture the necessary formal symmetry. Thus, for example, even though the phonetic mechanisms needed to produce a voiced intervocalic stop in Korean are not

exactly the same for all the Korean places of articulation, the fact that all of the places participate in parallel in an intervocalic voicing process suggests that [voice] is an authentic phonological feature of Korean. I would expect that most of the relatively uncontroversial current features, such as [coronal], [round], [nasal], and [back], could be justified in this way.

16. Local Conclusions

To sum up the main content of this paper: I have suggested, following much earlier work, that phonological constraints are often phonetic in character. They are not phonetics itself, but could in principle be “read off” the phonetics. Most of what I have said has been an effort to specify what this “reading off” could consist of.

The hypotheses considered have been, in increasing order of specificity: (a) Learners extract phonological constraints from their own experience; (b) In constructing constraints, learners execute a trade-off between phonetic accuracy and formal simplicity; (c) Learners go through the logical space of possible phonological constraints, seeking local maxima of good phonetic fit, and at each point comparing candidate constraints only with rivals of equal or greater simplicity.

I have further suggested that the data of child language support the view that children can and do create constraints by inductive grounding, and made suggestions regarding how feature theory might work under an inductive grounding approach.

17. General Conclusions

In principle, the approach taken here to functional factors in language is applicable elsewhere in linguistics. The basic idea has been that functional factors are represented indirectly: they enter in at the level of language design, leading to the construction of formal grammars that are functionally good, with a bias toward formal symmetry. I have posited that the functional factors make themselves apparent in “maps,” compiled from the experience of the language learner. Inductive grounding creates constraints that reflect the functional principles, in a way that is somewhat indirect, due to their formal character. Finally, constraint ranking molds the raw set of constraints into a full and explicit grammar. If the approach of Optimality Theory is correct, such grammars will do full justice to the amazing intricacy of linguistic phenomena.

If this view of things is right, there are a number of things we should expect to find in the linguistic data.

First, grammar should emerge rather consistently as functionally good. In the area of phonology, I am encouraged in this respect by my reading of the literature cited in section 4: by consistently examining their data with the question “why” in mind, the authors of this work have been able to expand considerably the domain of phonological facts that have plausible phonetic explanations.

Second, we should find that functional goodness appears in grammar not directly, but mediated by grammatical constraints, with a strong bias toward formal symmetry.

Third, we should find a pervasive role for *violable* grammatical constraints, as Optimality Theory claims, since constraints based on functional principles have no *a priori* claim to inviolability.

As noted earlier, very little in what is assumed here need be posited as innate knowledge. In principle, only the procedure for inductive grounding and the mechanisms of Optimality Theory itself need be innate, the rest being learned. But I am not at all *a priori* opposed to positing that parts of grammar and phonology are genetically encoded. This view seems especially cogent in domains of grammar that abound in “projection puzzles” (Baker 1979).

However, I do have a suggestion regarding research strategy: arguments for innate principles can only be made *stronger* when inductive alternatives are addressed and refuted. By this I mean both induction from internally-generated maps, as discussed here, and also ordinary induction from the input data. When induction has been explicitly shown to be inadequate, innateness is left in a much stronger position as the only non-mystical alternative.

18. Coda: “Good Reductionism” in Linguistics

Dennett (1995) has recently written a book that combines an excellent tutorial in evolutionary theory with interesting discussion of the relationship of evolution to cognitive science. Dennett suggests that an appropriate stance for a cognitive scientist to take is a form of “Good Reductionism,” which may be characterized as follows:

(a) Good Reductionism acknowledges the incredible richness and complexity of cognitive phenomena, and thus is the opposite of the trivializing “Greedy Reductionism.”

(b) Good Reductionism takes engineering, not physics, as its physical-science model. The reason is that natural selection tends to produce incrementally-engineered solutions, rather than proceeding with bold, fundamental moves.

(c) But on rare occasions, natural selection produces a “crane,” a particular trick that can make the apparently-miraculous phenomena of biology emerge from mundane origins. Examples of cranes include the “Baldwin Effect,” described by Dennett (1991, 184-7; 1995, 77-80); or sexual reproduction (Dennett 1995, 323).

(d) Cranes are opposed by Dennett to “skyhooks,” which explain the apparently miraculous by positing actual miracles. Skyhooks are obviously scientifically inappropriate, but have been proposed by scientists surprisingly often, he claims.

The approach taken here might be construed as an attempt to engage in Good Reductionist phonology, steering between the twin perils of reckless Skyhook Seeking and head-in-the-sand Greedy Reductionism. The two cranes I have posited are Optimality Theory and inductive grounding. These, and only these, must be assumed to be innate. Elsewhere, the approach has been incrementalist: the goal is to reconstruct the miraculous complexities of phonological systems incrementally, using materials that are directly accessible to the language learner.

The approach proposed is formalist in that it seeks to attain utterly explicit and complete phonological description. It is functionalist in that it seeks to obtain much of the content of phonology from external, functional principles, by means of inductive grounding. What emerges, I hope, is somewhat different from what has dominated either traditional formalist or traditional functionalist thinking.

References

- Anderson, Stephen R. (1981) "Why phonology isn't 'natural,'" *Linguistic Inquiry* 12, 493-539.
- Archangeli, Diana and Douglas Pulleyblank (1994) *Grounded Phonology*, MIT Press, Cambridge, MA.
- Bach, Emmon and Robert Harms (1972) "How do languages get crazy rules?" in Robert P. Stockwell and Ronald K. S. Macaulay, eds., *Linguistic Change and Generative Theory*, Indiana University Press, Bloomington, pp. 1-21.
- Baker, C. L. (1979) "Syntactic theory and the Projection Problem," *Linguistic Inquiry* 10, 533-581.
- Baroni, Marco (1996) "Morphological conditions on the distribution of [s] and [z] in Standard Northern Italian," ms., Department of Linguistics, UCLA, Los Angeles, CA.
- Braine, Martin S. (1974) "On what might constitute learnable phonology," *Language* 50, 270-299.
- Chen, Matthew (1973) "On the formal expression of natural rules in phonology," *Journal of Linguistics* 9, 223-249.
- Choi, John-Dongwook (1992) *Phonetic underspecification and target-interpolation: an acoustic study of Marshallese vowel allophony*, Ph.D. dissertation, UCLA, Los Angeles, CA.
- Chomsky, Noam and Morris Halle (1968) *The Sound Pattern of English*, Harper and Row, New York.
- Dennett, Daniel C. (1995) *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Simon and Schuster, New York.
- Dennett, Daniel C. (1991) *Consciousness Explained*, Little, Brown, Boston.
- Devine, A.M. and Laurence D. Stephens (1977) *Two Studies in Latin Phonology*, Anma Libri, Saratoga, CA.
- Donegan, Patricia Jane and David Stampe (1979) "The study of Natural Phonology," in Daniel A. Dinnsen, ed., *Current Approaches to Phonological Theory*, Indiana University Press, Bloomington, pp. 126-73.
- Eimas, Peter (1996) "The perception and representation of speech by infants," in James L. Morgan and Katherine Demuth, eds., *Signal to Syntax: Bootstrapping from Speech to Grammar in Early Acquisition*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 25-39.
- Ferguson, Charles (1975) "Sound patterns in language acquisition," in Daniel P. Dato, ed., *Developmental Psycholinguistics: Theory and Applications, Georgetown University Round Table on Languages and Linguistics*, 1975, Georgetown University Press, Washington, DC, pp. 1-16.
- Flemming, Edward (1995) *Perceptual Features in Phonology*, Ph.D. dissertation, UCLA, Los Angeles, CA..
- Gordon, Matt (forthcoming) "A phonetically-driven account of syllable weight," submitted; to appear in *UCLA Working Papers in Phonology* 3.
- Hamp, Eric P. (1974) "Wortphonologie," *Journal of Child Language* 1, 287-288.
- Harrell, Richard, Laila Y. Tewfik, and George D. Selim (1963) *Lessons in Colloquial Egyptian Arabic*, Georgetown University Press, Washington, DC.

- Hayes, Bruce (1989) "Compensatory lengthening in moraic phonology," *Linguistic Inquiry* 20, 253-306.
- Hayes, Bruce (1995) *Metrical Stress Theory: Principles and Case Studies*, University of Chicago Press, Chicago.
- Hayes, Bruce and Aditi Lahiri (1991) "Bengali intonational phonology," *Natural Language and Linguistic Theory* 9, 47-96.
- Hayes, Bruce and Tanya Stivers (1996) "The phonetics of postnasal voicing," ms., Dept. of Linguistics, UCLA, Los Angeles, CA.
- Hock, Hans Henrich (1986) *Principles of Historical Linguistics*, Mouton de Gruyter, Berlin.
- Hombert, Jean-Marie (1978) "Consonant types, vowel quality, and tone," in Victoria Fromkin, ed., *Tone: A Linguistic Survey*, Academic Press, New York, pp. 771-111.
- Hyman, Larry M. (1973) "The role of consonant types in natural tonal assimilations," in Larry M. Hyman, ed., *Consonant Types and Tone*, Southern California Occasional Papers in Linguistics No. 1, pp. 151-176.
- Hyman, Larry M. (1976) "Phonologization," in Alphonse Juillard, A.M. Devine, and Laurence D. Stephens, eds., *Linguistic Studies Offered to Joseph Greenberg on the Occasion of His Sixtieth Birthday*, Anna Libri, Saratoga, CA.
- Hyman, Larry M. (1982) "The representation of nasality in Gokana," in Harry van der Hulst and Norval Smith, eds., *The Structure of Phonological Representations, Part I*, Foris, Dordrecht, pp. 111-130.
- Hyman, Larry M. (1985) *A Theory of Phonological Weight*, Foris, Dordrecht.
- Ingram, David (1989) *First Language Acquisition: Method, Description, and Explanation*, Cambridge University Press, Cambridge.
- Ito, Junko, Armin Mester, and Jaye Padgett (1995) "Licensing and underspecification in Optimality Theory," *Linguistic Inquiry* 26, 571-613.
- Jakobson, Roman, C. Gunnar M. Fant, and Morris Halle (1951) *Preliminaries to Speech Analysis: The Distinctive Features and their Correlates*, MIT Press, Cambridge, MA.
- Jun, Jongho (1995a) "Place assimilation as the result of conflicting perceptual and articulatory constraints," to appear in *Proceedings of the West Coast Conference on Formal Linguistics* 14.
- Jun, Jongho (1995b) *Perceptual and Articulatory Factors in Place Assimilation: An Optimality Theoretic Approach*, Ph.D. dissertation, UCLA, Los Angeles, CA.
- Kager, René (to appear) "Stem disyllabicity in Guugu Yimidhirr," to appear in Marina Nespors and Norval Smith (eds.), *HIL Phonology Papers II*.
- Kaun, Abigail (1995a) *An Optimality-Theoretic Typology of Rounding Harmony*, Ph.D. dissertation, UCLA, Los Angeles, CA.
- Kaun, Abigail (1995b) "An Optimality-Theoretic account of rounding harmony typology," *Proceedings of the Thirteenth West Coast Conference on Formal Linguistics*, Center for the Study of Language and Information, Stanford, CA, pp. 78-92.
- Keating, Patricia (1984) "Aerodynamic modeling at UCLA," *UCLA Working Papers in Phonetics* 54, 18-28.
- Keating, Patricia A. (1985) "Universal phonetics and the organization of grammars," in Fromkin, Victoria, ed. *Phonetic Linguistics: Essays in Honor of Peter Ladefoged*, Academic Press, Orlando, FL, pp. 115-132.

- Kelly, Michael H. and Susanne Martin (1994) "Domain-general abilities applied to domain-specific tasks: sensitivity to probabilities in perception, cognition, and language," *Lingua* 92, 105-140.
- Kiparsky, Paul (1995) "The phonological basis of sound change," in John Goldsmith, ed., *The Handbook of Phonological Theory*, Blackwell, Oxford, pp. 640-670.
- Kiparsky, Paul and Lisa Menn (1977) "On the acquisition of phonology," in J. Macnamara, ed., *Language Learning and Thought*, Academic Press.
- Kirchner, Robert (in progress) *Faithfulness and Contrastiveness in Phonetically-Based Optimality-Theoretic Phonology*, Ph.D. dissertation, UCLA, Los Angeles, CA.
- Koutsoudas, Andreas, Gerald Sanders, and Craig Noll (1974) "The application of phonological rules," *Language* 50, 1-28.
- Liljencrants, Johan, and Björn Lindblom (1972) "Numerical simulation of vowel systems: the role of perceptual contrast," *Language* 48, 839-862.
- Lindblom, Björn (1983) "Economy of speech gestures," in Peter F. MacNeilage, *The Production of Speech*, Springer-Verlag, New York, pp. 217-245.
- Locke, John L. (1983) *Phonological Acquisition and Change*, Academic Press, New York.
- Locke, John L. and Dawn M. Pearson (1990) "Linguistic significance of babbling: evidence from a tracheostomized infant," *Journal of Child Language* 17, 1-16.
- Macken, Marlys A. (1980a) "The child's lexical representation: the 'puzzle-puddle-pickle' evidence," *Journal of Linguistics* 16, 1-17.
- Macken, Marlys A. (1980b) "Aspects of the acquisition of stop systems: a cross-linguistic perspective," in Grace H. Yeni-Komshian, James F. Kavanagh, and Charles A. Ferguson, eds., *Child Phonology, Volume 2: Production*, Academic Press, New York, pp. 143-168.
- Macken, Marlys A. (1995) "Phonological acquisition," in John Goldsmith, ed., *The Handbook of Phonological Theory*, Blackwell, Oxford, pp. 671-696.
- Maddieson, Ian (1984) *Patterns of Sounds*, Cambridge University Press, Cambridge.
- Maiden, Martin (1995) *A Linguistic History of Italian*, Longman, London.
- Mangold, Max (1962) *Duden Aussprachewörterbuch*, Bibliographisches Institut, Mannheim.
- McCarthy, John J. and Alan S. Prince (1994) "The emergence of the unmarked: Optimality in Prosodic Morphology," *Proceedings of the North East Linguistic Society* 24, 333-379.
- McCarthy, John J. and Alan S. Prince (1995) "Faithfulness and reduplicative identity," in Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, eds., *Papers in Optimality Theory*, University of Massachusetts Occasional Papers 18, Dept. of Linguistics, University of Massachusetts, Amherst, pp. 249-384.
- Myers, Scott (1996) "Perceptual constraints and neutralization: a direct approach," this volume.
- Odden, David (1988) "Anti antigemination and the OCP," *Linguistic Inquiry* 19, 451-475.
- Ohala, John J. (1974) "Phonetic explanation in phonology," in *Papers from the Parasession on Natural Phonology*, Chicago Linguistic Society, pp. 251-74.
- Ohala, John J. (1978) "Production of tone," in Victoria Fromkin, ed., *Tone: A Linguistic Survey*, Academic Press, New York, pp. 5-39.

- Ohala, John (1981) "The listener as a source of sound change," *Papers from the Parasession on Language and Behavior*, Chicago Linguistic Society, pp. 178-203.
- Ohala, John J. (1983) "The origin of sound patterns in vocal tract constraints," in MacNeilage, Peter F. (ed.) *The Production of Speech*. New York: Springer, pp. 189-216.
- Ohala, John J. and Manjari Ohala (1993) "The phonetics of nasal phonology: theorems and data," in Huffman, Marie K. (ed.) Krakow, Rena A. (ed.) *Nasals, Nasalization, and the Velum*, Academic Press, San Diego.
- Orr, Carolyn (1962) "Ecuadorian Quichua phonology," in Benjamin Elson, ed., *Studies in Ecuadorian Indian Languages I*, Summer Institute of Linguistics, Norman, Okla.
- Pater, Joe (1995) "Austronesian nasal substitution and other NC effects," to appear in René Kager, Harry van der Hulst, and Wim Zonneveld, eds., *Proceedings of the Workshop on Prosodic Morphology*.
- Pater, Joe (1996) "*NC," *Proceedings of the North East Linguistic Society* 26, Graduate Linguistic Student Association, University of Massachusetts, Amherst, pp. 227-239.
- Pierrehumbert, Janet B. (1980) *The Phonology and Phonetics of English Intonation*. Ph.D. dissertation, Massachusetts Institute of Technology. Distributed 1987 by Indiana University Linguistics Club, Bloomington.
- Pinker, Steven and Paul Bloom (1990) "Natural language and natural selection," *Behavioral and Brain Sciences* 13, 707-784.
- Prince, Alan and Paul Smolensky (1993) *Optimality Theory: Constraint Interaction in Generative Grammar*, ms. To appear: MIT Press.
- Rischel, Jørgen (1974) *Topics in West Greenlandic Phonology*, Akademisk Forlag, Copenhagen.
- Rood, David (1975) "The implications of Wichita phonology," *Language* 51, 315-337.
- Rothenberg, Martin (1968) *The Breath-Stream Dynamics of Simple-Released-Plosive Production*, Bibliotheca phonetica, no. 6, Karger, Basel.
- Sagey, Elisabeth (1986) *The representation of features in non-linear phonology*, Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA. Published 1990: Garland Press, New York.
- Schachter, Paul (1969) "Natural assimilation rules in Akan," *International Journal of American Linguistics* 35, 342-355.
- Selkirk, Elizabeth (1980) "Prosodic domains in phonology: Sanskrit revisited," in Mark Aronoff and Mary-Louise Kean, eds., *Juncture: A Collection of Original Papers*, Anma Libri, Saratoga, CA, pp. 107-29.
- Silverman, Daniel (1995) *Acoustic Transparency and Opacity*, Ph.D. dissertation, UCLA, Los Angeles, CA.
- Smith, Neilson (1973) *The Acquisition of Phonology*, Cambridge University Press.
- Smith, Bruce L. (1978) "Temporal aspects of English speech production: A developmental perspective," *Journal of Phonetics* 6, 37-67.
- Smolensky, Paul (1993) "Harmony, markedness, and phonological activity," ROA-37, Rutgers Optimality Archive, <http://ruccs.rutgers.edu/roa.html>.
- Stampe, David (1973) *A Dissertation on Natural Phonology*, Ph.D. dissertation, University of Chicago. Distributed 1979 by Indiana University Linguistics Club, Bloomington.

- Steriade, Donca (1993) "Positional neutralization," presentation at the 23rd meeting of the Northeastern Linguistic Society, University of Massachusetts, Amherst..
- Steriade, Donca (1995a) "Underspecification and markedness," in John Goldsmith, ed., *The Handbook of Phonological Theory*, Blackwell, Oxford, pp. 114-174.
- Steriade, Donca (1995b) *Positional Neutralization*, ms., UCLA, Los Angeles, CA.
- Steriade, Donca (in progress) "Licensing laryngeal features," ms., UCLA, Los Angeles, CA.
- Tesar, Bruce and Paul Smolensky (1993) "The learning of Optimality Theory: An algorithm and some basic complexity results," ROA-52, Rutgers Optimality Archive, <http://ruccs.rutgers.edu/roa.html>.
- Tesar, Bruce and Paul Smolensky (1995) "The learnability of Optimality Theory," *Proceedings of the Thirteenth West Coast Conference on Formal Linguistics*, Center for the Study of Language and Information, Stanford, CA, pp. 122-137.
- Traill, Anthony (1981) *Phonetic and Phonological Studies of !Xõõ Bushman*, Ph.D. dissertation, University of the Witwatersrand, Johannesburg.
- Vance, Timothy J. (1987) *An Introduction to Japanese Phonology*, State University of New York Press, Albany.
- Wells, John (1982) *Accents of English I: An Introduction*, Cambridge University Press, Cambridge.
- Westbury, John R. (1979) *Aspects of the Temporal Control of Voicing in Consonant Clusters in English*, *Texas Linguistic Forum* 14, Department of Linguistics, University of Texas at Austin.
- Westbury, John (1983) "Enlargement of the supraglottal cavity and its relation to stop consonant voicing," *Journal of the Acoustical Society of America* 74, 1322-36.
- Westbury, John and Patricia Keating (1986) "On the naturalness of stop consonant voicing," *Journal of Linguistics* 22, 145-166.

Dept. of Linguistics
UCLA
Los Angeles, CA 90095-1543

bhayes@humnet.ucla.edu