WHAT PEOPLE KNOW ABOUT SOUNDS OF LANGUAGE

Janet Pierrehumbert
Northwestern University
jbp@northwestern.edu

Generative phonology has the goal of developing a comprehensive and predictive model of implicit knowledge of sound structure. In this paper, I review the nature and import of the discoveries about implicit knowledge which were made possible by modern methods for gathering experimental data and for analyzing large data sets. First, it is known that languages differ in extremely fine phonetic detail. These systematic differences are learned by speakers and represent part of their implicit knowledge. Second, knowledge of phonological grammar gradiently reflects the frequencies of patterns in the lexicon. Third, knowledge of morphophonological alternations reflects the frequencies of word relations in the lexicon. These findings point to a model of phonology in which gradient phonetic resources are organized and exploited by languages in their lexical inventories, and the phonological grammar arises as generalizations over the lexicon.

0. Introduction

A central goal of generative linguistics is to describe and explain the implicit knowledge of language. Implicit knowledge is implicit because it is not generally accessible to conscious introspection. It is a form of knowledge because it encapsulates in some form the general principles of language which are evidenced by a speaker's ability to produce and understand language fluently and productively.

Implicit knowledge of language sound structure (that is, of phonology and phonetics) is reflected in the speaker's ability to understand novel utterances in real time, to produce novel phrases with native allophonic details, to evaluate the well-formedness of neologisms, to assimilate loan words to native sound patterns, and to extend the vocabulary through new morphological collocations. In the early days of generative linguistics, limited data sets on all of these fronts led to theoretical proposals which all shared two key assumptions, though they differed in many other respects. First, implicit knowledge of sound structure is viewed as a grammar (in the sense of formal language theory) consisting of a reasonably small number of categories and a logic for combining these categories into well-formed wholes. In a highly simplified model, the categories of a language are phonemes. A grammar specifies how the phonemes may be combined into well-formed syllables, and a word is viewed as a sequence of one or more syllables. In more sophisticated views, distinctive features rather than phonemes are the most elementary building blocks of sound structure; regularities at other prosodic levels such as the foot and the intonation phrase also

contribute to well-formedness; and morphophonological alternations are also described in grammatical terms.

A second key assumption is that the relationship between outputs of the phonological grammar to physical properties of the speech signal lies outside of implicit phonological knowledge proper, being instead describable as a universal — language-independent — interface. A particularly clear statement of this assumption is found in Chomsky & Lasnik's 1995 synopsis of the Minimalist Program. In the Minimalist framework , the level of representation PF ('Phonetic Form') is characterized as symbolic, universal, and supporting a uniform interface to the sensorimotor system (Chomsky 1995:21). Discussion in Chomsky 1993 reiterates that characterizing PF means discovering what objects have a uniform language-independent interpretation. He states 'We must spell out explicitly what are the legitimate objects at PF and LF. At PF, this is the problem of universal phonetics'. (Chomsky 1993: 26-7). In taking this stand, Chomsky reiterates the position of Chomsky & Halle 1968.

In this talk, I will show how these two key assumptions must yield in face of the large- scale data sets which can now be analyzed using modern technology and experimental methods. These data sets — which differ more in scale than in kind from the data sets traditionally analyzed — display general patterns which any model of implicit knowledge of language must account for. Specifically, they demonstrate the existence of language-specific phonetic patterns down to extremely fine levels of detail, most naturally described using continuous mathematics rather than an inventory of phonetic categories such as the IPA. Second, they demonstrate that implicit knowledge of both phonotactics and morphophonological alternations is stochastic rather than purely grammatical. Both of these findings come together in a picture in which language patterns are learned through statistical generalizations over numerous examples.

1.   Phonetic detail

As speech workstations have become more and more widespread, more and more comparative studies of phonetic implementation have been published. A very large literature now permits a definitive assessment of the claim that universal phonetics is a critical characteristic of the human capability for language. In this section, I present only a few highlights of this literature.

Bradlow 1995 describes a comparative study of English and Spanish vowels. Target vowels were located in the the most comparable consonantal and prosodic contexts possible. Formant measures of even point vowels differed systematically, with English /i/ being more high and front than Spanish /i/ while Spanish /u/ is more close to cardinal /u/ than English /u/ is. Bradlow provides clear arguments that these differences cannot be explained by overall differences in vocal tract size between her English subjects and her Spanish subjects, so they must instead reflect learned details of pronunciation.

VOT (voice onset time) of stops provides one of the stronger candidates for a universal system of phonetic categories. As described in Keating's 1984 cross-linguistic survey of stops, three qualitatively different outcomes are provided by prevoiced stops (stops with negative VOT values), short-lag stops (voiceless unaspirated stops) and long-lag stops (voiceless aspirated stops). In a language such as French, the phonologically critical

boundary is that between prevoiced stops and short-lag stops. For word-initial stops of English, the critical boundary is between long-lag stops and short-lag stops. So-called voiced stops of English can surface phonetically as either voiceless unaspirated or as prevoiced stops.

The existence of these qualitative patterns is not enough, however, to support the kind of universal uniform phonetic interface that Chomsky & Lasnik specify. If a uniform phonetic interface really existed, the outcomes would be identical not just qualitatively, but even quantitatively. In fact, the exact extent of prevoicing and aspiration differs across languages. Differences are found both in the distribution of values in analogous prosodic positions and in the way that values are affected by prosodic and segmental context. All of these details have been learned by anyone who is perceived to have a native accent in a language. An example is provided by Caramazza & Yeni-Komshian's 1974 study of voice onset time in European French, Canadian French, and Canadian English. Their data on European French and Canadian English correspond to the patterns generally reported. The distribution of values for Canadian French, however, matched neither the European French pattern nor the Canadian English pattern. For /p/, the VOT values of Canadian French speakers fell into the voiceless unaspirated range of European French, but for /b/, the Canadian French speakers had more tokens with zero to small positive VOT values, and fewer examples of large negative VOT values than did the European French speakers. This means that the combined distribution of VOT values for /b/ and /p/ is more concentrated towards the middle of the total available range in Canadian French than in European French or Canadian English.

A comparison of French and English also shows that the exact extent of contextually induced allophony is language-specific, even when the allophonic effect in question is broadly analogous across languages. Flege & Hillenbrand 1986 collected production and perception data on durational reflexes of the post-vocalic /s/-/z/ distinction. In the English word *peas*, the vowel is longer and the fricative is shorter than in the word *peace*. These durational differences are sometimes reported to be universal, and provide effective cues in speech perception to the phonological category of the final fricative. Flege & Hillenbrand showed that French speakers have the same kind of durational differences as English speakers but that the exact quantitative character of these differences is not the same. In English, the extent of vowel lengthening before voiced obstruents is greater than in French. In French, the length of the fricative itself carries more of the information about voicing status than in English. The results of a perception test on a *peace/peas* stimulus set mirrors these differences. For the same stimulus set, English speakers rely more on the vowel duration and French speakers rely more on the duration of frication in categorizing the stimulus.

Turning to the domain of tone and intonation, a comparison of results on English from Liberman & Pierrehumbert 1984 with results on Japanese in Pierrehumbert & Beckman 1988 also reveals language-specific principles of phonetic implementation. Both English and Japanese have phrase final L% boundary tones in the default declarative phrasal intonation. This L% is temporally aligned at the very edge of the phrase after the last accent. When Liberman & Pierrehumbert asked subjects to produce declarative patterns in many different overall voice levels, they found that the L% achieved a fixed value for each speaker regardless of the (considerable) effect of voice level on any preceding H tone. They

hypothesized that the floor of the speaker's voice is a physiologically determined parameter of the phonetic system (the baseline) and that implementation rule for L% realizes the tone on the baseline. A comparable experimental paradigm for Tokyo Japanese yielded contrasting results, according to Pierrehumbert & Beckman. The L% did not exhibit a fixed realization for each speaker. Instead, its f0 value was a cumulative function of the overall pitch range and the occurrence of downstep-inducing accents earlier in the same phrase. Thus, the L% in Tokyo Japanese appears to be realized in relation to the preceding H level rather than in an invariant position. Pierrehumbert & Beckman also discuss differences between Tokyo and Osaka Japanese, by which the same underlying phonological tonal sequences turn out differently because of systematic differences in L tone scaling.

I would like to take a strong stand on phonetic implementation at this point. The studies I have mentioned are not isolated examples. On the contrary, I believe that every thorough study which has looked for a difference between two languages in details of phonetic implementation has found one. These differences concern both detailed outcomes for analogous phonemes in the most analogous available positions, and — to an even greater extent — principles of allophonic variation in context. As a result of these findings, the level of representation posited in Chomsky & Lasnik is not a scientific possibility at this point. Not only do some phonological entities fail to meet the conditions they lay out, there is no known case of a phonological entity which does meet these conditions. To explain the extremely detailed but extremely systematic patterns which characterize the native phonetics of any language, it is necessary to posit learning mechanisms which can aquire quantitative distributions of phonetic outcomes. These mechanisms form part of the human endowment for language. A key feature of such mechanisms is that they acquire patterns by generalizing statistically over many examples. As we will see in the next sections, this same feature is also found at more abstract levels of description.

2.   Phonotactics

A central concept in generative phonology is that of the 'possible word'. In an introductory phonology course, the instructor may introduce the entire subject matter by pointing out that any native speaker of English can judge a form such as *blick* to be a possible word of English — even though it does not happen to exist in the current English lexicon — whereas *bnick* is not judged to be a possible word. In the theoretical treatment, *blick* is possible because it is a wellformed combination of well-formed subparts, but *bnick* is not possible because it contains an onset cluster which does not exist in English. In constructing the specifics of such a theory, two lines of evidence are brought into convergence. One is the intuitive well-formedness judgments of native speakers. The other is the existence of systematic gaps in the lexicon. The fact that no words at all begin in /bn/, although many other sequences of obstruents and sonorants are found, provides evidence that /bn/ is a systematic gap to which the phonological theory is accountable.

The idea of a systematic gap already brings us to the realization that phonotactic constraints are generalizations over what exists (and what doesn't exist) in the lexicon. In this section, I first discuss evidence that such generalizations are stochastic rather than purely categorical. Then, I will draw a connection between well-formedness judgments and these lexical statistics, showing that they are highly related.

The calculations presented in Pierrehumbert 1994 provide an inventory of phonological patterns which differs only in scale, and not in kind, from the data in a traditional phonology problem set. The study concerns the inventory of long medial clusters (containing three or more consonants) within monomorphemic words. Examples of words containing such clusters include *velcro* and *doldrums*. The number of attested clusters is remarkably small and it is easy to make up nonsense forms involving long medial clusters which anyone would judge to be impossible except as a compound. A particular theoretical challenge is posed by the fact that many of the missing long clusters contain only subsequences which are well-formed under any reasonable syllabic parse. For example, the hypothetical word *pelskra* contains a medial sequence whose well-formedness is supported by words such as *else, ascii, screw,* and *crow*. A standard model — in which words are made up of syllables — would predict that any medial cluster containing an allowable coda followed by an allowable onset would be possible, modulo word-level constraints on the syllable contact.

The analysis in Pierrehumbert 1994 is based on a complete inventory of medial clusters of three or more consonants contained in the on-line Collins dictionary distributed by the Linguistics Data Consortium; this dictionary is a very large one, containing about 70,000 words. Words with long clusters were extracted automatically and the monomorphemic words were identified by individual scrutiny. A rather generous definition of monomorphemicity was used, including words such as *constrain* and *abstract* which are probably not synchronically decomposed by average speakers. This list of occurring clusters was compared to the list of clusters which would be generated as the crossproduct of allowable codas and allowable onsets. The expected likelihood of each complex cluster was estimated as the likelihood of the random combination of the coda and onset, given their rates of occurrence as word onsets or codas of final syllables (with appendices stripped off). This extremely crude method of estimating the likelihoods nonetheless led to a surprising conclusion.

More than 8700 medial clusters arise from the crossproduct. Only about 40 are found, and these are almost all in the most likely 200. Within the most likely 200, the percentage of candidate clusters that actually occur peaks at 47% (for the top 40 candidates) and dwindles steadily thereafter. Overall, probability was the single most successful predictor by far of whether a complex cluster will occur or not, eliminating more than 8500 candidates. A standard nonstochastic model would need to rule out these numerous and diverse cases through a battery of constraints which would miss the main generalization, which is so straightforward it is almost tautological. If a cluster is too unlikely to occur in the 70,000 item sample which represents an adult vocabulary, then it probably will not occur. This overarching generalization provides a straightforward argument that probabilities are inside the grammar.

Turning now to well-formedness judgments, a experiment described in Treiman et al. 2000 demonstrates a connection between these judgments and the statistics of the lexicon. Through computations over an on-line dictionary, Treiman et al. identified pairs of VC combinations (V1 C1 versus V2 C2) which could be transformed into less frequent, but still attested, combinations by switching the consonants (e.g. V1 C2 and V2 C1). By constructing monosyllabic nonsense words involving these combinations, they were able to control strictly for phoneme frequency and determine whether subjects had implicit knowl-

edge of the collocational frequencies. Results of both well-formedness judgments and performance on a blending task revealed that subjects did indeed treat the higher frequency combinations as better than lower frequency combinations.

Since the experimental design in Treiman et al. involves only a two way comparison, the nature of the function relating lexical statistics to wellformedness is not mapped out. A systematic exploration of nasal-obstruent clusters (NO clusters) is presented in Hay et al. (forthcoming). These clusters were selected both because they span the range rather evenly from impossible to extremely frequent. Although linguistic texts may undertake to describe the patterns using a nasal homorganic rule which requires the nasal to agree with a following obstruent in place of articulation, the situation is actually far more complicated. Lexical statistics show that this requirement is stronger for stops than for fricatives. Among nasal-fricative combinations, the inhomorganic /nf/ is actually more likely than the homorganic /mf/. Three series of nonsense words were constructed which varied only in the medial NO cluster. (e.g. *strinty*, *strinsy*, *strimpy*, etc.; *zanter*, *zanser*, *zamper*, etc; *krenter*, *krenser*, *kremper*, etc.). The same large range of NO clusters was used in all series. None of these words had a real word imbedded at the beginning, though it is impossible to avoid medially embedded words (e.g. *imp*, *ant*). Recordings of the words were generated by cross-splicing syllables recorded in homorganic environments.

Subjects transcribed and rated each nonsense word. Both the transcription data and the rating data showed gradient effects of lexical frequency. In the transcription data, subjects showed a tendency to misperceive infrequent clusters as acoustically similar but more frequent ones. For example, /np/ is often misheard as /mp/. However, infrequent clusters were often transcribed correctly. The wellformedness judgments then proved to be a linear function of the statistically best morphological parse. For example, for *strinpy*, the best morphological parse is one which imputes a word boundary between the /n/ and /p/; the rating of such a (compound) form depends on the likelihood of /n/ as a word-final coda and /p/ as a word beginning. For *strinty*, the best morphological parse is as a monomorphemic word, and the perceived wellformedness reflects the frequency of /nt/ as a medial cluster in a monomorpheme. The relationship of likelihood to perceived well-formedness appeared to be linear.

In summary, then, including pattern statistics inside the grammar allows us to capture generalizations about systematic gaps in the lexicon, and to make predictions about possible and impossible words. Experiments indicate that well-formedness judgments are gradiently related to lexical statistics. Implicit knowledge of lexical statistics is also demonstrated by its impact on morphological parsing, speech perception, and performance on creative tasks such as word blending.

3.   Morphophonological alternations

Phonotactic and prosodic constraints represent generalizations over words which describe the relative well-formedness of potential word shapes. As such, they do not cover one of the main areas of phonological investigation, namely the morphophonological alternations which arise in words related to each other through affixation or other morphological processes. For example, *serenity* and *seronity* are both perfectly well-formed as potential English words; prosodic and phonotactic theory does not in any way specify that *serenity* is

a complex form related to *serene*, whereas *seronity* is not. Relations such as these are the centerpiece of classical generative phonology, going back to Chomsky a Halle 1968.

A basic point to note about such phenomenona is that they are relations amongst words. Insofar as they are general and productive, we have evidence for implicit knowledge of relationships amongst things, a higher level of abstraction than mere implicit knowledge of things. Having moved one level up, however, we find that experimental results support exactly the same sort of observations that we just made about lower levels of description.

A pioneering paper in this area is Cena 1978, who explored the psychological reality of the vowel shift rule proposed in Chomsky & Halle 1968 using a concept generalization paradigm. He found that vowel shifting is most readily extended for the more frequent vowel pairings than for the least frequent pairing (the *profound/profundity* pairing, for which only 6 examples exist). The possible intrusion of orthographic factors make this result difficult to interpret, however. A more clearcut finding resulted from Bybee & Pardo's 1981 experimental study of verb paradigms in Spanish. They found that Spanish subjects extend to novel verbs the inflectional patterns which are attested for numerous real verbs, but not ones attested only for a few high-frequency verbs. Thus, what matters is the number of different examples of the alternation in the lexicon, and not the token frequency of the actual forms. This supports the idea that morphophonological processes are generalizations over entries in the lexicon. Further support for this viewpoint may be adduced from Marchman & Bates 1994 finding that English speaking children begin to use overregularized past tense forms (such as *teached*) at the time at which they have acquired a critical number of different regular past tense forms.

A clever experimental study described in Ohala & Ohala 1987 brings together results on a number of different morphophonological alternations. The stimuli in this experiment were 20 word pairs exemplifying fairly commmon alternations (such as *particle/particular*, *substance/substantial*) and 20 words pairs representing isolated patterns (such as *thumb/ thimble*, *strong/stringent*, and *slay/slaughter*). 16 subjects provided three ratings for each pair on a five-point scale. They rated phonetic similarity, semantic relatedness, and derivational relatedness, where the concept of derivational relatedness was operationalized by asking how likely the words were to have a common historical ancestor. Phonetic similarity did not prove to be a predictor of derivational relatedness in this study, possibly because the range of variation in this parameter was not sufficiently great or because it was not possible to control it well enough given the many other constraints on the stimulus set. Derivational relatedness was strongly predicted by semantic relatedness. But the derivational relatedness scores differed depending on whether the alternation in question was common or isolated; for the same degree of semantic relatedness, derivational relatedness was about one point higher for common pairings than for isolated ones. Thus, we see that the number of different examples of a word relationship influences the extent to which it is cognitively real.

If morphophonological alternations are relationships amongst words, and if their cognitive status critical depends on the number of examples of the relationship in the lexicon, it follows that there should be considerable individual differences in this area of phonology. Different people have different vocabularies, due to differences in upbringing,

work environment, and verbal ability. Many morphophonological alternations that have occupied linguists are attested mainly in low frequency words, which a given speaker may or may not have learned. This prediction, though possibly unwelcome to some phonologists, is borne out in two important studies.

A careful and ingenious experiment described in McCawley 1986 explores the status of the English Vowel Shift. The stimuli in this study were pairs of morphological related words, such as *serene/serenity*, *explode/explosion* and *rot/rotten*, and a control set of unrelated words of comparable phonological similarity, such as *bone/bonnet*, *lie/light*, and *mouse/mustard*. The morphologically related words included word pairs exemplifying Chomsky & Halle's 1968 vowel shift rule, word pairs exhibiting an anti-vowel-shift relationship (such as *peace/pacify*), and baseline pairs in which the vowel was identical in the two words, such as *rot/rotten*. Subjects judged the semantic relatedness of the words in each pair and also whether the putative base was 'contained in' the longer form. The data analysis concentrates on words with reasonably high semantic relatedness, because in the absence of such relatedness speakers do not judge the shorter form to be contained in the longer one. The morphological relatedness is measured by the 'contained in' judgments.

McCawley takes the Vowel Shift rule to be psychologically real if the pairs exhibiting this relationship had morphological relatedness scores comparable to pairs with identical vowels. In this case, he argues, the difference in vowel quality is fully transparent and therefore the pairing of the vowels is fully active. This was the case for the majority group of his subjects. For a minority group, the morphological relatedness was significantly less for vowel shift pairs than for pairs with identical vowels. The combined data set is strongly bimodal, indicative of a situation in which some people have the rule and others do not. Since McCawley's subjects were University of Chicago undergraduates, who generally display a high level of verbal ability, we would not necessarily expect the proportion of vowel-shifters in the general population to be as large. The existence of two modes in the distribution, however, strongly confirms the prediction of individual variation in the morphophonological part of grammar.

Lastly, I'd like to present an experimental study by Steriade 2000 which both furthers the point of this section and connects it to the starting point of this paper. Steriade investigated the alternation between aspirated /t/ and tap in English. This alternation is regular in word pairs such as *platonic/Plato*. In *platonic*, the /t/ begins a stressed syllable and is aspirated. In *Plato*, the /t/ is in post-stress position (intervocalic between a stressed vowel and an unstressed one), and it is tapped in fluent speech. (It may of course be aspirated in the most hyperarticulated or corrective possible pronunciations in which the second syllable is not fully unstressed). In post-post-stress position, however, tapping is variable. In a word such as *positive*, the /t/ may or may not be tapped in fluent speech.

Steriade assessed patterns of tapping for twelve subjects. The experimental word pairs in the study were morphologically related pairs in which /t/ occurred in post-post-stress position in both words. An example is *positive/positivistic*. Control pairs had the /t/ in the obligatory tapping environment, e.g., *rotary/rotaristic*. Steriade found that subjects varied considerably in which experimental words they produced with taps. However, 11 out of twelve subjects had identical tapping within every word pair.

This result brings us back to the beginning of the paper, in that it shows systematicity in a area of fine phonetic detail. The nature of this systematicity — allophonic effects specific to members of a morphological family — cannot be captured by a model with uniform, universal phonetic implementation rules. Indeed, the result goes beyond those previously presented by demonstrating a pattern specific not just to a language, but to words within language. It also ties in with the present section by illustrating again individual variation in relations amongst words. Even an apparently low-level alternation — namely tapping — appears to be organized cognitively in terms of relations amongst words, and in consequence it exhibits individual variation.

4.   Conclusion

In conclusion, then, languages can differ systematically in arbitrarily fine phonetic detail. This means we do not want to think about universal phonetic categories, but rather about universal phonetic resources which are organized and harnessed by the cognitive system. Water is a physical resource. It is used in different ways in different cultures. A desert culture may transport water in jars and develop a custom of using water to wash feet in a ritual of welcome or deference. Americans move water in pipes and hoses and have developed a ritual of using it to wash cars. The vowel space — a continuous physical space rendered useful by the connection it establishes between articulation and perception — is also a physical resource. Cultures differ in the way they divide up and use this physical resource.

Learning of fine phonetic detail requires the ability to form statistical generalizations over large classes of speech tokens. The ability to form statistical generalizations is recapitulated at higher levels. Phonotactic constraints arise as statistical generalizations over the lexicon, and lexical statistics are gradiently reflected in well-formedness judgments. The knowledge of morphophonological alternations similarly reflects the frequency of word relations in the lexicon.

These findings all point towards a theory in which gradient phonetic resources are organized and exploited by languages in their lexical inventories, and the phonological grammar arises as generalizations over the lexicon.

REFERENCES

BRADLOW, A. 1995. A comparative acoustic study of English and Spanish vowels. *Journal of the Acoustical Society of America* 97:3.1916-24.

BROE, M., & J. PIERREHUMBERT (eds.). 2000. *Papers in Laboratory Phonology* V: *Acquisition and the Lexicon*. Cambridge: Cambridge University Press.

BYBEE, J., & E. Pardo. 1981. Morphological and lexical conditioning of rules: Experimental evidence from Spanish. *Linguistics* 19.937-68.

CARAMAZZA, A., & G. H. YENI-KOMSHIAN. 1974. Voice onset time in two French dialects. *Journal of Phonetics* 2.239- 45.

CATFORD, J. C. 1977. *Fundamental Problems in Phonetics*. Bloomington: Indiana University Press.

CENA, R. M. 1978. When is a phonological generalization psychologically real? Indiana University Linguistics Club. As summarized in McCawley 1986.

CHOMSKY, N. 1993. A minimalist program for linguistic theory. *The View from Building 20*, ed. by K. Hale & S. J. Keyser. Cambridge, MA: MIT Press.

——. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.

——, & H. LASNIK. 1995. The theory of principles and parameters. In Chomsky 1995, 13-128.

——, & M. HALLE. 1968. *The Sound Pattern of English*. New York: Harper & Row.

FLEGE, J. E., & J. HILLENBRAND. 1986. Differential use of temporal cues to the /s/ - /z/ contrast by native and non-native speakers of English. *Journal of the Acoustical Society of America* 79:2.508-17.

KEATING, P. (1984) Phonetic and phonological representation of consonant voicing. *Language* 60.286-319.

LIBERMAN, M., & J. PIERREHUMBERT. 1984. Intonational invariance under changes in pitch range and length. *Language Sound Structure*, ed.by M. Aronoff & R. Oehrle, 157-233. Cambridge MA: MIT Press.

MARCHMAN, V., & E. Bates. 1994. Continuity in lexical and morphological development: A test of the critical mass hypothesis. *Journal of Child Language* 21.339-66.

MCCAWLEY, J. D. 1986. Today the world, tomorrow phonology. *Phonology Yearbook* 3. 27-45.

OHALA, J. 1987. Experimental phonology. *Proceedings of the Berkeley Linguistic Society* 13:207-22.

OHALA, M., & J. OHALA. 1987. Psycholinguistic probes of native speakers phonological knowledge. *Phonologica 1984*, ed. by W. U. Dressler (3). Cambridge: Cambridge University Press.

PIERREHUMBERT, J. 1994. Syllable structure and word structure. *Papers in Laboratory Phonology* III, ed. by Keating & Hayes, eds. Cambridge: Cambridge University Press.

——, & M. BECKMAN. 1988. *Japanese Tone Structure* (Linguistic Inquiry Monograph 15.) Cambridge, MA: MIT Press.

——, J. HAY, and M. BECKMAN. [Forthcoming.] Speech perception, well-formedness, and lexical frequency. *Papers in Laboratory Phonology* VI. Cambridge: Cambridge University Press.

STERIADE, D. 1999. Paradigm uniformity and the phonetics-phonology interface.In Broe & Pierrehumbert 2000: 313-34.

TREIMAN, R., B. KESSLER, S. KNEWASSER, R. TINCOFF, & M. BOWMAN. 2000. English speakers' sensitivity to phonotactic patterns. In Broe & Pierrehumbert 2000:269-82.