Phonological Naturalness and Phonotactic Learning

Bruce Hayes
James White

We investigate whether the patterns of phonotactic well-formedness internalized by language learners are direct reflections of the phonological patterns they encounter, or reflect in addition principles of phonological naturalness. We employed the phonotactic learning system of Hayes and Wilson (2008) to search the English lexicon for phonotactic generalizations and found that it learned many constraints that are evidently unnatural, having no typological or phonetic basis. We tested 10 such constraints by obtaining native-speaker ratings of 40 nonce words: 10 violated our unnatural constraints, 10 violated natural constraints assigned comparable weights by the learner, and 20 were control forms. Violations of the natural constraints had a powerful effect on ratings, violations of the unnatural constraints at best a weak one. We assess various hypotheses intended to explain this disparity, and conclude in favor of a learning bias account.

Keywords: phonotactics, nonword acceptability, maximum entropy, learnability, learning bias

1 Introduction: The Problem of Unnatural Constraints

Our starting point is a classic phonological problem, the origin of phonotactic knowledge (Chomsky and Halle 1965). Speakers can rate novel words of their language, judging them to be fully acceptable (e.g., *blick* [blɪk]), intermediately acceptable (*bwick* [bwɪk]), or fully ill-formed (*bnick* [bnɪk]). Not only their intuitive judgments, but also their behavior, reflects such hierarchies of well-formedness, as is shown by experimental evidence from speech production and perception (Massaro and Cohen 1980, Dupoux et al. 1999, Mattys and Jusczyk 2001, Moreton 2002, Berent et al. 2007, Berent et al. 2008, Wilson and Davidson 2009). Patterns of phonotactic well-formedness are partly language-specific and therefore must, at least to some extent, be learned during the period of phonological acquisition.

   The problem of phonotactic learning is particularly suited to the strategy of computational modeling (Hayes 2004, Prince and Tesar 2004, Jarosz 2006, Coetzee and Pater 2008, Hayes and Wilson 2008, Albright 2009, Heinz 2009, 2010). A model can be set up to implement a variety

of hypotheses about the language faculty as it relates to phonology. Ideally, it should be fed a representative lexicon given in phonetic transcription, approximating the experience of the language-learning child. The model learns a phonotactic grammar, which can then be tested by comparing the well-formedness values it assigns to novel stimuli with well-formedness measures obtained experimentally.

An informative strategy for such work is the ''inductive baseline'' approach (e.g., Gildea and Jurafsky 1996, Hayes and Wilson 2008). The idea is to start with very simple models, embodying few a priori principles, and see where they fail. When augmenting the models with principles of phonological theory produces success instead, we obtain insight into the usefulness of such principles for learning.

A baseline model of this type is proposed by Hayes and Wilson (2008). We review the details of this model below; for present purposes, the crucial aspect of the model is that in its core rendition, the a priori knowledge that it brings to phonological learning is largely confined to the feature system. This serves as the basis for phonological constraints, which the model constructs by concatenating feature matrices that denote natural classes. For example, the constraint that bans prevocalic lax vowels in English could be stated as *[+syllabic,−tense][+syllabic].

Hayes and Wilson (2008) argue that their model does a fairly good job of locating in some form the generalizations that linguists find when they inspect phonotactic patterns. For example, it finds the principles of featural agreement that govern Shona vowel harmony, and it replicates essentially all the phonotactic generalizations proposed by Dixon (1981) in a meticulous phonotactic study of Wargamay, an Australian language. The model is also successful in matching human phonotactic intuitions: it achieves a close match to the experimentally gathered intuitions on English onset well-formedness collected by Scholes (1966); and more recently, it has achieved a reasonably good match for the experimental data reported by Albright (2009), Colavin, Levy, and Rose (2010), and Daland et al. (2011).

However, another aspect of the Hayes/Wilson model is potentially far more controversial. When fed with Wargamay data, the model did not learn just the constraints that were needed to recapitulate Dixon's well-motivated analysis; it also learned a number of phonological constraints that would strike experienced phonologists as unnatural. One example is given in (1), stated first in features and then in prose. The symbol ''^'' may be read ''unless.''

(1) *A puzzling constraint learned for Wargamay*

   a. $*\begin{bmatrix} ^\wedge -\text{son(orant)} \\ -\text{ant(erior)} \end{bmatrix}$ [+long][−cons(onantal)]

   b. ''If a long vowel is followed by a glide, it must be preceded by a palato-alveolar obstruent.''

Hayes and Wilson point out two possibilities concerning constraints like (1). One is that they are indeed valid for Wargamay: were it possible to access native-speaker intuitions in this language, forms violating them would be judged as ill-formed to a degree corresponding to the weight of the applicable constraint. Another possibility, however, is that these constraints reflect

a defect in the learning model: a constraint could be entirely exception-free in the Wargamay lexicon, yet fail to be implemented by native speakers as part of their phonological grammar. The purpose of this article is to offer evidence from a phonological experiment that bears on which of these two hypotheses is more likely to be correct. For practical reasons, we shift our focus to English, where the Hayes/Wilson model also finds unnatural-seeming constraints like (1).

Before starting in, we clarify two terms to be used below.

Phonological constraints are usually defended on two grounds: either typological or phonetic. The typological criterion can be expressed on the basis of Greenbergian implicational universals (the presence of sequences that violate the constraint implies the presence of closely similar sequences that do not; Greenberg 1966, 1978). The phonetic criterion is that a constraint should be functionally effective, serving to form a phonological system in which words are easier to articulate or in which possible words are perceptually distinct from one another. We will refer to constraints that satisfy one or the other criterion as *natural*, and to other constraints as *unnatural*. It is the unnaturalness of constraint (1) that would render it suspect for many phonologists, we think, as a valid constraint of Wargamay phonology.

We will call a constraint *accidentally true* if it holds true of a language's lexicon but experimental investigation indicates that it is not part of the phonotactic knowledge of native speakers. One possible ground for suspecting a constraint of being accidentally true is that it is unnatural. Thus, in these terms, the purpose of our experiment is to test whether some of the unnatural constraints learned by the Hayes/Wilson learner are accidentally true.

2 Research Background

The problem of unnatural phonotactic constraints is relevant to a current debate in phonology. Is phonological learning the result of an unbiased, inductive search for generalizations (see, e.g., Blevins 2004:chap. 9)? Alternatively, are language learners limited to learning only generalizations that are expressible with a limited, universal set of natural constraints (see, e.g., Becker, Ketrez, and Nevins 2011)? If the former position is correct, then language learners should in principle be able to access and employ unnatural generalizations.[1]

Experimental evidence bearing on this issue is steadily accumulating. Our reading of this literature is that the evidence is quite mixed and gives no comfort to advocates of either of the two possible extreme positions (all constraints are a priori knowledge/all learning is purely inductive).

2.1 Evidence for Learnability of Unnatural Generalizations

A fairly clear case of an evidently learnable generalization is the English rule of Velar Softening (Chomsky and Halle 1968:219–221). This rule is argued to be unnatural by Pierrehumbert (2006);

---

[1] In the inductivist view, the typological patterns that manifest natural principles are attributed instead to diachronic factors: languages change phonologically through phonetic shifts and misperceptions that are sensitive to phonetic or other naturalness principles (see, e.g., Ohala 1981, Blevins 2004).

notably, it derives [s], rather than the phonetically/typologically expected [tʃ], from /k/. Pierre-humbert's (2006) experiments demonstrate that Velar Softening is surprisingly productive.

Further, processes of phonology or allomorph selection have been shown to apply differentially in a way that is moderated by unnatural factors in the phonological environment. One such case is found in Hungarian vowel harmony (Hayes et al. 2009). The harmony pattern is mostly predictable, but with certain vowel sequences harmony is partly arbitrary, with front or back harmony occurring on a stem-by-stem basis. Corpus data for Hungarian show a statistical skewing, favoring front harmony for stems ending in bilabial stops, an environment that would not qualify as natural by either phonetic or typological criteria. Hayes et al. find that Hungarian speakers are tacitly aware of this unnatural pattern and several others, respecting them when they apply harmony in words with nonce stems.

Competition between morphological processes is likewise often affected by unnatural factors present in the phonological environment, part of a phenomenon that Albright (2002) calls ''islands of reliability.'' For instance, every English verb stem that ends in a voiceless fricative takes the regular past tense ending. When tested with nonce stems, English speakers show that they particularly prefer regular past tense suffixation (/-d/) for stems of this type (Albright and Hayes 2003).

Finally, experimenters have created novel languages and tested whether unnatural phonological patterns could be learned from them. Often, such studies strengthen their case by comparing the learnability of a particular phonological pattern with its opposite (both cannot be natural). For instance, Onishi, Chambers, and Fisher (2002) compared artificial languages in which {b, k, m, t} were limited to onset and {p, g, n, tʃ} to coda—or vice versa. Both phonotactic systems were learnable by adults, and also by 16.5-month-old infants (Chambers, Onishi, and Fisher 2003). Related work (Dell et al. 2000, Warker and Dell 2006, Warker et al. 2008) found similar learning patterns using a speech-error testing paradigm. Adult learners have been able to learn vowel disharmony about as well as (typologically far more common) vowel harmony, in an artificial-language study (Pycha et al. 2003) and in a study where the vowel harmony or disharmony rule was used to create a novel ''dialect'' of French (Skoruppa and Peperkamp 2011). Unnatural consonant alternations (/p, g/ → [ʒ, f] / V _____ V and /ʃ, v/ → [b, k] / V _____ V) were successfully learned by participants in the artificial-language learning study of Peperkamp and Dupoux (2007). Seidl and Buckley (2005) found that 9-month-old infants could learn both phonetically natural patterns and similar unnatural patterns.

2.2 Evidence for the Role of Naturalness

We next address the opposite possible extreme position: that *all* phonological learning is purely inductive and that naturalness considerations play no role. This strong position is likewise contradicted by evidence in the literature. For instance, Wilson (2006) showed that learners of an artificial language extended a palatalization alternation in one direction (learn forms with /ke/ → [tʃe], test on forms with /ki/ → [tʃi]) but not in the other (learn /ki/ → [tʃi], test on /ke/ → [tʃe]); this finding matches both language typology and the predictions of Wilson's phonetic model. Experiments have also shown that participants learn a dependency between the height of two vowels more easily than a dependency between the voicing of a consonant and the height

of a vowel (Moreton 2008) and that they learn a directional vowel harmony pattern over a ''majority rules'' pattern when presented with ambiguous training data (Finley 2008, Finley and Badecker 2008). Other experiments that have supported enhanced learnability for natural phonological patterns are reported by Pater and Tessier (2003), Wilson (2003), Peperkamp, Skoruppa, and Dupoux (2006), Berent et al. (2007), Berent et al. (2008), Berent et al. (2009), and Hayes et al. (2009).

In several of the experiments just cited, the findings support a *bias* effect: the unnatural patterns are learnable but take longer to learn, or yield weaker experimental effects than comparable natural patterns. We will return to the question of bias below.

2.3 Overview

To sum up: at present, purist ''all naturalness'' and ''no naturalness'' positions seem ill-supported, but the articulation of a theory explaining how and when naturalness plays a role in phonological learning lies in the future. We hope to contribute to this debate by addressing one particular angle of the problem, one for which the Hayes/Wilson learning model can play a useful role. As noted above, the model learns unnatural-seeming constraints when applied to Wargamay. We have since found the same for English (see below) and believe the model would almost certainly find similar constraints when applied to most other languages. Our interest here is not so much the details of the Hayes/Wilson learner but its characteristic behavior: through extensive search, it discovers phonological ''gaps'' (unpopulated regions) in a lexical corpus, and the constraints it uses to describe these gaps often appear to be unnatural. The existence of such gaps is in one sense a fact about the lexicon, rather than about the learner itself. The data patterns are there to be discovered; the question is whether native speakers find them.

In what follows, we first review the workings of the Hayes/Wilson phonotactic learner (section 3), then discuss how we employed it with English data to discover a variety of unnatural constraints (section 4). We then describe how we selected the test words and obtained ratings of them in an experiment (section 5). Our main result, that the unnatural constraints have little or no effect on native-speaker intuition, is given in section 5.2. We defend our claim against possible confounding effects in section 6. In section 7, we assess what our results mean for the naturalness debate and suggest how research might proceed from here.

3 The Hayes/Wilson Phonotactic Learner

In broad outline, the Hayes/Wilson learner forms a space of possible constraints using a feature system given to it in advance. It selects constraints from this set using ranked heuristics, and weights them using the criterion of maximum likelihood. The final grammar learned can assign a likelihood value to any given string, forming a quantitative prediction about phonotactic well-formedness. We elaborate this picture below.

In its simplest form, the model uses *SPE*-style phonological representations (Chomsky and Halle 1968) consisting of sequences of feature matrices, and assumes that constraints likewise consist of matrix sequences banning particular sequences of natural classes. A key observation is that although the number of possible feature matrices in any reasonable feature system is

extremely large,[2] the number of natural classes these matrices define on a segment inventory is far smaller, typically in the hundreds. This makes it feasible to do exhaustive searching of the class of possible constraints, provided the maximum number of matrices in a constraint is not too high. We assume this number is at least three (e.g., constraints applying to intervocalic consonants are abundant); for discussion, see Hayes and Wilson 2008:sec. 4.1.2 and Kager and Pater 2012.

The Hayes/Wilson model iteratively searches for new constraints to add to the grammar, selecting from the full set of possibilities using a hierarchy of heuristics. The top-ranked heuristic is matrix count: unigram constraints are favored over bigram, bigram over trigram. The second-ranked heuristic is accuracy: constraints are preferred if they are violated by very few forms, relative to the number of forms that would be expected to occur given whatever constraints and weighting have been learned so far. The lowest-ranked heuristic is generality: constraints are favored that rule out a larger fraction of the possible strings. Constraint search continues until no further constraints are available that meet the minimal criterion for accuracy, or, where this is convenient, when a user-specified maximum is reached.

As they are selected, constraints are weighted. Intuitively speaking, weights express the relative strength of a constraint; in the completed grammar, higher-weighted constraints play a greater role in lowering the predicted well-formedness of the words that violate them. Weighting follows mathematical procedures, backed by proof, that find the weights that respect the maximum likelihood criterion (i.e., that assign the most probability to the observed data)—hence allocating the smallest probability to unobserved data, insofar as this can be done with the available constraints. The result is phonotactic grammars that are restrictive.[3]

Both the weight-setting process and the predictions of the learned grammar depend on the basic formula for maxent grammars (Berger, Della Pietra, and Della Pietra 1996, Della Pietra, Della Pietra, and Lafferty 1997, Goldwater and Johnson 2003), which assigns a probability to a word ω based on its profile of the constraint violations and on the weights of the violated constraints. The computation is summarized in (2).

(2) *Maxent probability computation (Della Pietra, Della Pietra, and Lafferty 1997:380)*

$$p(\omega) = \exp(-\Sigma_i \lambda_i \chi_i(\omega))/Z, \text{ where } Z = \Sigma_j \exp(-\Sigma_i \lambda_i \chi_i(\omega_j))$$

The elements of the formula are as follows:

$p(\omega)$	predicted probability of word ω
$\exp(x)$	e to the power of x
Σ_i	summation across all constraints
λ_i	weight of the ith constraint
$\chi_i(\omega)$	number of times ω violates the ith constraint
Σ_j	summation across all possible words

---

[2] Assuming a feature may be $+$, $-$, or absent, the number of feature matrices is $3^n - 1$, where n is the number of features. The feature set we use here has 22 features and thus implies about 31 billion matrices.

[3] Following general practice in maxent modeling, we employed a Gaussian prior (Goldwater and Johnson 2003: ex. (3)), $\sigma = 1$. The main effect of this is to avoid infinite weights for never-violated constraints.

A variety of other approaches to constraint weighting are explored and compared by McClelland and Vander Wyk (2006).

The end product is a probability value for ω.[4] If one's goal is simply a comparison of different words, it suffices to calculate just the expression $\Sigma_i \lambda_i \chi_i(\omega)$, which can be regarded as a kind of penalty score.

4 Finding Candidates for Accidentally True Constraints

We began our inquiry by using the Hayes/Wilson learner to search for candidate accidentally true constraints in English similar to (1) for Wargamay. For this purpose, we needed a training set, ideally as close as possible to what is encountered by the experimental participants during language acquisition. Since our participants were Americans, we used the pronunciations of the Carnegie-Mellon Pronouncing Dictionary (http://www.speech.cs.cmu.edu/cgi-bin/cmudict). We selected the words that in the CELEX lexical database (Baayen, Piepenbrock, and Gulikers 1995) have a frequency of at least one; inspection suggested that this would achieve a reasonably good fit with the words known to our participants. We removed from the list as well as we could all compounds, inflected forms, and forms created by highly transparent processes of morphological derivation, since these tend to have special phonotactic properties; the assumption we made is that simple forms presented to participants as possible words are interpreted as monomorphemic and that the relevant phonotactics is that of Level I, as defined in the theory of Lexical Phonology (Kiparsky 1982). This assumption is defended in section 6.1. We also attempted to correct as many errors as possible in the Carnegie-Mellon transcriptions of the words we used, including incorrect markings of primary and secondary stress. Finally, we syllabified our training data following the Maximal Onset Principle (Selkirk 1982), so that constraints could refer to onset and coda position.[5]

Preliminary exploration indicated that the model was not guaranteed to learn constraints that were obviously natural. Since our interest was in examining constraints that had clear typological and/or phonetic support, we fed 36 such constraints into the grammar in advance, then let it continue until it had learned 160 weighted constraints.[6] The resulting grammar gave reasonably good (not perfect) descriptive performance, ruling out most impossible onset clusters, coda clusters, and medial clusters, and we chose it as our base grammar.

---

[4] Probability is used here in a rather abstract sense: a total probability mass of one is allocated among all possible phonological strings, and the probability of a string is its share in the total.

[5] Implementation: We assigned the feature values [+ rhyme] and [− rhyme] to consonants in coda and onset position, respectively. The onsets used for maximal onset syllabification, which follow Hayes and Wilson 2008, are posted at the article website. We avoided ''exotic'' onsets like [km] (*Khmer*), since when maximized they result in implausible syllabifications like *acme* ['æ.kmi]. In the experimental stimuli, we largely avoided questions of syllable division by placing all sequences that violated syllable-based constraints at word edge, where syllabification is unambiguous.

[6] In retrospect, this procedure strikes us as having been too cautious. Following a reviewer's suggestion, we reran the learning simulation without using any prior constraints. The resulting grammar ended up including 9 out of our 10 Unnatural constraints (shown in (6)). Most of our Natural constraints also showed up in recognizable form, either as a notational variant of one of the constraints in (3)–(5) or as a bundle of more complex constraints having a function similar to that of the constraints in (3)–(5). Like our main test grammar, the no-prior-constraints grammar assigned near-identical (and high) penalties to our Natural and Unnatural Violating forms and near-zero penalties to our Natural Control forms. It gave penalties to our Unnatural Control forms about one-eighth the size of that assigned to the Unnatural Violating forms. It thus appears that our results would have been essentially the same had we used the no-prior-constraints grammar. This grammar, and the scores it assigns, may be inspected at the website for this article.

4.1 Selecting the Relevant Constraints

Guided by our knowledge of phonological typology and phonetic naturalness, we picked from
the 160 constraints 10 fairly clearly natural constraints and 10 fairly clearly unnatural ones. Five
of the 10 natural constraints were manifestations of the well-known Sonority Sequencing Principle
(Sievers 1881, Greenberg 1978, Berent et al. 2007). We adopted the feature-based implementation
of the Sonority Hierarchy proposed by Clements (1990) and set up constraints that penalize
consonant clusters that have less than ideal sequencing for a particular sonority feature. These
constraints are listed in (3); for sample violations, see the list of experimental stimuli in (7).

(3) *Natural constraints I: Sonority-based*

    a. *[−son][+son] IN CODA

    b. *[+cons][−cons] IN CODA

    c. *[−cons][+cons] IN ONSET

    d. *[−cont(inuant)][−cont(inuant)] IN ONSET[7]

    e. *[−cont][+nasal] IN ONSET

       where

       [−son] = [p t tʃ k b d dʒ g f θ s ʃ h v ð z ʒ]

       [+son] = [m n ŋ l ɹ w j]

       [+cons] = [p t tʃ k b d dʒ g f θ s ʃ h v ð z ʒ l m n ŋ]

       [−cons] = [ɹ w j]

       [−cont] = [p t tʃ k b d dʒ g]

       [+nasal] = [m n ŋ]

Three constraints reflected the common pattern for coda segments to be homorganic with
what follows and for onset segments to be heterorganic (see Harris 1983:31–35, Kager 1999:
131).

(4) *Natural constraints II: Homorganicity/Heterorganicity-based*

    a. *[+labial][+dorsal] IN CODA   *Heterorganicity in codas*

    b. *[+dorsal][+labial] IN CODA   *Heterorganicity in codas*

    c. *[+labial][+labial] IN ONSET   *Heterorganicity in onsets*

       where [+labial] = [p b f v m w], [+dorsal] = [k g ŋ]

One of the two remaining constraints was straightforward, (5a): the commonplace require-
ment (Lombardi 1999) that obstruent clusters agree in voicing (in this case, only in one particular
direction). The other constraint was the only one that we did not design into the grammar, but
on reflection it seemed a plausible constraint, (5b): it forbids the glides [j, w] in syllable coda.[8]

---

[7] For this constraint, see Morelli 1999; on the basis of a typological survey, Morelli suggests a general constraint
banning obstruent clusters whose first element is a stop.

[8] The Americanist tradition of phonetic transcription (Pullum and Ladusaw 1996:22–24) often depicts the diphthongs
of English with glide letters (e.g., [ay, aw]) to describe what IPA transcription more accurately notates with vowel symbols
([aɪ, aʊ]). Our recorded tokens of forms violating (5b), given in (7j), employ true glides with full constriction.

This is a plausible restriction for a language like English that has multiple diphthongs; the ban keeps the diphthongs distinct from what would be similar vowel + glide sequences. For the principle of phonological dispersion underlying this view, see for example Flemming 2004.

(5) *Natural constraints III: Other*

a. $*\begin{bmatrix}-\text{son}\\-\text{voice}\end{bmatrix}\begin{bmatrix}-\text{son}\\+\text{voice}\end{bmatrix}$ *Voicing assimilation*

b. $*\begin{bmatrix}-\text{syllabic}\\+\text{high}\end{bmatrix}$ IN CODA   *Glides in coda, in a diphthongal language*

where

$$\begin{bmatrix}-\text{son}\\-\text{voice}\end{bmatrix} = [\text{p t tʃ k f θ s ʃ h}]$$

$$\begin{bmatrix}-\text{son}\\+\text{voice}\end{bmatrix} = [\text{b d dʒ g v ð z ʒ}]$$

$$\begin{bmatrix}-\text{syllabic}\\+\text{high}\end{bmatrix} = [\text{j w}]$$

For the unnatural constraints, we combed through the output of the grammar looking for constraints that met several criteria: that they should have, at most, weak typological or phonetic support, that they should have weights similar to those learned for the natural constraints above, and that they should have few or no exceptions in the training data. The constraints are given with prose descriptions in (6).

(6) *10 unnatural constraints*

a. $*\begin{bmatrix}+\text{round}\\+\text{high}\end{bmatrix}\begin{bmatrix}-\text{cons}\\-\text{son}\end{bmatrix}$ *No* [u, ʊ, w] *before* [h]

b. $*\begin{bmatrix}+\text{cons}\\-\text{ant}\end{bmatrix}[-\text{son}]$ [ʃ, ʒ, tʃ, dʒ] *may not precede obstruents*

c. $*[-\text{back}]\ [+\text{diphthong}]$ *No* [j] *before* [aɪ, aʊ, ɔɪ]

d. $*[_{\text{word}}\begin{bmatrix}-\text{diphthong}\\+\text{round}\\+\text{high}\end{bmatrix}$ *No word-initial* [u, ʊ]

e. $*[+\text{diphthong}]\begin{bmatrix}+\text{cont}\\-\text{ant}\end{bmatrix}$ *No* [aɪ, aʊ, ɔɪ] *before* [ʃ, ʒ]

f. $*\begin{bmatrix}+\text{coronal}\\+\text{cont}\\-\text{strident}\end{bmatrix}[-\text{son}]$ *No* [θ, ð] *before obstruents*

g. $*\begin{bmatrix}+\text{cont}\\-\text{strident}\end{bmatrix}\begin{bmatrix}-\text{stress}\\+\text{round}\end{bmatrix}$ *No* [θ, ð] *before stressless rounded vowels*

h. $*\begin{bmatrix} +\text{diphthong} \\ +\text{round} \\ -\text{back} \end{bmatrix} [-\text{ant}]$ *No* [ɔɪ] *before* [ʃ, ʒ, tʃ, dʒ]

i. $*\begin{bmatrix} +\text{cont} \\ +\text{voice} \\ -\text{ant} \end{bmatrix} [+\text{stress }] [-\text{son}]$ *No* [ʒ] *before stressed vowel + obstruent*

j. $*[_{\text{word}} \begin{bmatrix} +\text{diphthong} \\ +\text{round} \end{bmatrix} \begin{bmatrix} -\text{son} \\ +\text{voice} \end{bmatrix}$ *Initial* [aʊ, ɔɪ] *may not precede a voiced obstruent*

For sample violations, see the list of experimental stimuli in (8).

The natural constraints have, in the aggregate, very similar weights to the set of unnatural constraints. The average weight of the natural constraints is 3.79 (range 2.68–4.65), and the average weight of the unnatural constraints is 3.96 (range 3.51–5.22). In terms of testing the model, the minor discrepancy goes in the direction desired: if it turns out that the unnatural constraints have the weaker effect on native-speaker judgment, we do not want to attribute this to the weight difference. Weights for all constraints are given in appendix A.

4.2 Diachronic Origin of Unnatural Constraints

We digress to address the question of why languages should have any unnatural constraints at all. Some of our constraints have a clear diachronic basis in what could be called ''constraint telescoping,'' analogous to the ''rule telescoping'' observed by Kenstowicz and Kisseberth (1977: 64–65). The idea is that an originally natural constraint can be obscured by a sequence of natural historical changes while retaining its effects, simply by inertia, in the inherited lexicon. Constraint (6e), banning [aɪ, aʊ, ɔɪ] before [ʃ, ʒ], is one such case. Ignoring the very rare sounds [ʒ] and [ɔɪ] for simplicity, we observe that /ʃ/ originated in English from historical *sk, and [aɪ] and [aʊ] from historical *iː, *uː. Thus, (6e) is the historical descendent of a constraint that originally banned long vowels before a consonant cluster, a highly natural pattern. This history is discussed in detail by Iverson and Salmons (2005), who suggest that for English a synchronic ban on long vowels before /ʃ/ is (in the terms of this article) accidentally true.[9]

Not all of the constraints of (6) have such a clear diachronic origin, and some may indeed be true entirely by accident. Still others may result from a blend of diachronically motivated and accidental factors. For (6c), the absence of [jaɪ] has a clear diachronic origin, in that [aɪ] descends from [iː], and bans on [j] before high front vowels are common typologically (Kawasaki 1982: sec. 2.7.2; for English, see Jespersen 1909:sec. 58). The lack of [jaʊ], however, may be accidental.

---

[9] In support of this, Iverson and Salmons point out the vulnerability of their constraint to acquiring new counterexamples through borrowing (e.g., *pastiche, cartouche*). The more specific constraint we use here, (6e), has been less vulnerable to counterexamples because the likely donor language, French, lacks [aɪ, aʊ, ɔɪ].

5 Magnitude Estimation Experiment

We used the constraints described in the preceding section to design the nonce words used in the following word acceptability experiment. We used the magnitude estimation technique, following methods described by Lodge (1981) and Bard, Robertson, and Sorace (1996). In this task, participants increase or decrease the magnitude of their response on the basis of the relative increase or decrease in some property of the stimuli. In our case, participants were rating the relative goodness of nonwords as potential words of English. We used number estimation and line drawing as response modalities because these two tasks are easy to implement and their relationship to each other is well-understood.

5.1 Method

5.1.1 Participants   Twenty-nine UCLA undergraduate students participated in the experiment for partial course credit. All participants were native speakers of English with no hearing or speech impairments.

5.1.2 Materials   For each constraint in section 4.1 (both natural and unnatural), we invented two stimulus pairs. Each pair consisted of a Violating word that, in most cases, violated only the constraint in question, and a Control word that, in most cases, violated no other constraints. In a couple of cases, the Control word violated a very low-weighted constraint, which was also violated by the Violating word. Aside from the target violation, we tried to make the words as phonotactically bland as possible, and also to avoid strong resemblances to particular existing words. We found that satisfying all of these requirements at once was not easy, and for this reason the pairs were not statistically controlled for resemblance to existing words.

Since there were 10 Natural and 10 Unnatural constraints, and each constraint was tested with two Violating/Control pairs, there were a total of 80 stimuli: 20 Natural Violating forms, 20 Natural Control forms, 20 Unnatural Violating forms, and 20 Unnatural Control forms. They are listed in (7) and (8). We give the forms both in the orthography we employed and in IPA notation.

(7) *Stimulus pairs for the Natural constraints*

Constraint	Violating - Control	IPA
a. *[−son] [+son] IN CODA	*kipl - kilp*	[ˈkɪpl] - [ˈkɪlp]
	canifl - canift	[kəˈnɪfl] - [kəˈnɪft]
b. *[+cons] [−cons] IN CODA	*tilr - tilse*	[ˈtɪlɹ] - [ˈtɪls]
	shapenr - shapent	[ʃəˈpɛnɹ] - [ʃəˈpɛnt]
c. *[−cons] [+cons] IN ONSET	*hlup - plup*	[ˈhlʌp] - [ˈplʌp]
	hmit - smit	[ˈhmɪt] - [ˈsmɪt]
d. *[−cont] [−cont] IN ONSET	*cping - sping*	[ˈkpɪŋ] - [ˈspɪŋ]
	ctice - stice	[ˈktaɪs] - [ˈstaɪs]
e. *[−cont] [+nasal] IN ONSET	*cnope - clope*	[ˈknoʊp] - [ˈkloʊp]
	pneck - sneck	[ˈpnɛk] - [ˈsnɛk]

f. *[+labial] [+dorsal] IN CODA   *trefk - treft*   ['tɹɛfk] - ['tɹɛft]

                                         rufk - ruft   ['ɹʌfk] - ['ɹʌft]

g. *[+dorsal] [+labial] IN CODA   *bikf - bimf*   ['bɪkf] - ['bɪmf]

                                         sadekp - sadect   [sə'dɛkp] - [sə'dɛkt]

h. *[+labial] [+labial] IN ONSET   *bwell - brell*   ['bwɛl] - ['bɹɛl]

                                         pwickon - twickon   ['pwɪkən] - ['twɪkən]

i. *$\begin{bmatrix} -son \\ -voice \end{bmatrix} \begin{bmatrix} -son \\ +voice \end{bmatrix}$ *esger - ezger*   ['ɛsgɚ] - ['ɛzgɚ]

                                         trocdal - troctal   ['tɹɑkdəl] - ['tɹɑktəl]

j. *$\begin{bmatrix} -syllabic \\ +high \end{bmatrix}$ IN CODA   *jouy - jout*   ['dʒaʊj] - ['dʒaʊt]

                                         tighw - tibe   ['taɪw] - ['taɪb]

(8) *Stimulus pairs for the Unnatural constraints*

       Constraint   *Violating - Control*   *IPA*

a. *$\begin{bmatrix} +round \\ +high \end{bmatrix} \begin{bmatrix} -cons \\ -son \end{bmatrix}$ *luhallem - laihallem*   [lu'hæləm] - [leɪ'hæləm]

                                         tuheim - towheim   [tu'heɪm] - [toʊ'heɪm]

b. *$\begin{bmatrix} +cons \\ -ant \end{bmatrix}$ [−son]   *ishty - ishmy*   ['ɪʃti] - ['ɪʃmi]

                                         metchter - metchner   ['mɛtʃtɚ] - ['mɛtʃnɚ]

c. *[−back] [+diphthong]   *youse - yoss*   ['jaʊs] - ['jɑs]

                                         yout - yut   ['jaʊt] - ['jʌt]

d. *$[_{word} \begin{bmatrix} -diphthong \\ +round \\ +high \end{bmatrix}$ *utrum - otrum*   ['utɹəm] - ['oʊtɹəm]

                                         ooker - ocker   ['ʊkɚ] - ['ɑkɚ]

e. *[+diphthong] $\begin{bmatrix} +cont \\ -ant \end{bmatrix}$ *pyshon - pyson*   ['paɪʃən] - ['paɪsən]

                                         foushert - fousert   ['faʊʃɚt] - ['faʊsɚt]

f. *$\begin{bmatrix} +coronal \\ +cont \\ -strident \end{bmatrix}$ [−son]   *hethker - hethler*   ['hɛθkɚ] - ['hɛθlɚ]

                                         muthpy - muspy   ['mʌθpi] - ['mʌspi]

g. *$\begin{bmatrix} +cont \\ -strident \end{bmatrix} \begin{bmatrix} -stress \\ +round \end{bmatrix}$ *potho - pothy*   ['pɑθo] - ['pɑθi]

                                         taitho - taithy   ['teɪθo] - ['teɪθi]

h. *$\begin{bmatrix} +diphthong \\ +round \\ -back \end{bmatrix}$ [−ant]   *noiron - nyron*   ['nɔɪɹən] - ['naɪɹən]

                                         boitcher - boisser   ['bɔɪtʃɚ] - ['bɔɪsɚ]

i. *$\begin{bmatrix} +cont \\ +voice \\ -ant \end{bmatrix}$ [+stress] [−son]   *zhep - zhem*   ['ʒɛp] - ['ʒɛm]

                                         zhod - zhar   ['ʒɑd] - ['ʒɑɹ]

j. *$[_{word} \begin{bmatrix} +diphthong \\ +round \end{bmatrix} \begin{bmatrix} -son \\ +voice \end{bmatrix}$ *ouzie - oussie*   ['aʊzi] - ['aʊsi]

                                         oid - oit   ['ɔɪd] - ['ɔɪt]

   Our experiment also included a set of filler words, partly as a way of distracting the partici-
pants from the fact that the stimuli were paired, and partly to provide an independent check on

our method. For the fillers, we selected 20 forms each from two earlier phonotactic rating studies: Scholes 1966 (Experiment 5) and Albright 2009. In both cases, the forms selected represented the full range of forms found in each study, from highly well-formed to highly ill-formed. The fillers are listed in appendix B.

Stimuli were presented auditorily as well as orthographically. To create the auditory stimuli, the nonwords were recorded in a random order by an English-speaking, female, trained phonetician in a sound booth, assisted by monitoring and feedback from the experimenters. The speaker read from a transcript containing both orthographic and IPA renderings of the words. From this recording, the tokens judged by the experimenters to be the clearest rendering of each intended phonemic sequence were selected, then equalized for volume.[10]

We presented the stimuli both auditorily and orthographically in order to maximize the chance that participants would internalize the intended phonemic representations of the nonwords represented by the IPA transcriptions in (7) and (8). The auditory presentation provided the intended pronunciation in cases where orthography may be ambiguous. However, studies have shown that nonnative sequences of sounds may be misperceived by listeners (Dupoux et al. 1999); thus, we chose to provide orthography as well in order to aid participants in parsing the intended sequence of phonemes.

5.1.3 Procedure The magnitude estimation procedure consisted of three blocks: a calibration block, a number estimation block, and a line-drawing block. All participants began with an identical calibration phase, following Lodge (1981) and Bard, Robertson, and Sorace (1996). They were told that they would see multiple lines on the computer screen and that they would be assigning each one a number based on the length of the line. They were shown a horizontal line approximately 35 mm in physical length; this was designated as the reference line and assigned a numerical value of 100. Participants were told to enter numerical values for subsequent lines based on their lengths relative to the reference line; if a line was twice as long as the reference line, they were to enter a number twice as high as 100 (200), and so on. Participants entered numbers using the keyboard and pressed the ''next'' button to begin the next trial. The reference line was not displayed while the participants were giving their estimations.

After giving numbers for eight lines ranging from 6 to 600 units, participants were given eight numbers of equivalent values (6 to 600) and asked to draw lines. The number 100 was once again used as a reference value. Participants drew horizontal lines by clicking in a rectangular box on the computer screen, dragging the mouse cursor to another part of the box, then releasing the mouse button. If they clicked in the box again, the old line would disappear and a new line could be drawn. When a participant was satisfied with the line in the box, he or she pressed the

---

[10] Discussing how illegal clusters can be rendered by English speakers, Davidson (2007) provides a useful taxonomy: speakers insert a full schwa, insert a shorter ''transitional schwa,'' or leave the members of the cluster adjacent. We requested our speaker to avoid inserted schwas of any sort, and spectrographic inspection indicates that indeed no inserted schwas of any sort appear in any of the tokens used in the experiment. We also asked our speaker to avoid rendering sonorants in sonority-reversed clusters as syllabic, and careful listening suggests that the speaker likewise succeeded in this task. The sound files for the experimental tokens may be downloaded from the website for this article.

''next'' button to move on to the next word. An experimenter watched the participants perform the calibration phase to make sure that they understood the task. If a participant was not giving a reasonable response (e.g., by entering a number that was less than 100 for a line that was obviously longer than the reference line), then the experimenter would repeat the task instructions until the participant understood. Otherwise, the experimenter gave no further instruction on how to draw lines or give numbers.

After the calibration block had ended, participants were told that they would be performing a similar task but would be rating made-up words. Participants were randomly assigned to perform the number estimations first or to draw lines first.[11] Those who did the number estimation block first were told that they would be entering numbers for made-up words based on how good the words sounded as new words of English. To familiarize the participants with the full range of words they would be looking at, they were given *bzarshk* [ˈbzɑɹʃk] and *kip* [ˈkɪp] as examples of (respectively) strange-sounding and normal-sounding English words. In addition, they were given the word *poik* [ˈpɔɪk] as an example of an intermediate word. All words in the experiment were displayed in English orthography on the screen as well as played through headphones.

The participants were then instructed that *poik* would serve as their reference word and that it should be assigned the number 100. Words that they thought sounded better than *poik* as words of English should be given a number higher than 100 and analogously for words that sounded worse. Participants were encouraged to use a proportional scale: for example, if they thought a word was twice as good a word of English as *poik*, then they would enter a number twice as high as 100 (200); similarly, they would enter 50 for a word that sounded only half as good as *poik*. The rationale for this procedure is that (unlike with ratings scales that use a fixed set of values), participants are free to extend their scale upward or downward when they encounter new items that are unprecedentedly good or bad; it also makes available essentially unlimited granularity for their responses, useful when they encounter new words that seem intermediate between two previous words.

The participants completed four practice words before beginning an experimental phase with the 40 fillers and 80 experimental words described in section 5.1.2. Once the ''next'' button was pressed, the next word appeared on the screen and the sound file was played once automatically. The order of the words was randomized for each participant. The experimenter stayed in the room for the practice trials but left before the participants began the experimental trials.

After completing the number estimation block, the participants were instructed to perform the same task except with line drawing instead of numbering. *Poik* was again used as the reference word, presented with a line of 100 units. If a word seemed twice as good as *poik*, the participants were instructed to draw a line twice as long, and so on. Participants drew lines for the same set of practice items and experimental items, in a newly randomized order. This block completed the experiment. Participants who were assigned to perform the line-drawing block first completed the same tasks with the same stimuli, but with the blocks in reverse order.

---

[11] Because of a programming error, more participants were given the number estimation task first than were given the line-drawing task first. However, post hoc tests did not find any significant differences between the two groups.

5.2 Results and Discussion

5.2.1 Calibration   Studies using magnitude estimation can be calibrated to assess their validity. We first examine if participants are self-consistent in the training phase described in the previous section: do the lines they draw match up to the numbers they are attempting to match, and vice versa? We can check this by performing a regression analysis, comparing a participant's numerical response to a line of a particular length against the same participant's line length for the same number. This analysis (carried out with log values) yields a strong positive correlation ($r = .96$). The slope of the regression line is almost exactly one. This showed that, as in previous work, our participants had no trouble performing the basic magnitude estimation task. In addition, as a group, they neither underestimated nor overestimated in either modality.

We also examined how participants' responses to the nonword items compared across the two modalities. Regression analysis for these values indicated nearly perfect correlation ($r = .98$) and a perfect slope of one. This indicates that participants were consistent in their nonword ratings across the two modalities. Therefore, we may assume that these values are valid and reliable (for discussion, see Lodge 1981 and Bard, Robertson, and Sorace 1996).

5.2.2 Replication of Scholes 1966 and Albright 2009   We found that the mean log ratings of the borrowed fillers correlated strongly with log ratings from Scholes 1966 and Albright 2009 ($r = .90$ and $r = .86$, respectively), indicating that our experiment succeeded in eliciting similar phonotactic well-formedness intuitions.

5.2.3 Main Results   For the following analyses, data from the line-drawing task and the number estimation task, which yielded very similar results, have been collapsed. As a check, we ran all of the analyses on the line data and numerical data separately, and the results showed the same basic pattern as with the combined data.

Figure 1 shows the mean log ratings for nonwords according to the Naturalness of the constraint being tested (Natural or Unnatural) and to the nonwords' status as Control or Violating
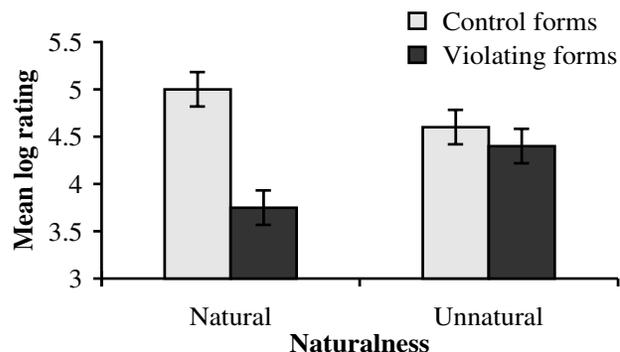

Figure 1
Mean log ratings for combined line drawing and number estimation data by Naturalness and Control/Violating Status

Table 1

Results of mixed-effects model for Naturalness and Control/Violating Status. (CI = confidence interval)

Fixed effects	Estimate	95% CI		t-value	p-value
Intercept	5.00	4.79	5.21	42.85	<.001
Status = Violating form	−1.33	−1.58	−1.09	−9.49	<.001
Naturalness = Unnatural	−0.40	−0.65	−0.17	−2.87	.004
Naturalness = Unnatural & Status = Violating	1.13	0.80	1.48	5.70	<.001

Random effects	Standard deviation
Subject (intercept)	0.33
Item (intercept)	0.43
Residual	0.76

forms (error bars represent the standard error of the mean). The figure shows that for the Natural constraints, the ratings for Violating forms ($M = 3.67$, $SD = 1.02$) were much lower than those for Control forms ($M = 5.00$, $SD = 0.87$). For the Unnatural constraints, the ratings for Violating forms ($M = 4.40$, $SD = 0.89$) were also lower than those for Control forms ($M = 4.60$, $SD = 0.92$), but this difference was much smaller—less than one-sixth of the difference found for the Natural constraints.

To evaluate these differences, we created linear mixed-effects models in R (R Development Core Team 2008) using the *lmer*() function of the *lme4* package (Bates, Maechler, and Dai 2008), following Baayen (2008a:chap. 7). As a baseline model, we began with the factors that we were interested in—Naturalness and Control/Violating Status—as fixed effects with an interaction term. Random intercepts were included for Subject and Item because they significantly improved model fit.[12] The results of this model are presented in table 1. P-values and 95% confidence intervals (CI) were computed by a Monte Carlo Markov chain (MCMC) sampling method, using the *pvals.fnc*() function of the *languageR* package (Baayen 2008b) with 10,000 samples.

Each factor contributed significantly to the model. A potential form begins with the baseline intercept log score of 5.00. (This is in fact the Natural Control mean rating, since Natural Control forms are not further modified by other factors in the model.) The row below *Intercept* indicates that Violating forms had significantly lower ratings in general than Control forms, by about 1.33. In the next row, forms selected for an Unnatural constraint (Control or Violating) received a significantly lower rating than those selected for a Natural constraint, by 0.40. Finally, the last

---

[12] Following a reviewer's suggestion, we also tried models containing random slopes for Subject according to Naturalness and the interaction between Naturalness and Status. These random slopes also significantly improved model fit; however, estimating *p*-values using MCMC sampling for models with random slopes is not currently implemented in the *languageR* package. The *t*-values of the model with random slopes were very similar to those in table 1: Intercept 42.03, Status = Violating −7.30, Naturalness = Unnatural −2.77, and Naturalness/Status interaction 5.03. Moreover, with a large number of degrees of freedom, it can be estimated that a *t*-value greater than 2 represents a significant value (Baayen 2008a:270). We conclude that adding random slopes to the model does not change the overall pattern or magnitude of significance presented in table 1.

row of fixed effects shows that being a Violating form of an Unnatural constraint resulted in a significantly higher rating as compared with forms violating Natural constraints, by 1.13. This final factor is the crucial interaction term: it indicates that violating a Natural constraint is much worse than violating an Unnatural constraint.

To confirm that adding the fixed effects for Naturalness and its associated interaction term improves model fit, the model in table 1 was compared with an analogous model containing the random effects and a fixed effect for only Control/Violating Status using a log likelihood test, performed with the *anova*() function in R (see Baayen 2008a). The fixed effects for Naturalness and Naturalness × Control/Violating Status significantly improved model fit (negative log likelihoods: *model with naturalness* = −5430 vs. *model without naturalness* = −5446),[13] $\chi^2(2) =$ 30.18, $p < .001$. In other words, a model appealing to naturalness fits the data better than a model that treats all of the constraints alike.

5.2.4 Individual Constraints We next examine the results on a constraint-by-constraint basis. We estimate the magnitude of the effect of individual constraints by taking the ratio *log rating of Control form/log rating of Violating form* and average over both data types (line, number) and both word pairs (from (7) and (8)) used to test the constraints. By this measure, with just one exception every Natural constraint had a stronger effect on ratings than every Unnatural constraint. This is shown in table 2.

The exception was *[−son,−voice][−son,+voice], exemplified by *esger* versus *ezger* and *trocdal* versus *troctal*. This exception is easily explained: although our speaker was in general able to produce our forms with high accuracy, measurement indicates that she did not succeed in producing a voiced closure in either *esger* or *trocdal*; hence, what we had intended as phonotactically illegal forms were very close to ordinary (['ɛskɚ] and ['tɹɑktəl]).

5.2.5 Did the Unnatural Constraints Have Any Effect? The mixed-effects model establishes that the Unnatural constraints did not have as strong an effect as the Natural constraints, but the question remains whether the Unnatural constraints had any effect at all. To investigate this question, we created another linear mixed-effects model on a subset of the data containing only the Unnatural constraint forms, with Control/Violating Status as a fixed effect and random intercepts for Subject and Item. The model (using *pvals.fnc*() as above) found that the small difference between Violating and Control forms, though trending in the right direction, did not reach significance, *Estimate* = −0.20, *t-value* = −1.54, $p = .12$. A second version of this experiment using only orthographic forms (not reported here) also found that the Control forms were rated only slightly better, but the difference reached significance in that version. We conclude that the Unnatural constraints had, at best, only a small effect on participant ratings.

---

[13] A negative log likelihood closer to 0 indicates a better fit. The fixed effects for Naturalness result in a similar increase in log likelihood when random slopes are included in the model.

Table 2
Effects of individual constraints

Constraint	Status	Pairs	Effect size
*[−cont] [−cont] IN ONSET	natural	*cping/spring, ctice/stice*	1.65
*[glide] IN CODA	natural	*jouy/jout, tighw/tibe*	1.56
*[−cons] [+cons] IN ONSET	natural	*hlup/plup, hmit/smit*	1.51
*[−cont] [+nasal] IN ONSET	natural	*cnope/clope, pneck/sneck*	1.44
*[+labial] [+dorsal] IN CODA	natural	*rufk/ruft, trefk/treft*	1.44
*[+dorsal] [+labial] IN CODA	natural	*bikf/bimf, sadekp/sadect*	1.36
*[+cons] [−cons] IN CODA	natural	*shapenr/shapent, tilr/tilse*	1.34
*[−son] [+son] IN CODA	natural	*canifl/canift, kipl/kilp*	1.31
*[+labial] [+labial] IN ONSET	natural	*bwell/brell, pwickon/twickon*	1.23
$*\begin{bmatrix} +\text{coronal} \\ +\text{cont} \\ -\text{strident} \end{bmatrix} [-\text{son}]$	unnatural	*hethker/hethler, muthpy/muspy*	1.14
$*\begin{bmatrix} +\text{cont} \\ -\text{strident} \end{bmatrix} \begin{bmatrix} -\text{stress} \\ +\text{round} \end{bmatrix}$	unnatural	*potho/pothy, taitho/taithy*	1.10
$*\begin{bmatrix} +\text{diphthong} \\ +\text{round} \\ -\text{back} \end{bmatrix} [-\text{ant}]$	unnatural	*boitcher/boisser, noiron/nyron*	1.10
$*[+\text{diphthong}] \begin{bmatrix} +\text{cont} \\ -\text{ant} \end{bmatrix}$	unnatural	*foushert/fousert, pyshon/pyson*	1.08
$*[_\text{word} \begin{bmatrix} -\text{diphthong} \\ +\text{round} \\ +\text{high} \end{bmatrix}$	unnatural	*ooker/ocker, utrum/otrum*	1.03
*[−back] [+diphthong]	unnatural	*youse/yoss, yout/yut*	1.02
$*[_\text{word} \begin{bmatrix} +\text{diphthong} \\ +\text{round} \end{bmatrix} \begin{bmatrix} -\text{son} \\ +\text{voice} \end{bmatrix}$	unnatural	*oid/oit, ouzie/oussie*	1.02
$*\begin{bmatrix} +\text{cont} \\ +\text{voice} \\ -\text{ant} \end{bmatrix} [+\text{stress}] [-\text{son}]$	unnatural	*zhep/zhem, zhod/zhar*	1.01
$*\begin{bmatrix} +\text{cons} \\ -\text{ant} \end{bmatrix} [-\text{son}]$	unnatural	*ishty/ishmy, metchter/metchner*	0.99
$*\begin{bmatrix} -\text{son} \\ -\text{voice} \end{bmatrix} \begin{bmatrix} -\text{son} \\ +\text{voice} \end{bmatrix}$	natural	*esger/ezger, trocdal/troctal*	0.98
$*\begin{bmatrix} +\text{round} \\ +\text{high} \end{bmatrix} \begin{bmatrix} -\text{cons} \\ -\text{son} \end{bmatrix}$	unnatural	*luhallem/laihallem, tuhaim/towhaim*	0.97

5.2.6 Considering Additional Factors   To check whether factors other than those discussed above played a role in our experiment, we also considered a number of additional variables post hoc. These included the following: (a) for each form, the weight assigned by the phonotactic learning model to the constraint it violates;[14] (b) the score assigned by Albright's (2009) phonotactic learner;[15] (c) two measures of length: number of syllables and number of segments; and (d) two measures of assessing the simplicity or generality of constraints: number of features in the constraint (roughly following Chomsky and Halle 1968:334), and the proportion of all logically possible *n*-grams that violate each *n*-gram constraint (Hayes and Wilson 2008:394).[16]

These additional factors were examined by adding them to the model in table 1 one at a time, both with and without interaction terms. The resulting models were compared with the original model using log likelihood tests to determine whether adding the additional term(s) would significantly improve model fit.[17]

Only one of these factors resulted in a significant increase to model fit: number of features in the constraint. However, the effect was in the wrong direction: violations of constraints with more features had a stronger effect on nonword ratings than violations of constraints with fewer features. This goes against traditional views of generality, in which constraints with fewer features are simpler, and simpler constraints are more highly valued. Thus, we judge that this effect was most likely an accident. Our factors of interest (i.e., those in table 1) remain highly significant in the model even when number of features is included, meaning that this additional factor does not confound the main findings of this study.

It is important to keep in mind the following when considering these additional factors: the current study was designed to compare the Unnatural constraints with the Natural constraints, so we attempted to control other factors as much as possible. As a result, the nonwords varied minimally with respect to these other factors (e.g., constraints were chosen such that their weights were similar, and nonwords varied little in their length). Therefore, any effects of these factors (or the lack thereof) are not very meaningful for this experimental design, provided that they do not confound the results of interest. A study intended to test for these additional factors would vary them systematically rather than controlling for them.

6 Possible Objections

In this section, we consider various alternative interpretations of our results.

---

[14] Using the constraint weights in the model is mostly redundant, since they closely match our Violating versus Control factor. Our interest was whether the small differences in weights among the test stimuli would have a significant effect beyond what would result from the binary factor.

[15] We would like to thank Adam Albright for assistance in computing these scores.

[16] Complexity is worth examining because there is a considerable body of evidence indicating a bias for simple generalizations in phonological learning; see Moreton and Pater, in preparation. For purposes of computing complexity, we included the ad hoc features we used for word boundaries and syllabification.

[17] Constraint weights and Albright (2009) scores were centered before we ran the models by subtracting the mean from each value to reduce collinearity (see Baayen 2008a:276–277).

6.1 The Effect of Training Data

Our training data (section 4) were chosen because we felt that they offered the best chance of matching the mental lexicons of our experimental participants. However, it is possible that our training set was inadequate in two ways. First, we excluded many morphologically complex forms, under the assumption that these forms have their own phonotactics and would not affect the judgments of new monomorphemic forms. This assumption may not have been well-founded; that is, perhaps complex forms did affect the well-formedness judgments of our participants. Some of our constraints are indeed potentially affected; thus, although (6f) *[+cont,−strident][−son] is violation-free in simplex forms, it is violated in past tenses such as *bathed* [beɪðd].

Second, it is possible that we underestimated the number of words in the mental lexicons of our participants that are very rare and violate the Unnatural constraints: if these are included in the training data, the weights of the Unnatural constraints would go down, perhaps explaining the experimental findings. For instance, it seems plausible that many of our participants do not know the name *Pushkin* (it is not in our training set); but if they do, and consider it to be an English word, it would produce a slightly lower weight for constraint (6b), *[+cons,−ant][−son].

To test these possibilities, we created three new training sets. The first contained all of the affixed forms that we had previously excluded. For the second, we tried to include all exceptions to our unnatural constraints (such as *Pushkin*) that the participants might plausibly have been familiar with; most of these we found by consulting the full Carnegie-Mellon database. The total number added was 24. The third training set combined the first two, including both the affixed forms and the rare exceptions. We then reweighted the 160-constraint grammar using each of the new training sets. The mean weights for the Natural and Unnatural constraints using each training set are shown in table 3.

As the means in table 3 demonstrate, constraint weights did vary to some extent depending on the training set. The means in boldface mark the largest changes. As expected, the weights for the Unnatural constraints fell slightly when relevant exceptions were added to the training set. The Natural constraints, on the other hand, received higher weights when the affixed forms were added, probably as a result of the larger set of data for which they could ''prove their worth'' by remaining exceptionless.

However, the changes in mean constraint weight were relatively small for both the Unnatural constraints (3.96 → 3.58) and the Natural constraints (3.79 → 4.07). Even though the Unnatural

Table 3
Mean constraint weights for Unnatural and Natural constraints using the original training set and each of the three modified training sets

	Original	With affixed forms	With exceptions	With exceptions and affixed forms
Unnatural constraints	3.96	3.84	**3.65**	**3.58**
Natural constraints	3.79	**4.07**	3.79	**4.07**

constraint weights become smaller than the Natural constraint weights, they are still quite similar. It is unlikely that this small difference could explain the large difference in effect between Natural and Unnatural constraints found in our experiments. In our testing using these weights, constraint weight continued to have no statistically significant effect.

6.2 Have We Correctly Classified Our Constraints for Naturalness?

It is not easy to establish firmly the naturalness of constraints by either of the criteria laid out in section 1. The reviewers for this article were helpful in offering their input concerning some of our naturalness claims. No one seems to have objected to the classification of our Natural constraints as natural, but some of our Unnatural constraints may have been prematurely classified as such. Thus, (6f) excludes the nonstrident fricatives [θ, ð] before obstruents; the Latin fricative [f], phonetically similar to [θ], was likewise excluded before obstruents.[18] Constraint (6b) forbids palato-alveolars before obstruents; this was a productive phonological constraint in Sanskrit (see Whitney 1889:72–75). Constraint (6c), forbidding [j] before diphthongs, might be assigned a phonetic rationale: it bans a *high - nonhigh - high* pattern within the syllable, perhaps analogous to the widespread ban on complex (triply linked) contour tones (Yip 2002:30).

In light of this, we ran an additional model in which (6b), (6c), and (6f) were recategorized as Natural constraints. The effects in the new model were somewhat smaller than in the original model, but the overall pattern remained the same and the crucial interaction term remained highly significant ($p < .001$). Indeed, our main result is fairly robust against further such reclassifications. We experimented with reclassifying as Natural not just (6b), (6c), and (6f), but also (6a), (6i), and (6j)—the three Unnatural constraints that had the smallest effect size (table 2) and thus contributed most to our statistical result. Even with just four Unnatural constraints still classified as such, the main result remained statistically significant ($p = .046$).

In the long term, finding better ways of assessing phonological naturalness (e.g., through typological surveys and modeling) is needed to allow us to pin down the concept of naturalness more precisely.

6.3 How Do Experimental Participants Interpret Ill-Formed Stimuli?

As noted earlier, our experiment faced the problem of how to present to experimental participants phonological sequences that are phonologically ill-formed, given that people sometimes hear phonologically illegal forms as perceptually similar legal forms. The question arises primarily with our Natural Violating forms, which, it seems clear, were by far the hardest to hear accurately. We must consider the following scenario: the participants may have perceptually repaired a stimulus (e.g., hearing our *hlup* as [flʌp]), but at the same time noticed that the stimulus was a phonetically poor rendition of the perceived phonemic intent. The very low scores assigned to

---

[18] We confirmed this by searching a Latin electronic corpus; [f] occurs only before vowels and the liquids [l, r].

our Natural Violating stimuli might reflect this phonetic factor, rather than the phonological ill-formedness of the phonemic sequences we had intended.[19]

   We carried out an informal post hoc test of this hypothesis by asking seven English-speaking undergraduate students who had had one term of phonetic training to transcribe our Natural Violating stimuli in IPA notation. Unlike our experimental participants, they listened without the aid of an orthographic form. We found that a number of the stimuli were indeed systematically misheard; the worst-case example was *jouy* ['dʒaʊj], heard by all seven consultants as disyllabic ['dʒaʊ.i]. For purposes of assessing our main result, we confined our attention to the opposite end of the spectrum: the three forms heard accurately by all seven consultants (*pneck*, *bwell*, and *rufk*) and the three forms heard accurately by six out of seven (*cping*, *sadekp*, and *trefk*).

   Redoing the statistical analysis with just these six Natural items and their Control forms, we found that the new model was very similar to the one in table 1. Most importantly, the interaction effect remained significant: violating an Unnatural constraint resulted in a smaller reduction in participant ratings than violating a Natural constraint. In fact, the model's estimate of the interaction effect (i.e., *how much* worse it is to violate a Natural constraint than to violate an Unnatural constraint) actually increased slightly from 1.13 in the original model to 1.38 in the present model. Moreover, the effect cannot be attributed to higher weights for the Natural constraints penalizing the accurately heard forms, because the average weight of these constraints was in fact lower than the average for our Natural constraints overall. We conclude that although misperception of stimuli may have occurred in our experiments, it is unlikely to provide an adequate alternative explanation of our results.

6.4 Could the Unnatural Constraints Have Been Excluded on Statistical Grounds?

It is possible that the magnitude of a constraint's weight is not a fully accurate reflection of the constraint's importance in accounting for the data. This hypothesis can be checked by carrying out a statistical assessment of a constraint's effect in improving the performance of a grammar. The rationale for doing this is the possibility that language learners might likewise be unconsciously savvy about the effectiveness of constraints, and evaluate them with a procedure analogous to statistical testing.

   Pursuing this possibility, we used the likelihood ratio test, which is commonly used to assess models that are in a subset relation.[20] For purposes of testing a constraint, we designate as the *subset model* a grammar (with optimized weights) formed with all of the constraints except the tested one, and as the *full model* the grammar (with optimized weights) that uses all the constraints. The likelihood ratio test computes the value $-2 \times \log(\text{probability}_{\text{D,full model}}) / (\text{probability}_{\text{D,subset model}})$, where $\text{probability}_{\text{D}}$ is the probability that the model assigns to the training data. The distribution of this value can be approximated by a chi-square distribution with one degree of

---

[19] That detailed phonetic properties of experimental stimuli can strongly affect phonotactic ratings is demonstrated in Wilson and Davidson 2009.
   [20] For a clear description of the test, see Pinheiro and Bates 2000:83.

freedom, from which one can determine the probability of the hypothesis that the improvement in accuracy due to including the target constraint could arise by accident.

Using software provided to us by Colin Wilson, we achieved an approximation of this statistic for both our Natural and our Unnatural constraints.[21] All constraints tested as highly significant by this test; no p-value was greater than .007, and many were much smaller. We conclude that the Unnatural constraints are as well justified by the lexical data as the Natural ones.

7 General Discussion

To review, the original impetus for our study was a point made by Hayes and Wilson (2008:sec. 8.5) concerning their Wargamay simulation: that in the course of learning the system, their model generated a large set of constraints that are evidently phonologically unnatural. Hayes and Wilson suggested that either (a) language learners are actually very adept at learning such generalizations, so that these constraints would turn out to be valid if tested against Wargamay native intuition, or (b) the constraints reveal a defect in the model. Our findings point to the latter conclusion. Colavin, Levy, and Rose (2010), applying the Hayes/Wilson model to a corpus of Amharic roots, obtained a similar result, finding that the model was able to provide only limited improvement over a core model of hand-created constraints. We conjecture that in this case the model was likewise finding unnatural constraints, which are undervalued by Amharic learners.

Our findings do not suffice to identify with certainty where the model is going astray. We consider three possibilities here.

7.1 Naturalness

Our original hypothesis was that natural constraints are learned more easily than unnatural constraints. As we noted earlier, this hypothesis takes two flavors, one of which is that unnatural constraints are simply inaccessible to language learners. We take the extensive evidence reviewed in section 2.1 as indicating that this possibility is unlikely; and indeed it is possible that in our own experiment, the Unnatural constraints did have a modest effect on ratings (see section 5.2.5).

A more plausible theory is that learners are *biased* to favor natural generalizations, a view suggested by Wilson (2006), Albright (2007), Berent et al. (2007), Finley (2008), Kawahara (2008), Moreton (2008), Finley and Badecker (2009), Hayes et al. (2009), and others. A simple way to check for bias is to examine the output of maxent grammars in which the weights of the Unnatural constraints have been ''hobbled'' (i.e., given a lower weight than would be justified simply by fit to the data). We experimented with this by modifying the grammar described in section 4, multiplying the weights of the Unnatural constraints by a factor that varied from 0 to 1 and examining the correlation of the resulting scores with the log average participant ratings for each experimental stimulus. The best-fit value of this ''hobbling'' factor was .33—a substantial weakening, but not elimination, of the effect of the Unnatural constraints.

---

[21] The approximation was that we could only use the 98 first-learned constraints as our base grammar; memory limitations prevented our testing with the full 160-constraint grammar.

Language learners must have access to some basis for a learning bias. We think a plausible basis would be phonetically based phonology (see, e.g., Myers 1997, Boersma 1998, Hayes 1999, Steriade 1999, 2001, Côté 2000, Flemming 2001, Hayes, Kirchner, and Steriade 2004). Under this approach, language learners evaluating phonotactic generalizations would evaluate them not just for their degree of fit to the learning data, but also for their effectiveness in avoiding articulatory difficulty and in maintaining perceptual distance between contrasting forms in perception.

7.2 Naturalness Again: Are Consonant-Vowel Generalizations Harder to Learn?

Moreton (2008) has provided experimental evidence suggesting that phonotactic generalizations that require access to both vowel and consonant identity can in some cases be phonetically natural but nevertheless harder to learn: a bias exists, but it is a general learning bias rather than one based on phonetic naturalness. Becker, Ketrez, and Nevins (2011) suggest that the same principle could also explain their data. As a reviewer points out, this explanation might be applicable here. All but one of our Natural constraints (see (3)–(5)) evaluate consonant sequences, and all but one of our Unnatural constraints (see (6)) evaluate consonant-vowel sequences. We have no data that could distinguish whether it is a general learning bias or a phonetic one that favors the Natural constraints in cases where the two principles are both applicable. We add that the world's phonologies do include a great many cases of vowel-consonant interaction, as in palatalization, nasalization of nasal-adjacent vowels, influence of secondary consonant articulation on vowel quality, and intervocalic lenition, so we think that further research would be needed to establish a learning bias for consonant-vowel patterns more firmly.

7.3 The Search Heuristics

Another possibility is that the Hayes/Wilson learner might learn fewer (or no) unnatural-seeming constraints if it were modified to use different search heuristics, so that ''accidentally true'' constraints became less tempting to it. The existing heuristics, reviewed in section 3, favor constraints that are exceptionless or nearly so. Yet exceptionlessness is not necessarily as helpful a criterion as it might seem. For instance, in seeking exceptionlessness the model favored constraint (6i), *ʒ [+stress][−son], which reduced the impetus to learn a more general ban on pretonic [ʒ] (*ʒ [+stress]); indeed, there was no such ban in the 160-constraint grammar we used.[22] If the model were altered to give more priority to generality and less to exceptionlessness, then it might have acquired *ʒ [+stress] first, which would then have devalued (6i) (given it a lower weight). It might even have prevented the selection of (6i) entirely, since the model favors only constraints whose violation counts are below the *expected* value; and learning *ʒ [+stress] would lower the expected value for (6i).

The plausibility of this scenario is increased by the fact that our Unnatural Control forms generally received lower ratings than our Natural Control forms (see figure 1). A possible reason

---

[22] The reason the model included [−son] is that none of the six words in our training set that had pretonic [ʒ] (e.g., *luxuriant, regime, genre*) included an obstruent after the stressed vowel; thus, adding [−son] reduces the number of exceptions from six to zero.

for this is that these forms violate simple constraints unlearned by the model, such as *ʒ [+stress]; these are the constraints that might have preempted the learning of our Unnatural constraints had they been learned first.[23]

7.4 For Future Work

In conclusion, we suggest that the modeling research strategy pursued by Hayes and Wilson could be informative concerning the effectiveness of a naturalness bias approach. As Wilson (2006) showed, a bias can be formalized in maxent through the use of constraint-specific prior terms, which militate against the assignment of high weights to particular constraints, such as (as Wilson suggests) the phonetically natural ones. In this approach, the hobbling of unnatural constraints (as in section 7.1) would take place as part of the learning system itself, rather than being a post hoc procedure.

Thus, in principle, all the ingredients for exploring the role of natural versus unnatural constraints in phonotactic learning are at hand. Candidate theories of Universal Grammar must be formalizable as constraint sets (or systems that construct constraints) and must come with a mechanism for imposing biases, perhaps based on more than one mechanism (here we have considered phonetic naturalness, simplicity (section 5.2.6), and simultaneous reference to consonants and vowels (section 7.2)). Such systems could be tested by the method we have used here, comparing model predictions with native-speaker ratings of carefully chosen stimulus words. Our hope is that such a program would facilitate progress on the question of naturalness in phonology by making it possible to test specific hypotheses and mechanisms.

---

[23] A modified learner created by Wilson that appears promising on these lines is reported by Wilson and Davidson (2009), Berent et al. (2012), and Hayes, Wilson, and Shisko (to appear). This learner uses the principle of ''gain'' (Della Pietra, Della Pietra, and Lafferty 1997:381) to select constraints. In a preliminary examination of this modified system, we found that the constraints it selects were indeed more general and less idiosyncratic than those chosen by the Hayes/ Wilson (2008) learner; more specifically, it learned none of our 10 Unnatural constraints. However, it may still be learning *different* ones: for instance, it posited a constraint banning stressed lax vowels before coronal codas as well as a constraint against word-final sonorants. Our testing was tentative, owing to memory limitations, and more serious evaluation of the revised model awaits further research.

Appendix A

The following table lists mean log experimental ratings by nonword, with constraint and constraint weight. Phonetic transcriptions of the stimuli are given in (7) and (8).

Natural constraints					
Violators		**Controls**		**Constraint**	**Weight**
tilr	3.51	*tilse*	4.77	*[+cons] [−cons] IN CODA	3.01
shapenr	3.58	*shapent*	4.71		
trefk	3.80	*treft*	5.23	*[+labial] [+dorsal] IN CODA	2.68
rufk	3.64	*ruft*	5.47		
bikf	3.23	*bimf*	3.76	*[+dorsal] [+labial] IN CODA	3.03
sadekp	3.13	*sadect*	4.88		
esger	4.50	*ezger*	4.20	*$\left[\begin{array}{c} -\text{son} \\ -\text{voice} \end{array}\right]\left[\begin{array}{c} -\text{son} \\ +\text{voice} \end{array}\right]$	3.14
trocdal	5.12	*troctal*	5.24		
bwell	4.23	*brell*	5.27	*[+labial] [+labial] IN ONSET	4.07
pwickon	4.09	*twickon*	4.98		
cnope	3.49	*clope*	5.38	*[−cont] [+nasal] IN ONSET	4.21
pneck	3.54	*sneck*	4.76		
hlup	3.56	*plup*	5.06	*[−cons] [+cons] IN ONSET	4.27
hmit	3.41	*smit*	5.49		
cping	3.23	*sping*	5.40	*[−cont] [−cont] IN ONSET	4.33
ctice	3.29	*stice*	5.33		
kipl	3.77	*kilp*	5.12	*[−son] [+son] IN CODA	4.54
canifl	3.58	*canift*	4.51		
jouy	3.69	*jout*	5.20	*$\left[\begin{array}{c} -\text{syllabic} \\ +\text{high} \end{array}\right]$ IN CODA	4.65
tighw	3.04	*tibe*	5.29		
Mean	**3.67**		**5.00**		**3.79**

Unnatural constraints					
Violators		**Controls**		**Constraint**	**Weight**
ouzie	4.23	*oussie*	4.55	*[$_{\text{word}}$ $\left[\begin{array}{c} +\text{diphthong} \\ +\text{round} \end{array}\right]\left[\begin{array}{c} -\text{son} \\ +\text{voice} \end{array}\right]$	3.51
oid	4.28	*oit*	4.09		
pyshon	4.71	*pyson*	5.30	*[+diphthong] $\left[\begin{array}{c} +\text{cont} \\ -\text{ant} \end{array}\right]$	3.58
foushert	4.37	*fousert*	4.49		
potho	4.68	*pothy*	5.05	*$\left[\begin{array}{c} +\text{cont} \\ -\text{strident} \end{array}\right]\left[\begin{array}{c} -\text{stress} \\ +\text{round} \end{array}\right]$	3.65
taitho	4.34	*taithy*	4.88		
zhep	3.75	*zhem*	3.76	*$\left[\begin{array}{c} +\text{cont} \\ +\text{voice} \\ -\text{ant} \end{array}\right]$ [+stress] [−son]	3.71
zhod	3.75	*zhar*	3.84		

luhallem	4.21	*laihallem*	4.01	$*\begin{bmatrix} +\text{round} \\ +\text{high} \end{bmatrix}\begin{bmatrix} -\text{cons} \\ -\text{son} \end{bmatrix}$	3.78
tuheim	4.27	*towheim*	4.18		
noiron	4.33	*nyron*	5.26	$*\begin{bmatrix} +\text{diphthong} \\ +\text{round} \\ -\text{back} \end{bmatrix}[-\text{ant}]$	4.01
boitcher	5.02	*boisser*	4.98		
youse	4.57	*yoss*	4.74	$*[-\text{back}]\ [+\text{diphthong}]$	4.02
yout	4.70	*yut*	4.69		
hethker	4.45	*hethler*	4.94	$*\begin{bmatrix} +\text{coronal} \\ +\text{cont} \\ -\text{strident} \end{bmatrix}[-\text{son}]$	4.04
muthpy	4.35	*muspy*	5.07		
ishty	3.92	*ishmy*	3.94	$*\begin{bmatrix} +\text{cons} \\ -\text{ant} \end{bmatrix}[-\text{son}]$	4.11
metchter	4.71	*metchner*	4.64		
utrum	4.65	*otrum*	4.97	$*[_{\text{word}}\begin{bmatrix} -\text{diphthong} \\ +\text{round} \\ +\text{high} \end{bmatrix}$	5.22
ooker	4.68	*ocker*	4.61		
Mean	**4.40**		**4.60**		**3.96**

Appendix B

Shown here are the filler items used in the experiment.

From Scholes 1966:

stin ['stɪn], *smat* ['smæt], *blung* ['blʌŋ], *frun* ['fɹʌn], *glung* ['glʌŋ], *shlurk* ['ʃlɚk], *skeep* ['skip], *vrun* ['vɹʌn], *srun* ['sɹʌn], *vlurk* ['vlɚk], *shtin* ['ʃtɪn], *shnet* ['ʃnɛt], *zrun* ['zɹʌn], *shmat* ['ʃmæt], *zlurk* ['zlɚk], *znet* ['znɛt], *fnet* ['fnɛt], *zhmat* ['ʒmæt], *vnet* ['vnɛt], *vkeep* ['vkip]

From Albright 2009:

wiss ['wɪs], *stip* ['stɪp], *trisk* ['tɹɪsk], *preek* ['pɹik], *nace* ['neɪs], *spling* ['splɪŋ], *bize* ['baɪz], *gude* ['gud], *drit* ['dɹɪt], *skick* ['skɪk], *kweed* ['kwid], *blig* ['blɪg], *gwenge* ['gwɛndʒ], *twoo* ['twu], *sfoond* ['sfund], *smeerg* ['smiɹg], *trilb* ['tɹɪlb], *ploamf* ['ploʊmf], *smeenth* ['sminθ], *pwudge* ['pwʌdʒ]

References

Albright, Adam. 2002. Islands of reliability for regular morphology: Evidence from Italian. *Language* 78: 684–709.

Albright, Adam. 2007. Natural classes are not enough: Biased generalization in novel onset clusters. Paper presented at the 15th Manchester Phonology Meeting, Manchester, UK, 24–26 May.

Albright, Adam. 2009. Feature-based generalization as a source of gradient acceptability. *Phonology* 26: 9–41.

Albright, Adam, and Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90:119–161.

Baayen, R. Harald. 2008a. *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.

Baayen, R. Harald. 2008b. languageR: Data sets and functions with ''Analyzing Linguistic Data: A practical introduction to statistics.'' R package version 0.953.

Baayen, R. H[arald], R. Piepenbrock, and L. Gulikers. 1995. The CELEX lexical database (Release 2) [CD-ROM]. Philadelphia: University of Pennsylvania, Linguistic Data Consortium [Distributor].

Bard, Ellen Gurman, Dan Robertson, and Antonella Sorace. 1996. Magnitude estimation of linguistic acceptability. *Language* 72:32–68.

Bates, Douglas, Martin Maechler, and Bin Dai. 2008. lme4: Linear mixed-effects models using S4 classes. R package version 0.999375-28. Available at http://lme4.r-forge.r-project.org/.

Becker, Michael, Nihan Ketrez, and Andrew Nevins. 2011. The surfeit of the stimulus: Grammatical biases filter lexical statistics in Turkish voicing deneutralization. *Language* 87:84–125.

Berent, Iris, Tracy Lennertz, Jongho Jun, Miguel A. Moreno, and Paul Smolensky. 2008. Language universals in human brains. *Proceedings of the National Academy of Science* 105:5321–5325.

Berent, Iris, Tracy Lennertz, Paul Smolensky, and Vered Vaknin-Nusbaum. 2009. Listeners' knowledge of phonological universals: Evidence from nasal clusters. *Phonology* 26:75–108.

Berent, Iris, Donca Steriade, Tracy Lennertz, and Vered Vaknin. 2007. What we know about what we have never heard: Evidence from perceptual illusions. *Cognition* 104:591–630.

Berent, Iris, Colin Wilson, Gary F. Marcus, and Douglas K. Bemis. 2012. On the role of variables in phonology: Remarks on Hayes and Wilson 2008. *Linguistic Inquiry* 43:97–119.

Berger, Adam L., Stephen A. Della Pietra, and Vincent J. Della Pietra. 1996. A maximum entropy approach to natural language processing. *Computational Linguistics* 22:39–71.

Blevins, Juliette. 2004. *Evolutionary phonology: The emergence of sound patterns*. Cambridge: Cambridge University Press.

Boersma, Paul. 1998. *Functional phonology: Formalizing the interactions between articulatory and perceptual drives*. The Hague: Holland Academic Graphics.

Chambers, Kyle E., Kristine H. Onishi, and Cynthia Fisher. 2003. Infants learn phonotactic regularities from brief auditory experience. *Cognition* 87:B69–B77.

Chomsky, Noam, and Morris Halle. 1965. Some controversial questions in phonological theory. *Journal of Linguistics* 1:97–138.

Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. New York: Harper and Row.

Clements, George N. 1990. The role of the sonority cycle in core syllabification. In *Papers in laboratory phonology I*, ed. by John Kingston and Mary Beckman, 283–333. Cambridge: Cambridge University Press.

Coetzee, Andries, and Joe Pater. 2008. Weighted constraints and gradient phonotactics in Muna and Arabic. *Natural Language and Linguistic Theory* 26:289–337.

Colavin, Rebecca, Roger Levy, and Sharon Rose. 2010. Modeling OCP-place in Amharic with the Maximum Entropy phonotactic learner. Paper presented at the 46th meeting of the Chicago Linguistic Society. To appear in the proceedings.

Côté, Marie-Hélène. 2000. Consonant cluster phonotactics: A perceptual approach. Doctoral dissertation, MIT, Cambridge, MA.

Daland, Robert, Bruce Hayes, James White, Marc Garellek, Andrea Davis, and Ingrid Norrmann. 2011. Explaining sonority projection effects. *Phonology* 28:197–234.

Davidson, Lisa. 2007. The relationship between the perception of non-native phonotactics and loanword adaptation. *Phonology* 24:261–286.

Dell, Gary S., Kristopher D. Reed, David R. Adams, and Antje S. Meyer. 2000. Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 6:1355–1367.

Della Pietra, Stephen A., Vincent J. Della Pietra, and John D. Lafferty. 1997. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19:380–393.

Dixon, R. M. W. 1981. Wargamay. In *Handbook of Australian languages*, ed. by R. M. W. Dixon and Barry J. Blake, 2:1–144. Amsterdam: John Benjamins.

Dupoux, Emmanuel, Kazuhiko Kakehi, Yuki Hirose, Christophe Pallier, and Jacques Mehler. 1999. Epenthetic vowels in Japanese: A perceptual illusion? *Journal of Experimental Psychology: Human Perception and Performance* 25:1568–1578.

Finley, Sara. 2008. Formal and cognitive restrictions on vowel harmony. Doctoral dissertation, Johns Hopkins University, Baltimore, MD.

Finley, Sara, and William Badecker. 2008. Substantive biases for vowel harmony languages. In *Proceedings of the 27th West Coast Conference on Formal Linguistics*, ed. by Natasha Abner and Jason Bishop, 168–176. Somerville, MA: Cascadilla Press.

Finley, Sara, and William Badecker. 2009. Artificial language learning and feature-based generalization. *Journal of Memory and Language* 61:423–437.

Flemming, Edward. 2001. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18:7–44.

Flemming, Edward. 2004. Contrast and perceptual distinctiveness. In *Phonetically-based phonology*, ed. by Bruce Hayes, Robert Kirchner, and Donca Steriade, 232–276. Cambridge: Cambridge University Press.

Gildea, Daniel, and Daniel Jurafsky. 1996. Learning bias and phonological rule induction. *Computational Linguistics* 22:497–530.

Goldwater, Sharon, and Mark Johnson. 2003. Learning OT constraint rankings using a maximum entropy model. In *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*, ed. by Jennifer Spenader, Anders Eriksson, and Östen Dahl, 111–120. Stockholm: Stockholm University, Department of Linguistics.

Greenberg, Joseph H. 1966. Some universals of grammar with particular reference to the order of meaningful elements. In *Universals of language*, ed. by Joseph H. Greenberg, 73–113. Cambridge, MA: MIT Press.

Greenberg, Joseph H. 1978. Some generalizations concerning initial and final consonant clusters. In *Universals of human language, vol. 2*, ed. by Edith A. Moravcsik, 243–279. Stanford, CA: Stanford University Press.

Harris, James. 1983. *Syllable structure and stress in Spanish*. Cambridge, MA: MIT Press.

Hayes, Bruce. 1999. Phonetically-driven phonology: The role of Optimality Theory and inductive grounding. In *Functionalism and formalism in linguistics, vol. I*, ed. by Michael Darnell, Edith Moravcsik, Michael Noonan, Frederick Newmeyer, and Kathleen Wheatly, 243–285. Amsterdam: John Benjamins.

Hayes, Bruce. 2004. Phonological acquisition in Optimality Theory: The early stages. In *Fixing priorities: Constraints in phonological acquisition*, ed. by René Kager, Joe Pater, and Wim Zonneveld, 158–203. Cambridge: Cambridge University Press.

Hayes, Bruce, Robert Kirchner, and Donca Steriade, eds. 2004. *Phonetically-based phonology*. Cambridge: Cambridge University Press.

Hayes, Bruce, and Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39:379–440.

Hayes, Bruce, Colin Wilson, and Anne Shisko. To appear. Maxent grammars for the metrics of Shakespeare and Milton. *Language*.

Hayes, Bruce, Kie Zuraw, Peter Siptár, and Zsuzsa Londe. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85:822–863.

Heinz, Jeffrey. 2009. On the role of locality in learning stress patterns. *Phonology* 26:303–351.

Heinz, Jeffrey. 2010. Learning long-distance phonotactics. *Linguistic Inquiry* 41:623–661.

Iverson, Gregory K., and Joseph C. Salmons. 2005. Filling the gap. *Journal of English Linguistics* 33:207–221.

Jarosz, Gaja. 2006. Rich lexicons and restrictive grammars: Maximum likelihood learning in Optimality Theory. Doctoral dissertation, Johns Hopkins University, Baltimore, MD.

Jespersen, Otto. 1909. *A modern English grammar on historical principles*. Part I, *Sounds and spellings*. London: George Allen & Unwin.

Kager, René. 1999. *Optimality Theory*. Cambridge: Cambridge University Press.

Kager, René, and Joe Pater. 2012. Phonotactics as phonology: Knowledge of a complex constraint in Dutch. *Phonology* 29:81–111.

Kawahara, Shigeto. 2008. Phonetic naturalness and unnaturalness in Japanese loanword phonology. *Journal of East Asian Linguistics* 17:317–330.

Kawasaki, Haruko. 1982. An acoustical basis for universal constraints on sound sequences. Doctoral dissertation, University of California, Berkeley.

Kenstowicz, Michael, and Charles Kisseberth. 1977. *Topics in phonological theory*. New York: Academic Press.

Kiparsky, Paul. 1982. Lexical phonology and morphology. In *Linguistics in the morning calm*, ed. by In-Seok Yang, 3–91. Seoul: Hanshin.

Lodge, Milton. 1981. *Magnitude scaling: Quantitative measurement of opinions*. Beverly Hills, CA: Sage.

Lombardi, Linda. 1999. Positional faithfulness and voicing assimilation in Optimality Theory. *Natural Language and Linguistic Theory* 17:267–302.

Massaro, Dominic W., and Michael M. Cohen. 1980. Phonological constraints in speech perception. *Journal of the Acoustical Society of America* 67, S26.

Mattys, Sven L., and Peter W. Jusczyk. 2001. Phonotactic cues for segmentation of fluent speech by infants. *Cognition* 78:91–121.

McClelland, James L., and Brent C. Vander Wyk. 2006. Graded constraints on English word forms. Ms., Stanford University, Stanford, CA. Available at http://psychology.stanford.edu/~jlm/papers /GCEWFs_2_18_06.pdf.

Morelli, Frieda. 1999. The phonotactics and phonology of obstruent clusters in Optimality Theory. Doctoral dissertation, University of Maryland, College Park.

Moreton, Elliott. 2002. Structural constraints in the perception of English stop-sonorant clusters. *Cognition* 84:55–71.

Moreton, Elliott. 2008. Analytic bias and phonological typology. *Phonology* 25:83–127.

Moreton, Elliott, and Joe Pater. In preparation. Learning artificial phonology: A review. Ms., University of North Carolina, Chapel Hill, and University of Massachusetts, Amherst. Available at http://www.unc .edu/~moreton/Papers/MoretonPater.Draft.3.5.pdf.

Myers, Scott. 1997. Expressing phonetic naturalness in phonology. In *Derivations and constraints in phonology*, ed. by Iggy Roca, 125–152. Oxford: Oxford University Press.

Ohala, John. 1981. The listener as a source of sound change. In *Papers from the Parasession on Language and Behavior*, ed. by Carrie S. Masek, Roberta A. Hendrik, and Mary Frances Miller, 178–203. Chicago: University of Chicago, Chicago Linguistic Society.

Onishi, Kristine H., Kyle E. Chambers, and Cynthia Fisher. 2002. Learning phonotactic constraints from brief auditory experience. *Cognition* 83:B13–B23.

Pater, Joe, and Anne-Michelle Tessier. 2003. Phonotactic knowledge and the acquisition of alternations. In *Proceedings of the 15th International Congress of Phonetic Sciences*, ed. by Maria-Josep Solé, Daniel Recasens, and Joaquín Romero, 1177–1180. Barcelona: Universitat Autònoma de Barcelona.

Peperkamp, Sharon, and Emmanuel Dupoux. 2007. Learning the mapping from surface to underlying representations in artificial language learning. In *Laboratory phonology 9*, ed. by Jennifer Cole and José Hualde, 315–338. Berlin: Mouton de Gruyter.

Peperkamp, Sharon, Katrin Skoruppa, and Emmanuel Dupoux. 2006. The role of phonetic naturalness in phonological rule acquisition. In *Proceedings of the 30th Annual Boston University Conference on Language Development*, ed. by David Bamman, Tatiana Magnitskaia, and Colleen Zaller, 464–475. Somerville, MA: Cascadilla Press.

Pierrehumbert, Janet. 2006. The statistical basis of an unnatural alternation. In *Laboratory phonology 8: Varieties of phonological competence*, ed. by Louis Goldstein, D. H. Whalen, and Catherine T. Best, 81–107. Berlin: Mouton de Gruyter.

Pinheiro, José, and Douglas M. Bates. 2000. *Mixed-effects models in S and S-PLUS*. New York: Springer.

Prince, Alan, and Bruce Tesar. 2004. Learning phonotactic distributions. In *Fixing priorities: Constraints in phonological acquisition*, ed. by René Kager, Joe Pater, and Wim Zonneveld, 245–291. Cambridge: Cambridge University Press.

Pullum, Geoffrey K., and William A. Ladusaw. 1996. *Phonetic symbol guide*. 2nd ed. Chicago: University of Chicago Press.

Pycha, Anne, Pawel Nowak, Eurie Shin, and Ryan Shosted. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. In *Proceedings of the 22nd West Coast Conference on Formal Linguistics*, ed. by Gina Garding and Mimu Tsujimura, 423–435. Somerville, MA: Cascadilla Press.

R Development Core Team. 2008. R: A language and environment for statistical computing. Vienna: R Foundation for Statistical Computing. Available at http://www.R-project.org.

Scholes, Robert. 1966. *Phonotactic grammaticality*. The Hague: Mouton.

Seidl, Amanda, and Eugene Buckley. 2005. On the learning of arbitrary phonological rules. *Language Learning and Development* 1:289–316.

Selkirk, Elisabeth O. 1982. The syllable. In *The structure of phonological representations, part II*, ed. by Harry van der Hulst and Norval Smith, 337–383. Dordrecht: Foris.

Sievers, Eduard. 1881. *Grundzüge der Phonetik*. Leipzig: Breitkopf und Härtel.

Skoruppa, Katrin, and Sharon Peperkamp. 2011. Adaptation to novel accents: Feature-based learning of context-sensitive phonological regularities. *Cognitive Science* 35:348–366.

Steriade, Donca. 1999. Alternatives to syllable-based accounts of consonantal phonotactics. In *Proceedings of the 1998 Linguistics and Phonetics Conference*, ed. by Osamu Fujimura, Brian Joseph, and B. Palek, 205–245. Prague: The Karolinum Press.

Steriade, Donca. 2001. Directional asymmetries in place assimilation: A perceptual account. In *The role of speech perception in phonology*, ed. by Elizabeth Hume and Keith Johnson, 219–250. New York: Academic Press.

Warker, Jill A., and Gary S. Dell. 2006. Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32:387–398.

Warker, Jill A., Gary S. Dell, Christine A. Whalen, and Samantha Gereg. 2008. Limits on learning phonotactic constraints from recent production experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34:1289–1295.

Whitney, William Dwight. 1889. *Sanskrit grammar*. Cambridge, MA: Harvard University Press.

Wilson, Colin. 2003. Experimental investigation of phonological naturalness. In *Proceedings of the 22nd West Coast Conference on Formal Linguistics*, ed. by Gina Garding and Mimu Tsujimura, 533–546. Somerville, MA: Cascadilla Press.

Wilson, Colin. 2006. Learning phonology with substantive bias: An experimental and computational investigation of velar palatalization. *Cognitive Science* 30:945–982.

Wilson, Colin, and Lisa Davidson. 2009. Bayesian analysis of non-native cluster production. Paper presented at the 40th meeting of the North East Linguistic Society, Cambridge, MA. To appear in the proceedings.

Yip, Moira. 2002. *Tone*. Cambridge: Cambridge University Press.

Department of Linguistics
3125 Campbell Hall
UCLA
Los Angeles, CA 90095-1543

bhayes@humnet.ucla.edu
jameswhite@ucla.edu