The Effect of Complexity versus the Effect of Naturalness on Phonological Learning

By
Brandon Prickett

Honors Thesis
Linguistics
University of North Carolina at Chapel Hill

2014

Approved:

____
Elliott Moreton, Advisor

Jennifer Smith, Reader

Jeff Mielke, Reader

ABSTRACT

Brandon Prickett: The Effect of Complexity versus the Effect of Naturalness on Phonological Learning
(Under the direction of Elliott Moreton)

Recently a "surfeit of the stimulus" has been noticed in languages; specifically speakers do not seem to be learning all of the phonological patterns that exist in their language (Becker et al. 2007, 2008, 2011). Using the Hayes and Wilson (2008) phonotactic learner software to create a list of phonological patterns that are in English words, Hayes and White (2013) found that phonological constraints that were "natural"—that is typologically common and phonetically logical—seemed to have been learned better by speakers. Recent research using artificial languages has suggested that complexity has a stronger effect on phonological learning than naturalness (see Moreton and Pater 2012b for a review of experiments of this nature).

The current study seeks to use Hayes and White's methodology to test the effects of both naturalness and complexity on phonological learning. Subjects took part in an online experiment that presented an orthographic representation and an audio recording for each of the stimuli. Subjects were asked to numerically rate each of the words on how "good" they sounded in English and ratings of the experimental stimuli were compared with control group partners that did not break any constraints to determine how much of an effect each constraint had on the speakers' judgments. The stronger a constraint's effect was, the better learned it was assumed to be. Naturalness seemed to affect learning more than complexity—contradicting the findings of many recent studies on artificial language learning (see above citation) and suggesting that there could be significant differences in the processes that speakers use to learn real-world languages such as English and those used to learn artificial languages in a lab.

ACKKNOWLEDGEMENTS

First of all thanks goes to the Steve and Courtney Kapp Fund for Academic Excellence for funding the experiment for this thesis. Thanks to my family for all their support and to my father especially for his assistance and advice with programming some of the software needed for the thesis. For listening to me ramble and allowing me to discuss ideas with her on a fairly regular basis, thanks to Amanda Hayes. Thanks to Rachel Broad for teaching me how to run experiments over the web and the rest of P-Side for their helpful advice. Chris Wiesen at the Odum Institute deservers a large amount of thanks for his patience with someone who had never before dealt with a statistical analysis. Thanks to my thesis committee members, Jennifer Smith and Jeff Mielke, whose input was invaluable. And most of all, thanks to Elliott Moreton for facilitating this project, guiding me during it, and helping to refine it into what it is today. Professor Moreton has been an immense influence in not only making me a better thesis writer, but a better linguist in general.

1. Introduction

Recently, a phenomenon known as a "surfeit of the stimulus" has been noticed in language learning (Becker et al. 2008). This means that although children are presented with all of the phonological patterns within the language they are learning, they only seem to learn a subset of that information. For instance, Turkish consonants have laryngeal alternation that is correlated with three different factors: location in the word, word length, and vowel quality (Becker et al. 2008). While Turkish speakers are able to predict the laryngeal alternation based off of the first two variables, they could not predict the phenomenon using vowel quality. Becker et al. suggest that this could represent a bias in speakers for learning certain types of phonological patterns.

Hayes and White (2013) suggest that some phonological patterns could be considered "natural" and that language learners could have a bias towards learning these types of patterns. This bias could be how speakers subconsciously choose which phonological constraints to learn amidst the surfeit of possibilities. "Natural" constraints, according to Hayes and White, are those constraints which ban phonetically difficult forms (such as a voice assimilation constraint) or those constraints which are typologically common (such as a coda deletion constraint)[1]. Hayes and White (2013) experimentally tested the effect of naturalness on constraint learnability in order to gauge how much naturalness might be affecting speaker learning. If speakers were to be significantly biased towards natural constraints, it could explain the fact that children are not learning all of the constraints they are seemingly presented with. This is described in more detail in Section 2.2.

Another factor in phonological learning is the effect of complexity. Shepard et al. (1961) showed that people tend to learn simpler visual patterns better than complex ones. The study performed by Shepard et al. involved patterns in which shapes were placed into categories. Patterns that were simple involved categories that were based off of a shape's characteristics, such as "all the black shapes". Patterns that were more complex involved more arbitrary categories, such as "all triangles that are not both white and small, and big, black squares". When learning patterns, subjects made fewer mistakes when learning a categorically simpler pattern. This type of complexity and how it relates to the (somewhat different) form of complexity tested in this study will be discussed more in Section 2.3.

Linguists have shown that complexity, as defined by Shepard et al. (1961) also affects subjects' ability to learn artificial language phonology. When presented with a phonological pattern in an artificial language, both infants (Saffran and Thiessen 2003) and adults (Skoruppa and Peperkamp 2011) tend to learn simple phonological patterns but are worse at learning

[1] Although Hayes and White (2013) define naturalness as *either* being typological or phonetic, all of the natural constraints they test have both of these characteristics in common. The current study also required that constraints be both typologically attested and phonetically logical in order to be considered natural.

complex ones. Despite growing evidence for the effect of complexity in phonological learning, the experiment conducted by Hayes and White (2013) did not test for the effect of this variable.

The current study seeks to determine the individual effects of complexity and naturalness on phonological learning. Following the methods of Hayes and White (2013), grammaticality judgments of constraint-violating words were used to gauge how well speakers of American English learned the phonological constraints that seem to be present in their language. Unlike Hayes and White, this study systematically varies the naturalness and complexity of the constraints being tested.

Phonological constraints were considered to fall into one of four categories: Natural Simple, Natural Complex, Unnatural Simple, and Unnatural Complex. If complexity and naturalness both effect the learnability of constraints, then the Natural Simple constraints should be learned best by speakers and the Unnatural Complex constraints should be learned the worst. Alternatively, if one of the variables has a considerably larger effect than the other, then the less effective variable could be eclipsed by its counterpart. For instance, if naturalness has a much larger effect than complexity, then the Unnatural Complex and the Unnatural Simple constraints might not show any difference in their learnability, since their unnatural status could have made them equally hard to learn.

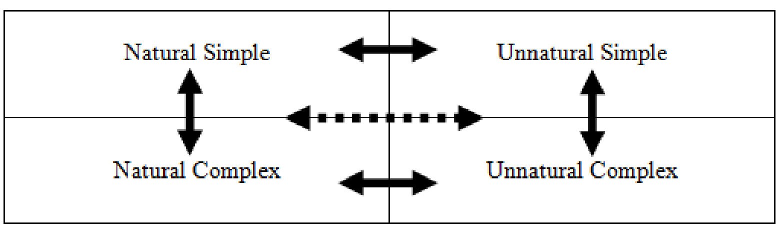

Figure 1. Four different categories that phonological constraints could fall under. The solid arrows represent comparisons made in the current study, while the dotted line represents the comparison made in Hayes and White (2013).

2. Research Background

In order to test how speakers learn their phonology, Hayes and Wilson created phonotactic learning software (see Section 2.1). Seemingly accidental constraints were learned by the Hayes and Wilson (2008) learner when it was applied to the language of Wargamay and were deemed "unnatural" Hayes and Wilson. A set of unnatural constraints such as these that were learned by the Hayes and Wilson learner when it was applied to English were tested by Hayes and White (2013). Hayes and White's findings suggest that speakers learn natural constraints better than unnatural ones (see Section 2.2); however many recent studies suggest that complexity has more of an effect than naturalness on phonological learning (see Section

2.3). The current study seeks to use Hayes and White's methodology to discover more about which learning bias—complexity or naturalness—is affecting phonological learning more.

2.1 Hayes and Wilson (2008)

In Hayes and Wilson (2008), the Hayes and Wilson Phonotactic Learner was used to simulate phonological learning in humans. The Learner is a computer program specifically designed to take a corpus of words in a language combined with a phonological feature system and use these to find phonotactic patterns within a language. The program uses a maximum entropy model of grammar to take in only positive evidence (the list of words in a language) and create a list of constraints as its output that are each assigned a weight. The weight of each constraint represents how unlikely it is for that particular constraint to be violated. In other words, the higher the weight of a constraint, the more often this constraint is shown to have an effect on the words in the language.

Figure 2 illustrates this, using an imaginary (and unrealistically simplistic) language and showing an output like what the Hayes and Wilson Learner might create. The language's lexicon and phoneme features are given to the learner as input (boxes "a" and "b" in Figure 2). The learner would then analyze the lexicon for constraints that seem to exist there. For instance, in the imaginary language illustrated below, there are no words that begin with a vowel. The learner would create the constraint *#[vowel][2] to describe this and then assign the constraint a weight. The weights in Figure 2 are just estimates as to what the real learner would assign to these constraints.

When creating constraints, the learner prioritizes accuracy of the constraint, which is why there is not a constraint like *[vowel] in the output. This constraint, while describing the fact that there are no vowels at the beginning of words, is also violated by all five of the words in the dataset. After accuracy of a constraint is ensured, more general constraints are prioritized above overly specific ones. This is why constraints like *#[i], *#[e], and *#[a] would not be created by the learner. These constraints are all more specific versions of the more general constraint *#[vowel]. The weight value of 5 for the first two constraints represents that these constraints are never violated in the input. The weight value of 4 for the *#[s] constraint represents the fact that one of the words in the data set violates this constraint, showing that the constraint is likely, but not certain, to have an effect in the imaginary language.

[2] The learner's constraint would not look exactly like those given in Figure 2 as an example. The learner describes segments and natural classes with binary features, so instead of [vowel], it would use [+syllabic], instead of # it would use [+word_boundary], and instead of [s] it would use something like [+strident, +anterior, -voice].

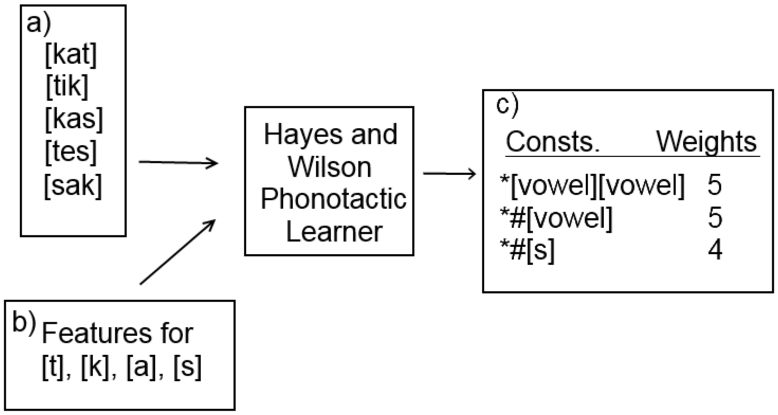

Figure 2. An illustration of the Hayes and Wilson (2008) phonotactic learner. In this example, the learner is given an imaginary language's lexicon (a) and phonetic features for its phonemes (b), and creates weighted constraints (c) to describe the input data. The weighting algorithm in this example is simplified for the sake of illustration.

In the Hayes and Wilson (2008) study, the Phonotactic Learner was used to analyze multiple languages in order to test if the software was using a phonological learning system that was analogous to that used by humans. The learner was used to create constraints for English onsets, a “lexicon” of English onsets was used as the input, along with the features for all English phonemes. The constraints created were then compared to a corpus of English words and to experiments that have been performed on English grammaticality judgments. Hayes and Wilson found that both the corpus and the experimental results supported the constraints that their learner had created. For example, the learner created the onset constraint *[+continuant, +voice, -anterior], essentially “*[ʒ] in the onset”. A corpus of English words would support this constraint since words beginning with [ʒ] are almost unattested. Hayes and Wilson then used results from Scholes’ (1966) grammaticality experiment to show that English speakers did not prefer onsets that began with [ʒ] (see Hayes and Wilson 2008: Figure 3).

However, when the Learner was applied, using a similar process, to the language Wargamay, it discovered a number of “accidentally true” constraints. Table 1 shows examples of these. These constraints were not recorded by any linguistics that had studied Wargamay in the past and also were not phonetically grounded. Hayes and Wilson suggested that this could be an error in the Phonotactic Learner, since these Wargamay constraints were not supported by data in the same way that the English onset constraints described above were.

Table 1. Examples of "Accidentally True" Constraints from Wargamay

*[+sonorant,+labial][-back,-stress]
*[+sonorant,+dorsal][-mainstress,+stress]
*[+back][+long][+approximant]

2.2 Hayes and White (2013)

Hayes and White (2013) sought to discover whether the Hayes and Wilson (2008) Phonotactic Learner's "accidentally true" constraints were a part of speakers' mental grammars. Since the Learner analyzes the same data that a newborn infant would have had his or her disposal (words from a language), humans should be learning the same patterns in a language that the Phonotactic Learner finds. If humans are not learning these patterns, it could mean that there is a bias that makes some patterns more favorable to learn than others.

In order to determine if humans have an innate bias towards some kinds of phonotactic constraints, Hayes and White used the Phonotactic Learner to analyze American English. The CELEX word database was used (Baayen, Piepenbrock and Gulikers 1995) to find words in English that were listed as having a frequency of at least 1. American pronunciations for these words were found in the Carnegie-Mellon Pronouncing Dictionary[3].

Finally, all compounds, inflected forms, and "forms created by highly transparent processes of morphological derivation" were removed from the database since these are known to have phonotactic patterns different from those in morphologically simple words. Subjects in their experiment would likely interpret stimuli as being morphologically simple, so this ensured that the model and the subjects were analyzing on the same morphological level.

The Hayes and Wilson Phonotactic Learner created 160 constraints from the above list of words, and out of these 160, 20 were chosen by Hayes and White for use in their study. 10 of these constraints that were well attested among languages and were phonetically logical—such as voicing assimilation—were called "natural" for the purposes of the study. The other 10 represented constraints that seemed to be "accidentally true", as in the case of some the Wargamay constraints from Hayes and Wilson (2008). Although these constraints did represent patterns from the English database, they did not seem to be phonetically motivated, nor were they typologically common and were called "unnatural" by Hayes and White.

Novel words were created that violated each of these constraints in order to test the effect of each constraint on speakers' grammaticality judgments. Twenty-nine native English speakers were presented with novel words that each violated one of the constraints and a control group of words that did not violate any constraints. The control group words were as similar to the violating group words as possible without actually violating the constraints. The participants in the study were asked to perform magnitude estimations of how "good" the words sounded to them. They expressed this judgment both numerically and through drawing lines of different lengths (the better sounding words received higher numerical ratings and longer lines). The line lengths and numerical ratings of violating words were compared with those of their control group partners to determine how much each constraint affected the speakers' judgments.

---

[3] http://www.speech.cs.cmu.edu/cgi-bin/cmudict

If natural constraints had more of an effect on speaker judgments than unnatural constraints, then the difference between judgments of words that violated natural constraints and their control group partners would be larger than the difference for words that violated unnatural constraints and their control group partners. For example, if a voicing assimilation constraint had a large affect on subjects' judgments, this would mean that the subjects would rate words that violated voicing assimilation as very low, causing a large difference between these words and their control group partners. Alternatively, if a constraint against diphthongs before [ɡ] did not affect judgments very much, ratings of words violating this constraint would be about the same as the ratings of their control group partners.

Hayes and White found that on average, the difference in subject ratings for words that violated unnatural constraints and ratings for the words' control group partners were not significantly different, meaning that people did not seem to be affected by the unnatural constraints. Words that violated the natural constraints, on the other hand, were given significantly lower ratings, on average, than their control group partners. Figure 3 shows both of these results. These findings suggest that natural constraints were being learned by speakers while unnatural constraints were not, or, at the very least, that unnatural constraints were not being learned as well as natural constraints.

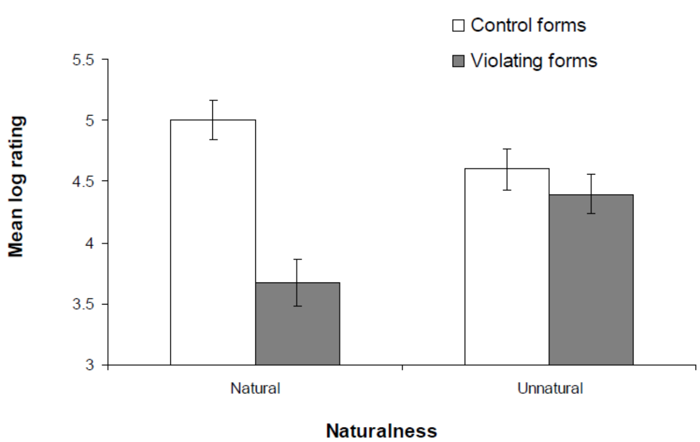

Figure 3. The mean log ratings of words that violate natural and unnatural constraints and their control group partners. The larger the difference between the white bar and the gray bar, the larger effect that group of constraints had on native speaker grammaticality judgments (from Hayes and White 2013).

Becker et al. (2008), mentioned in Section 1, discovered similar results in their experiments on Turkish laryngeal alternation. They attributed this preference for learning natural constraints to a universal set of learning biases that point people towards learning certain patterns within their language. Hayes and White suggest these biases could arise from a "phonetically driven" phonology (2013:24) which would mean that speakers are more prone to learning constraints that aid in the production and perception of speech sounds.

2.3 Studies on the Effects of Complexity on Artificial Language Learning

Within this section, two different types of complexity will be discussed: dimensional complexity and categorical complexity. These terms are not used in past literature. They are used here because of the relevance of both types and the need to easily differentiate between the two of them. Dimensional complexity is the idea that the more relevant variables there are in a pattern, the harder the pattern will be to learn. Bulgarella and Archer (1962) showed that a pattern involving one variable (in their case the pitch of a sound) was learned faster than a pattern involving two variables (in their case, pitch and loudness) which was easier than three variables, and so on. Chomsky and Halle (1968) refer to this kind of complexity in terms of phonology by saying that a phonological rule that requires more phonetic features to describe (such as "delete all [+voice, +Coronal, +continuant, +consonantal] segments") would be more complex than a rule that only requires a small number of features to describe (such as "delete all [+voice] segments").

Categorical complexity, described in Section 1 and used by Shepard et al. (1961) is the idea that a pattern involving categories is simpler if the categories refer to specific characteristics and are not arbitrarily arranged. Essentially, Shepard et al.'s version of complexity can be described as the idea that the more a category can be described by generalizations of characteristics, the simpler the pattern that matches it is. An example of these categories can be seen in Figure 4 (taken from Shepard et al. 1961). Each category in the below illustration differentiates between the shapes on the left of its box and the shapes on the right of its box with a pattern. Category 1 patterns only require a single, simple generalization: "black shapes on the left". Category 2 is more complicated, using the pattern "black triangles and white squares on the left". Category three increases pattern complexity even more with "large, black shapes and small triangles on the left", and so on with each category forcing less characteristic generalizations to be learned. Using this definition of complexity, Shepard et al. found that generally, simpler visual patterns were learned faster than complex visual patterns.

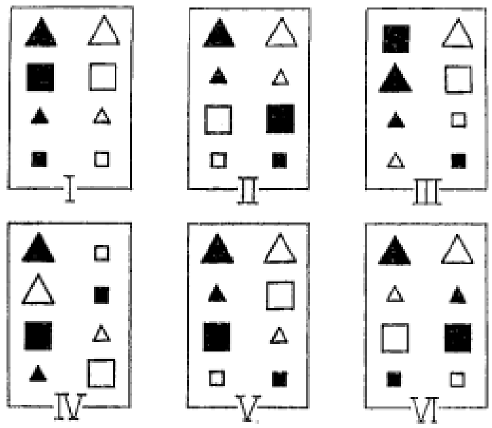

Figure 4. An illustration of the six possible categories of patterns in what is being called "categorical complexity" for the purposes of this study. Each category has a way of differentiating between the shapes on the left of the box and the shapes on the right. For category I, the pattern to differentiate between the two groups is simply "black shapes on the left". In category II, it becomes more complex: "black triangles and white squares on the left". Category III furthers this complexity with "large black shapes and small triangles" and so on for the remaining three categories. Each category forces less characteristic generalizations to be made than the category before it. Taken from Shepard et al. (1961).

Both categorical and dimensional complexity have been shown to affect non-linguistic pattern learning. Numerous studies have also shown that categorical complexity affects phonological learning (Saffran and Thiessen 2003 and Skoruppa and Peperkamp[4] 2011, for example). These studies, generally speaking, involved teaching an artificial phonological pattern to subjects and then testing them to see how well they learned it. In the artificial languages the categorically simpler phonotactic patterns were consistently easier for subjects to learn. This suggests that at least categorical complexity (and possibly dimensional complexity) could be useful for predicting how well people might learn phonotactic constraints in the real world, if learning artificial languages in a laboratory and learning natural languages use the same process.

For the purposes of this experiment, dimensional complexity will be used. Although artificial learning experiments have not been performed for this kind of complexity in phonology, the similar effects that the two types of complexity discussed in this section have on non-linguistic learning (seen in Shepard 1961 and Bulgarella and Archer 1962) suggest that this should not cause much discrepancy.

[4] It is important to note that not only did Skoruppa and Peperkamp find an effect of complexity in their experiment, but they also tested for and did not find an effect of naturalness. In fact, naturalness has not been seen to affect phonological learning in a number of artificial language learning experiments (reviewed in Moreton and Pater 2012b). This is discussed more in Section 6.

3. Research Design

Studies on real world languages suggest an effect of naturalness (such as Hayes and White 2013 and Becker et al. 2008) while experiments testing artificial language learning suggest that categorical complexity should have a large effect on phonological learning. To test the effects of these variables, the methodology of Hayes and White (2013) will be used. As in Hayes and White (2013), phonological constraints that were created by the Hayes and Wilson (2008) Phonotactic Learner will be categorized as natural or unnatural. For this study a constraint must be both typologically attested and phonetically motivated to e considered natural. In addition to this factor, constraints will also be categorized as complex or simple based on how many phonetic features are required to describe the constraint (for example, *[+voice][-voice] would be simple, while *[+round,-tense][-sonorant,+LABIAL,-rhyme] would be complex) as in Chomsky and Halle (1968). This represents the dimensional complexity (described in Section 2.3) of the constraints which, as discussed above, should have an analogous effect to categorical complexity.

As in Hayes and White (2013), subjects will be asked to give a grammaticality judgment of a word. Unlike Hayes and White (2013), in which there were four categories of stimuli (natural violating, natural control, unnatural violating, and unnatural control), the current study makes use of eight stimulus categories: natural simple violating, natural simple control, natural complex violating, natural complex control, unnatural simple violating, unnatural simple control, unnatural complex violating, and unnatural complex control. See Table 2 for examples of stimuli used in each of these eight categories. The four additional categories are a result of the added dimension in the groups of stimuli. Words that subjects will be judging will violate constraints that belong to one of four groups (natural simple, natural complex, unnatural simple, and unnatural complex) and then will be compared, as in Hayes and White (2013) to control words that do not violate any constraints, but that resemble one of the stimuli in the violating categories.

Table 2 - Examples of Stimuli Used for Each Category

Stimulus	Category	Constraint Violated
***esger* [ɛsgɹ]**	natural simple violating	*[-voice][+voice]
***ezger* [ɛzgɹ]**	natural simple control	n/a
***quinnk* [kwInk]**	natural complex violating	*[+nasal,+CORONAL][+DORSAL,+rhyme]
***quisk* [kwisk]**	natural complex control	n/a
***bowsh* [baʊʃ]**	unnatural simple violating	*[+diphthong][+continuant,-anterior]
***boosh* [buʃ]**	unnatural simple control	n/a
***moig* [moIg]**	unnatural complex violating	*[+diphthong,+round][+voice,+DORSAL]
***moid* [moId]**	unnatural complex control	n/a

If both naturalness and dimensional complexity affect phonological learning, then two results would be expected in subject grammaticality judgments: the difference between judgments of words that violate simple constraints and their control group partners should be greater than the difference between judgments of words that violate complex constraints and their control group partners. Also, the difference between judgments of words that violate natural constraints and their control group partners should be greater than the difference between judgments of words that violate unnatural constraints and their control group partners. Alternatively, if one of these factors is predominant above the other, than the other variable's effect could be eclipsed so that subjects group words into just two categories: very good or very bad.

4. Methods

4.1 Finding Constraints to Test

In order to determine what characteristics of phonological constraints might cause them to have a greater effect on speakers' judgments, 160 constraints produced by the Hayes and Wilson (2008) Phonotactic Learner were examined. Some constraints from this list were necessarily ignored, such as *[+word_boundary][+rhyme] ("no codas at the beginning of a word") since violating these constraints would be impossible. The rest were individually examined to determine their naturalness, complexity, and overall suitability for the experiment. This examination led to 19 constraints being chosen for the experiment, each being assigned to one of four categories: Natural Simple, Natural Complex, Unnatural Simple, and Unnatural Complex. Effort was also made to include both constraints that had prosodic information (the +/-rhyme feature in the system used by the Hayes and Wilson Phonotactic Learner) and those that did not in each category[5].

A constraint was considered natural if it could be reasonably assumed to be phonetically motivated (such as *[-voice][+voice]) and if constraints like it existed in other languages. Examples and discussion of this justification process for each constraint can be seen below. Checking for a constraint's existence in other languages involved searching through a database of phonological rules from P-Base, a database of phonological processes from 549[6] languages (Mielke 2008). These processes were analyzed to determine if they could be caused by constraints similar to the constraint in question. Individual constraint naturalness justifications will now be discussed in detail—for a summary of the discussion, see Table 3.

---

[5] An effect of prosody is not discussed in this paper, due to the somewhat uneven distribution of prosodic constraints in the four main categories.

[6] According to http://english.chass.ncsu.edu/faculty_staff/jimielke this is the number, P-Base data was obtained from Elliott Moreton, not directly from Jeff Mielke.

4.1.1. Natural Simple Constraints

The natural simple constraints were largely the same as the natural constraints used in Hayes and White (2013). *[+syllabic][+syllabic] was the only additional constraint used in this category. It was not chosen by Hayes and White (2013), but does represent a cross-linguistically common restriction on vowel hiatus (Casali 1997).

4.1.2. Natural Complex Constraints

These 5 constraints will be given individual attention, since they were almost all originally chosen for the current study.

*[+consonantal,-anterior][-sonorant], the first natural, complex constraint was chosen because it represented sonority ordering (a phonetically natural way of making consonant clusters easier to pronounce) and because a similar constraint has been observed in Sanskrit (Hayes and White 2013:21). The constraint does not directly refer to sonority in the first segment it describes, but in the Hayes and Wilson (2008) Learner's feature system, the description of [+consonantal,-anterior] only refers to sounds in English that have a certain level of sonority (such as the affricates [tʃ] and [dʒ]).

*[+diphthong,+round][+LABIAL] and *[+round,-tense][-sonorant,+LABIAL,-rhyme] were chosen because they demonstrated dissimilation between round vowels and the labial consonant that follows them. Dissimilation is described by Hayes and White (2013) as phonetically natural (for example, their natural constraint *[+LABIAL][+LABIAL]), and there is a similar rule to this constraint in the language of Kilivila (also known as Kiriwina) in which labial consonants are prohibited before round vowels (Mielke 2008).

*[-low,+tense][+sonorant,+DORSAL] was chosen because it represents assimilation of vowel and consonant height. This constraints prohibits –low vowels from coming before +DORSAL consonants. Dorsal consonants could also be considered +low because of the tongue's position in the mouth when articulating them. A similar rule which causes height assimilation in the presence of [-low,+tense] vowels can be seen in Ejagham (Mielke 2008).

*[+nasal,+CORONAL][+DORSAL,+rhyme] represents place of articulation assimilation, something that Hayes and White (2013) describe as phonetically natural (in their *[+labial][+dorsal] constraint). It is also attested in the language of Muna where nasal, coronal clusters are always homorganic.

4.1.3. Unnatural Constraints

Phonological rules similar to the unnatural constraints were also searched for in P-Base (a list of all constraints can be viewed in the Appendix). The constraints used in this experiment were not found in the database. In addition to this, each unnatural constraint was examined for phonetic justification and none was found.

Table 3 - A Summary of the Phonetic and Typological Justifications for Each Natural Constraint Used in the Experiment

Constraint	Simple/ Complex?	Phonetic Justification	Typological Justification
***[-voice][+voice]**	Simple	Voicing assimilation	See Hayes and White for a detailed typological justification
***[+syllabic][+syllabic]**	Simple	Restriction on vowel hiatus	"Many languages do not tolerate vowel hiatus." (Casali 1997)
***[-consonantal,-rhyme][+consonantal,-rhyme]**	Simple	Sonority ordering	See Hayes and White for a detailed typological justification
***[-continuant,-rhyme][-continuant,-rhyme]**	Simple	Sonority ordering	See Hayes and White for a detailed typological justification
***[-sonorant,+rhyme][+sonorant,+rhyme]**	Simple	Sonority ordering	See Hayes and White for a detailed typological justification
***[+consonantal,-anterior][-sonorant]**	Complex	Sonority ordering	Similar constraint in Sanskrit (Hayes and White 2013:21)
***[+diphthong,+round][+LABIAL]**	Complex	Dissimilation between round vowels and labial consonants	Labial sounds prohibited before +round vowels in Kilivila/Kiriwina
***[-low,+tense][+sonorant,+DORSAL]**	Complex	Assimilation between high vowels and -high consonants (+DORSAL)	[-low,+tense] vowels trigger assimilation in Ejagham
***[+nasal,+CORONAL][+DORSAL,+rhyme]**	Complex	Assimilation of place of articulation in the coda	Nasal coronal clusters are homorganic in Muna
***[+round,-tense][-sonorant,+LABIAL,-rhyme]**	Complex	Dissimilation between round vowels and labial consonants	Labial sounds prohibited before +round vowels in Kilivila/Kiriwina

4.1.4 Complexity

Constraints were considered complex if they took a larger number of phonetic features to describe (the dimensional complexity described in Section 2.3 above). To describe their complexity, constraints were each given a complexity score. This was found by calculating the mean number of phonetic features in each of the segments described in a constraint. The feature +/-rhyme was considered to be less relevant to the question of complexity because of its prosodic nature. To account for this, +/-rhyme was only counted as half of a feature when calculating scores and was not taken into consideration at all when grouping the constraints.

Table 4 – Dimensional Complexity Scores for Constraints Used[7]

Natural Simple	**1.30**	**Unnatural Simple**	**1.50**
*[-voice][+voice]	1.00	*[+diphthong][+continuant,-anterior]	1.50
*[+syllabic][+syllabic]	1.00	*[+round,-back][-anterior]	1.50
*[-consonantal,-rhyme][+consonantal, -rhyme]	1.50	*[+CORONAL,-rhyme][-back, -syllabic]	1.75
*[-continuant,-rhyme][-continuant, -rhyme]	1.50	*[-sonorant,+rhyme][-low]	1.25
*[-sonorant,+rhyme] [+sonorant,+rhyme]	1.50		
Natural Complex	**1.80**	**Unnatural Complex**	**2.00**
*[+consonantal,-anterior][-sonorant]	1.50	*[+round,+high][-consonantal, -sonorant]	2.00
*[+diphthong,+round][+LABIAL]	1.50	*[+diphthong,+round][+voice, +DORSAL]	2.00
*[-low,+tense] [+sonorant,+DORSAL]	2.00	*[+consonantal,-anterior][+high, -syllabic]	2.00
*[+nasal,+CORONAL][+DORSAL, +rhyme]	1.75	*[+diphthong,+round][-strident, -rhyme]	1.75
*[+round,-tense] [-sonorant,+LABIAL,-rhyme]	2.25	*[+round,-tense] [+continuant,+voice,+rhyme]	2.25

Table 4 shows the complexity scores that were found for each constraint. The mean complexity score for each simple group is .5 less than its complex counterpart.

4.2 Procedure

4.2.1 Participants. 77 participants' data were used in this study. 11 additional participants completed the experiment, although there were issues with their data due to the server that the experiment was being run on. These subjects were presented with each stimulus more than once, resulted in multiple judgments for each stimulus. Since there would be no objective way to choose which of the repeat judgments to use, the data from these 11 participants was discarded. Amazon.com's Mechanical Turk website (https://www.mturk.com) was used to anonymously recruit subjects and to compensate them for their time. Subjects were only able to participate if they lived in America and were told not to take part in the study unless American English was their first language.

4.2.2 Materials. A total of 79 stimuli were presented to each participant. 3 of these were fillers ("poik", "bzarshk", and "kip") and were used to ensure the participants were actively attempting the experiment (they had been told in the instructions what the relative ratings of these words should be). The other 76 stimuli were divided into two categories: control and experimental. The

[7] Scores that are in bold are the mean score for the category to their left.

experimental stimuli were created to specifically violate a constraint that was being testing in the experiment, while the control group words each differed from one of the experimental words by as little as possible while violating none of the constraints in the experiment[8]. Each stimulus was run through a computer program that searched the 160 constraints made by the Hayes and Wilson (2008) Phonotactic Learner mentioned above and ensured that each stimulus only violated the constraints from this list that it was supposed to (or none at all for the control group). Individual stimuli are listed in the appendix.

Audio recordings of the stimuli were made using a Logitech, Model Number A-00008 headset microphone. Words were read aloud into the microphone by the author, who is a native English speaker from the southeastern United States. The recordings were made in a sound proofed room straight onto a computer and were later checked for audio quality by the author. A sample of the recordings was also checked for quality by this study's research advisor, Elliott Moreton.

4.2.3 Apparatus. Subjects were directed from Mechanical Turk to a website created by the author. The website ran off of a perl CGI script that presented subjects with the following html pages: a page to verify that their computer had working audio output, a page with instructions on the experiment, a series of pages that asked the subjects for grammaticality judgments of the stimuli (which were presented in a random order), and then a page that thanked them for their time and provided them with a code to use for compensation from Mechanical Turk.

Audio Verification and Instructions: The audio verification involved the subjects first listening to a recording of the author saying the phrase "a cup of tea". After this, they were asked to pick which of four written phrases matched the recording. If they chose correctly, they were directed to the experiment instructions. If they chose incorrectly, they were taken to a webpage that stressed how important listening to the audio recordings was in the experiment and provided them with a link that could take them back to the verification page. This method of audio verification could have obviously been worked around by a determined or lucky participant guessing what that recording said without actually listening to it. However, making the task any more difficult could have also excluded honest participants that simply misunderstood what was said.

The instruction page for the website presented the following text:

> "Languages have rules that determine how well words sound in that language. For example, in English, 'bzarshk' would sound very odd but the word 'kip' would sound fine, even though neither of them are actual words. For the following words, rate numerically how acceptable they would be in English. Let

[8] There was one exception to this. The experimental stimuli used to test the effect of *[+nasal,+CORONAL][+DORSAL,+rhyme], "quinnk" and "zannk" had to also violate *[+CORONAL,+rhyme][+DORSAL,+rhyme]. To compensate for this, their control group partners, "quisk" and "zask" were also made to violate *[+CORONAL,+rhyme][+DORSAL,+rhyme].

‘poik’ be a value of 100. So anything better than this (for example, ‘kip’) would receive a higher score and anything worse (such as ‘bzarshk’) would receive a lower score. Be sure to listen to the audio clips and read the words so that you understand how the word is supposed to be pronounced.”

Below the text was a button for the participants to click labeled “Start Experiment”. This brought subjects to the first in a series of pages asking them to judge novel words.

Grammaticality Judgments: After reading the instructions, the subjects were presented with a series of pages that presented all of the stimuli along in an audio recording and an orthographic representation. Subjects were reminded on each page that “poik” would be given a rating of 100 and that better words would receive higher scores and worse words would receive lower scores. Below the audio recording was a place for the participant to input their grammaticality rating and beside this was a button labeled “Next Word”. See Figure 5 for a screenshot of what this would have looked like for the participants. Each rating was recorded after the participants pressed “Next Word”.

Figure 5. A screenshot of the webpage that subjects would have seen while participating in the experiment.

5. Results

5.1 Data for the Constraint Groups

The online magnitude estimation experiment resulted in the following data[9]. The mean log ratings for words that violated constraints were compared with the mean log ratings of words that did not in order to determine the effect that each constraint had on the speaker judgments. Figure 6 illustrates this comparison.

[9] Data from four constraints were not used in the data presented in Section 5.1 due to complications discussed in Section 5.2 below. The full data set represents similar patterns, but does not show them as clearly.

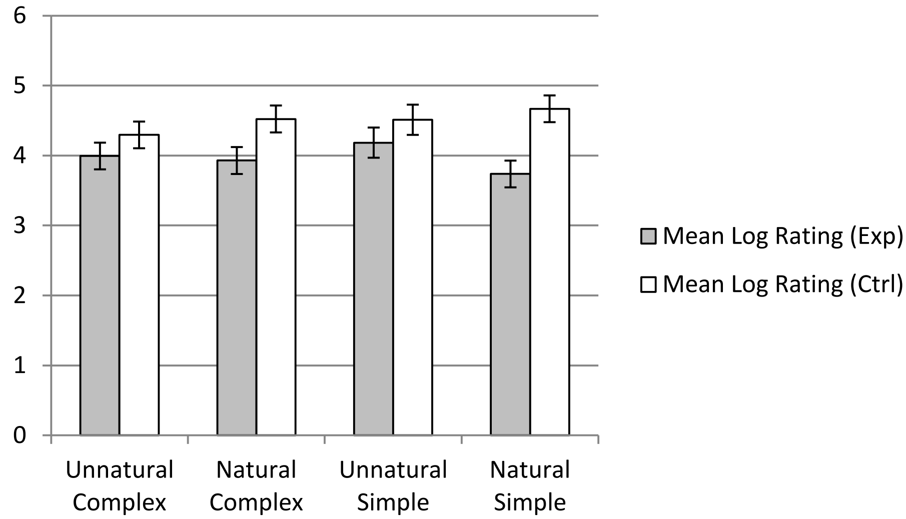

Figure 6. A comparison between the mean log ratings of each group. Control groups (stimuli that did not violate any constraints) are in white and experimental groups (those that did violate constraints) are in gray.

The above figure shows that naturalness, as in Hayes and White's (2013) experiment, has a significant effect on speaker judgments, since words that violated natural constraints were rated significantly lower than their counterpart words that did not. This effect is shown to be significant in both the simple and complex categories. Complexity does not seem to be showing as much of an effect. The natural complex category has a significant difference between its control and experimental groups while the unnatural simple category does not.

Table 5 shows a mixed effects model similar to the one in Table 1 of Hayes and White (2013). The effect of each variable that was relevant to the stimuli is shown as well as the upper and lower confidence intervals for each estimate. The baseline value, which represents the ratings for simple, natural, control stimuli is marked "intercept". The line below intercept shows that the experimental (that is, constraint-violating) stimuli had significantly lower ratings than the control stimuli (those that did not violate any constraints). This representation of the data shows the same trends as Figure 6, with the "Experimental & Unnatural" effect being much larger than the almost-insignificant "Experimental & Complexity" effect.

When compared to the results of Hayes and White (2013), the current results do not differ significantly in any of the effect sizes that both studies share. Table 6 demonstrates this by placing the two side-by-side for all of the relevant effects. The lower and upper confidence levels (CL in the table) overlap for each effect. This suggests that although there were some methodological differences between the two (for instance, the fact that the current experiment was run on the internet while Hayes and White performed their experiment in a lab) did not have an effect on the accuracy of either study's findings.

Table 5 – Mixed Effect Model for Naturalness, Complexity, and Control/Experimental Status

Effect	Estimate	Lower Confidence Limit	Upper Confidence Limit
Intercept	4.6691	4.4795	4.8587
Experimental	-0.9321	-1.1692	-0.6951
Unnatural	-0.1546	-0.4101	0.1008
Experimental & Unnatural	0.6052	0.2435	0.9669
Complex	-0.1503	-0.3869	0.08621
Unnatural & Complex	-0.0678	-0.416	0.2804
Experimental & Complex	0.3381	0.00307	0.6731
Experimental & Unnatural & Complex	-0.3127	-0.8055	0.1801

Random Effects (Mixed Effects Model)

Parameter	Estimate
Subject Number (Intercept)	0.1270
Stimuli (Intercept)	0.05384
Residual	0.3372

Table 6. Comparison of Results with Hayes and White (2013)

	Current Study			Hayes and White 2013		
Effect	**Estimate**	**Lower CL**	**Upper CL**	**Estimate**	**Lower CL**	**Upper CL**
Intercept	4.67	4.48	4.86	5	4.79	5.21
Experimental	-0.93	-1.17	-0.70	-1.33	-1.58	-1.09
Unnatural	-0.15	-0.41	0.10	-0.4	-0.65	-0.17
Experimental & Unnatural	0.61	0.24	0.97	1.13	0.8	.48

The above data was also analyzed to control for constraint weights (see Sections 2.1 and 4.1 for more information on the constraint weights). The patterns were the same in this analysis as the above figure. The weights did not significantly affect the overall group effects, probably because of the similarity in average weight that all of the groups shared.

5.2 Data for Individual Constraints

Table 7 shows each constraint's individual effect. No constraints are more than two standard deviations from their group's mean effect and most have mean effects that are fairly representative of their category. There are four, though, that show a larger amount of deviation than the rest. When excluded from the data set, these four do not change the general trend of the data (that is, the stronger effects of naturalness), but do make it more pronounced. Without these constraints, the mean complexity scores (discussed in Section 4.1.3) for each of the constraint categories remains representative of that category (1.38 and 1.42 for natural simple and unnatural simple, respectively; 1.89 and 1.94 for natural complex and unnatural complex, respectively).

The first of these deviating constraints is *[-voice][+voice] which had a much smaller effect than the rest of the Natural Simple constraints. This constraint also acted unpredictably for Hayes and White (having a smaller effect than the rest of their natural constraints), probably due to the fact that speakers have trouble hearing violations of it, and also because of the difficulty of expressing such a violation with English orthography. For instance, the word [dagz] *dogs* does not violate the above constraint when spoken aloud, even though the orthography suggests it ought to. The stimuli used in this experiment that violated this constraint (*esger* [ɛsgɹ] and *trocdal* [tʃɹakdəl]) could have been interpreted as orthographically ambiguous, like the word *dogs*, and could have then been given a high rating since they did not seem to violate any English constraint.

Another constraint with a surprising effect was *[+CORONAL,-rhyme][-back,-syllabic]. The effect of this group was basically zero. This could have been a result of the stimuli used. The experimental stimuli for this constraint were [djIn] *dyin* and [sjɛp] *syep* and the control stimuli were [dwIn] *dwin* and [vjɛp] *vyep*. The control stimuli did not violate any of the 160 constraints, and the onset clusters for the control words are attested in American English (for example, in the onsets of [dwɔɹf] *dwarf* and [vjuw] *view*). The problem with these stimuli was likely in their orthography—especially in the case of *vyep*. This irregular spelling could have caused participants to imagine these words as being inappropriate for English spelling and allowed this feeling of orthographic error to influence their grammaticality judgments, despite the fact that the words did not violate any phonological constraints.

Another group of constraints that behaved somewhat unexpectedly were the two constraints involving the high, back, lax vowel [ʊ]. These were *[+round,-tense][-sonorant,+LABIAL,-rhyme] and *[+round,-tense][+continuant,+voice,+rhyme]. These constraints deviated from their group's mean more than any of the other constraints in their group did. This complication likely arose from the difficulty in representing the vowel [ʊ] in English orthography and from difficulty that the author had in pronouncing [ʊ] in novel words when producing the auditory stimuli.

The remainder of the constraints were all much closer to their group's mean. Although these other constraints also showed variation within groups, it was much smaller than the constraints discussed above.

Table 7 - Individual Constraint Effects

Natural Simple	Effect:	Unnatural Simple	Effect:
*[-voice][+voice]	0.12	*[+diphthong][+continuant,-anterior]	0.42
*[+syllabic][+syllabic]	0.71	*[+round,-back][-anterior]	0.22
*[-consonantal,-rhyme][+consonantal,-rhyme]	1.11	*[+CORONAL,-rhyme][-back,-syllabic]	-0.01
*[-continuant,-rhyme][-continuant,-rhyme]	1.40	*[-sonorant,+rhyme][-low]	0.34
*[-sonorant,+rhyme][+sonorant,+rhyme]	0.46		
Natural Complex	Effect:	Unnatural Complex	Effect:
*[+consonantal,-anterior][-sonorant]	0.23	*[+round,+high][-consonantal,-sonorant]	0.21
*[+diphthong,+round][+LABIAL]	0.44	*[+diphthong,+round][+voice,+DORSAL]	0.39
*[-low,+tense][+sonorant,+DORSAL]	0.97	*[+consonantal,-anterior][+high,-syllabic]	0.29
*[+nasal,+CORONAL][+DORSAL,+rhyme]	0.74	*[+diphthong,+round][-strident,-rhyme]	0.29
*[+round,-tense][-sonorant,+LABIAL,-rhyme]	0.17	*[+round,-tense][+continuant,+voice,+rhyme]	1.09

6. Discussion

6.1 Reinforcement of Hayes and White (2013)

In Hayes and White (2013), a comparison was made between the effects of natural constraints and unnatural constraints. The primary conclusion drawn from their results was that learners of languages have an innate bias towards learning natural phonological constraints. In the current study, naturalness was again shown to have an effect. When presented with novel words that violated natural constraints, subjects rated the words as significantly worse sounding than when they were presented words that violated unnatural constraints. These lower ratings were interpreted, as in Hayes and White, as an indicator of the subjects having learned the natural constraints better than the unnatural ones when they first acquired English. This could mean, as Hayes and White suggest, that learners have a bias towards learning natural constraints.

6.2 Complexity's Effect?

Dimensional complexity seems to have had a fleeting effect in the ratings that subjects gave the stimuli. Its effect is present in the comparison between the Natural Complex and Natural Simple groups, but does not appear to have a significant effect in the comparison between the two unnatural categories. The mean complexity scores for the four categories (after the constraints discussed in Section 5.2 are removed) are 1.38 and 1.42 for natural simple and unnatural simple, respectively and 1.89 and 1.94 for natural complex and unnatural complex, respectively. This means the difference between the two natural categories is .51 and the difference between the unnatural categories is .52. With these scores being so similar, one would expect for complexity to have similar effects on both groups (or at least to have a slightly higher effect on the unnatural group) but instead the effect of complexity seems to have been negated for the unnatural constraints.

One possible explanation for this result is that naturalness had such a large effect on subjects' judgments that the unnaturalness of constraints eclipsed their complexity. Rather than distinguish between the unnatural simple and unnatural complex groups, the participants in the experiment could have given both groups low ratings indiscriminately. This would mean that dimensional complexity was still affecting judgment ratings, but only where naturalness did not already have an effect. Studies such as Saffran and Thiessen (2003) and Skoruppa and Peperkamp (2011) that show complexity's[10] effect on the learning of phonotactic patterns could be circumventing the issue of naturalness by using artificial languages. Since Hayes and White (2013) and the current study both tested speakers on patterns that were found within English, a real world language, the effects of naturalness could have taken a larger effect than with a language learned in a lab. This is discussed more in Section 6.4.

6.3 Future Directions for Research

6.3.1 Orthographic Constraints

As was mentioned briefly in Section 5.2, the orthography used in communicating some of the stimuli could have influenced participants' ratings. Specific rules that govern how phonemes are transformed into graphemes have been shown to exist (See, for example, Kreiner and Gough 1990). There is a possibility that when creating the orthographic representations for the stimuli, one of these rules of orthography could have been violated. If a participant were to see a word oddly spelled (for example "cping"), this could have caused their rating of the word to be lower which would have made the constraint that the word violated to seem more effective. In addition to this, orthography could have caused confusion about the pronunciation of some of the stimuli. The spelling of "esger", for instance, was meant to convey the word [ɛsgɹ] but could have been thought of as [ɛskɹ] since this is the pronunciation that would typically arise from that spelling. Perhaps a future experiment could test the effect of alternate spellings to determine if the orthographic representations had a significant effect on the participant ratings, or try to convey the words without orthography at all. An orthographic rule-learning software could also be used to create general guidelines for making well-spelled words. At the very least, software could be used to rate how "badly" a novel word is spelled (in terms of typical rules of orthography) and then use this numeric rating to control for the effect of orthography in the statistical analysis.

[10] These papers tested categorical complexity (as discussed in Section 2.3); however the basic principle of simpler constraints being easier to learn would still be expected to apply.

6.3.2 Better Constraint Complexity

Constraints within the set of 160 created by the Hayes and Wilson Phonotactic learner that were unnatural (See Table 7 for a list of all the constraints) had a tendency to also be more complex. Due to the finite number of constraints to choose from, the complex categories and simple categories did not differ in complexity more than half a point in their mean complexity scores. If there were a way to have a larger number of constraints (perhaps by making alterations to how the phonotactic learner created the constraints) to choose from, then perhaps this problem could be avoided in future research. Perhaps less of an emphasis on producing more general constraints could create more options for constraints to test in a perception experiment for naturalness, although this could also sacrifice some of the learner's value as an analogue to human learning.

6.3.4 Complexity as a Scalar Variable

The current study only looked at complexity as a binary variable. Constraints were considered to be either "simple" or "complex" without any steps in between the two groups. However, one could analyze the data with a statistical model that takes complexity into account as a continuum, rather than two distinct groups. The complexity scores discussed in Table 4 could be used to obtain a measurement for the effect of complexity that has a higher resolution than the current analysis provides. This could provide more insight on different the effects of complexity and naturalness are.

6.3.3 Artificial Language Studies

Experiments could also be performed with artificial language learning that involve infants learning and using an artificial phonology over time to see if, given time to "test drive", children would show a bias towards natural patterns in the lab. If infants were to develop these biases with artificial languages, then it could not only support and improve the use of artificial language learning for research, but also increase knowledge of the mind and learning as a whole.

Another artificial language study that could be performed would be one on adults that, after teaching the subjects a phonological pattern, asks them to make grammaticality judgments on words in a similar style to Hayes and White (2013). This could work out any issues that this particular methodology might be causing in the data.

It could also be beneficial to perform an experiment on dimensional complexity's effect in artificial language learning. There is a possibility that in phonological learning categorical complexity (which has been thoroughly studied) and dimensional complexity (which has not been looked at in the context of artificial language studies but was used in this study) have different effects on learning.

6.4 Conclusions

The data in the current study suggests that naturalness has a very strong effect on speaker judgments while dimensional complexity has a much smaller one. This seems to contradict many recent experiments using artificial languages (such as Skoruppa and Peperkamp 2011). These artificial language studies have not only found evidence for a large effect of categorical complexity, but they have failed to find a substantial effect of naturalness (for a review of these studies, see Moreton and Pater 2012a and b). This could be a result of an innate bias within speakers for certain types of phonotactic patterns, as suggested by Becker et al. (2008). Such an innate bias would explain the results in the current study, but would not necessarily explain why artificial language experiments are finding different results. Another possibility is that the human mind is tailored to only learn natural constraints on a linguistic level, with unnatural constraints only being learned cognitively and later in life.

Hayes (1999) presented the approach, called *inductive grounding*, that infants could be born with the knowledge of how to derive natural constraints. Children, Hayes said, experiment with the sounds and sequences of sounds that occur in the language they hear being spoken around them. They then calculate which of the sound sequences require more or less effort by articulating the sounds themselves. After finishing this test drive of their vocal tract, the children create a set of phonological constraints that strike a balance between being efficient on a phonetic level and being logically simple. The children then proceed to rank these constraints according to observations about how the adults around them are speaking until the children have constructed a grammar identical to that of the adults they were learning from.

This theory allows children to construct grammars based on information that is readily available to them—such as their own physiology. Since human beings have a generally universal physiology, the constraints that arise cross-linguistically would be basically the same; with the constraint ordering being what differed between languages, as in classical Optimality Theory (Prince and Smolensky 1993). The Hayes (1999) approach could account for speakers' preference of natural constraints, since these would be the original constraints people learn as children. Perhaps unnatural constraints are learned on a subconscious level as well as people grow older, but these would not be internalized into the grammar and would not affect grammaticality judgments as strongly.

So, for example, a constraint such as *[+voice][-voice] (voicing assimilation) would be learned by a child acquiring their native language. This constraint would not be learned by observing the adults speaking around them, but instead by producing consonant clusters to determine what the easiest and simplest way to create the clusters is. After discovering that voice assimilation assists with consonant clusters (and a number of other pieces of phonetic insight) the child would rank the constraints in a way that corresponded to the phonology of the language they were learning (so in English, for instance, the child would rank voicing assimilation very high, since this is very common in adult English phonology).

A less natural constraint, like [*+round,+diphthong][g] (no [aƱ] or [oi] before [g]), although seemingly present in English phonotactics, would not be learned by children at first, since it cannot be directly derived from phonetics. Eventually, as the child grows older, he or she might become accustom to the fact that round diphthongs do not often come before [g] in English (which could somewhat affect grammaticality ratings of words that violate this tendency) but they would never internalize this constraint into their grammar because it was not part of the original constraint list that the child derived from inductive phonetic grounding.

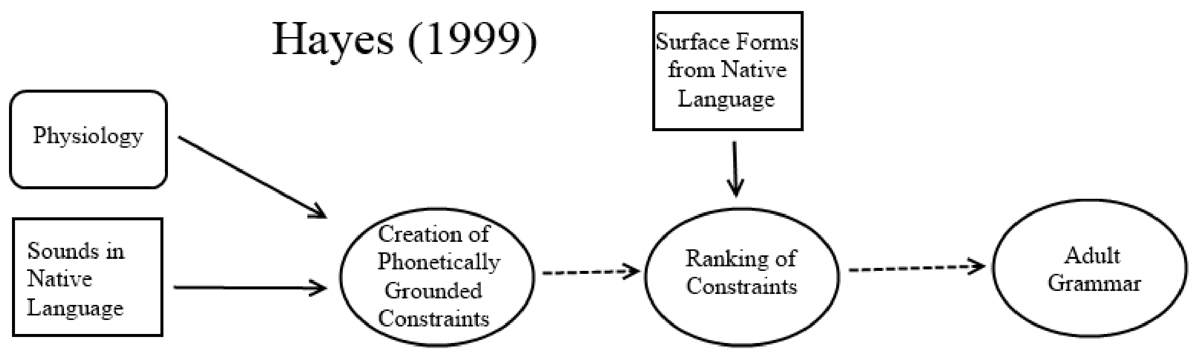

Figure 7 - An illustration of the process that infants could be going through to learn phonological constraints. By first constructing constraints that are phonetically grounded and structurally logical, children develop a preference for phonetically natural constraints that are also simple. These constraints are then ordered based on information that children hear in the language around them until the child has created a completely adult grammar.

This idea could also explain why experiments performed with artificial languages do not see an effect of naturalness on phonological learning. If the first step in learning a language's phonology is creating a set of personal phonological constraints, then this constraint creation and ranking could be a one-time process. Because subjects learning an artificial phonology are attempting to learn a new set of constraints (or re-rank old ones), a bias towards simple constraints could arise since these would be easier to learn on a purely cognitive level (as opposed to the kind of specialized learning that would happen as a young child). Since this learning would be different than natural linguistic learning, it would act more like the learning in studies such as Bulgarella and Archer (1962) and Shepard et al. (1961). Infants being tested in laboratory situations might also show a preference for simple rules over natural ones, but simply because they are not given the time to "test drive" the rules being presented to them.

In the learning of second language phonology (for real-world languages), a bias towards simplicity would again be predicted. Since a person's constraints would already be ranked, they would have to learn the constraint ranking for this new language on a more conscious level, which would likely create a stronger bias for simplicity. In addition to this, there would likely be certain aspects of the second language phonology that the individual would never be able to learn. Since, in this approach, the constraints in each person's grammar are created for the sounds in their native language, some constraints could be useless or missing for the second language the person is trying to learn, resulting in a permanently imperfect second language phonology.

Future work could look at creating a computational model like the Hayes and Wilson Phonotactic Learner that first creates constraints based on phonetic inductive grounding, then assigns each of the constraints a weight based on data from a language's surface forms. This would require an exhaustive model of human phonetics (such as the small example given on voicing in Hayes 1999) but would provide insight into the plausibility of this conclusion.

7. Appendix

Constraints and Their Corresponding Stimuli

Featural Constraints	***Stimuli***			
Natural Complex	**Violater 1**	**Control 1**	**Violater 2**	**Control 2**
***[+consonantal,-anterior][-sonorant]**	ishty	ishmy	metchter	metchner
***[+diphthong,+round][+LABIAL]**	frowp	frope	soib	soid
***[-low,+tense][+sonorant,+DORSAL]**	perng	pern	plieng	pline
***[+nasal,+CORONAL][+DORSAL,+rhyme]**	quinnk	quisk	zannk	zask
***[+round,-tense][-sonorant,+LABIAL,-rhyme]**	truhbin	troobin	yuhpase	yuhtase
Unnatural Complex				
***[+round,+high][-consonantal,-sonorant]**	pauhin	pausin	sloohite	sleehite
***[+diphthong,+round][+voice,+DORSAL]**	moig	moid	towg	towk
***[+consonantal,-anterior][+high,-syllabic]**	shween	sween	chwid	chrid
***[+diphthong,+round][-strident,-rhyme]**	boithin	biethin	owthat	owtat
***[+round,-tense][+continuant,+voice,+rhyme]**	wuhth	wooth	tuhz	tuss
Natural Simple				
***[-voice][+voice]**	esger	ezger	trocdal	troctal
***[+syllabic][+syllabic]**	keeane	klane	biate	brate
***[-consonantal,-rhyme][+consonantal,-rhyme]**	hloop	ploop	hmit	smit
***[-continuant,-rhyme][-continuant,-rhyme]**	cping	sping	ctice	stice
***[-sonorant,+rhyme][+sonorant,+rhyme]**	canifl	canift	kipl	kilp
Unnatural Simple				
***[+diphthong][+continuant,-anterior]**	bowsh	boosh	coish	kish
***[+round,-back][-anterior]**	coitch	keetch	goige	goik
***[+CORONAL,-rhyme][-back,-syllabic]**	dyin	dwin	syep	vyep
***[-sonorant,+rhyme][-low]**	pabyin	payblin	bidyon	binyon

8. References

Baayen, R. Harald, Richard Piepenbrock, and Rijn van H. "The {CELEX} lexical data base on {CD-ROM}." (1993).

Becker, Michael, Nihan Ketrez, and Andrew Nevins. "The surfeit of the stimulus: Analytic biases filter lexical statistics in Turkish devoicing neutralization." Web publication http://roa. rutgers. edu/view. php3 (2008).

Bulgarella, Rosaria, and E. James Archer. "Concept identification of auditory stimuli as a function of amount of relevant and irrelevant information." Journal of Experimental Psychology 63.3 (1962): 254.

Casali, Roderic F. "Vowel elision in hiatus contexts: Which vowel goes?" Language (1997): 493-533.

Chomsky, Noam, and Morris Halle. *The Sound Pattern of English*. New York: Harper & Row, 1968.

Hayes, Bruce. "Phonetically Driven Phonology." Functionalism and formalism in linguistics 1 (1999): 243-285.

Hayes, Bruce, and Colin Wilson. "A Maximum Entropy Model of Phonotactics and Phonotactic Learning." *Linguistic Inquiry* 39.3 (2008): 379-440.

Hayes, Bruce, and James White. "Phonological Naturalness and Phonotactic Learning." *Linguistic Inquiry* 44.1 (2013): 45-75.

Kreiner, David S., and Philip B. Gough. "Two ideas about spelling: Rules and word-specific memory." Journal of Memory and Language 29.1 (1990): 103-118.

Mielke, Jeff. "P-base." A database of sound patterns in 500 (2007).

Moreton, Elliott, and Joe Pater. "Structure and Substance in Artificial-phonology Learning, Part I: Structure." *Language and Linguistics Compass* 6.11 (2012): 686-701.

Moreton, Elliott, and Joe Pater. "Structure and Substance in Artificial-phonology Learning, Part II: Substance." *Language and Linguistics Compass* 6.11 (2012): 702--18.

Prince, Alan, and Paul Smolensky. "Optimality Theory in phonology." International encyclopedia of linguistics (2003).

Pycha, A., P. Nowak, E. Shin, and R. Shosted (2003). "Phonological rule-learing and its implications for a theory of vowel harmony." In M. Tsujimura and G. Garding (Eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)*, pp. 101-114.

Saffran, Jenny R., and Erik D. Thiessen. "Pattern induction by infant language learners." *Developmental psychology* 39.3 (2003): 484.

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

Shepard, R. N., C. L. Hovland, and H. M. Jenkins (1961). "Learning and memorization of classifications." *Psychological Monographs 75* (13, Whole No. 517).

Skoruppa, K. and S. Peperkamp (2011). "Adaptation to novel accents: feature-based learning of context-sensitive phonological regularities. *Cognitive Science 35,* 348-366.