From Phonetics to Phonology: Learning Epenthesis

Rebecca L. Morley
The Ohio State University

1 Introduction

This work investigates the question of how phonological systems emerge historically, with particular emphasis on the contribution of the language learner. A theoretical account of synchronic phonological patterns as the product of phonetically-based sound changes (e.g. Blevins 2004) is applied to a grammar of morphologically conditioned consonant epenthesis. Epenthetic material is hypothesized to emerge from listener misperception of the natural transition between two adjacent vowels: ratu+ək, pronounced as ratu^{w}ək, re-analyzed as ratuwək. This diachronic source entails that epenthesis will be, at least at first, contextually conditioned by the features of the adjacent vowels (what I will call Type 1). A distinction has been made, however, between such systems and those in which a unique segment is epenthesized regardless of context (Type 2) (see, e.g., Lombardi 2002, de Lacy 2006).

For each type of synchronic outcome a set of sufficient diachronic conditions is established. This is done through experiments that test how particular permutations of phonetic, phonological, and morphological information affect language learning. The results shed light on the largely unknown mechanisms by which lexical sound change is transformed across speakers into grammatical language change. In turn, a more explicit model of natural diachronic processes leads to tighter predictions regarding the typological distribution of natural synchronic languages.

2 The Model

The apparent universal dispreference for onsetless syllables is potentially an emergent consequence of processes that erode features of one or both of two immediately adjacent vowels. As is well known, naturally produced speech involves ordering the articulatory gestures for adjacent segments such that they overlap in time (e.g., Browman and Goldstein 1986, Byrd and Saltzman 1998, Zsiga 2000). Depending on the features of those adjacent segments, and the degree of overlap, varying acoustic and auditory outcomes result. Each arrow in Fig. 1 describes a hypothesized historic path arising from a different timing relationship between articulatory gestures for two vowels adjacent across a morpheme boundary (cf. Casali 1997). Of these, one route traces out a possible trajectory for the two types of epenthesis systems described in the introduction to this paper.

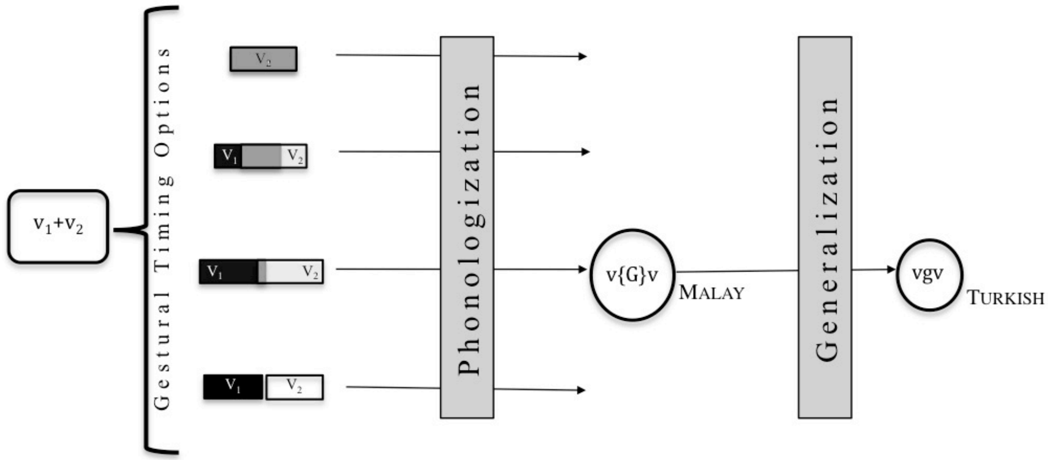

Figure 1 : An incremental learning model of the emergence of different grammars arising from historic vowel hiatus (v+v). The overlapping rectangles represent different timing options for the adjacent sets of vowel gestures. Each circle represents a possible language. v{G}v: Type 1 phonological epenthesis. vgv: Type 2 phonological epenthesis.

With respect to the model, a change in the synchronic grammar is defined as a disparity between the analysis of the speaker and the analysis of the listener. For the speaker, the gestural coordination relations are the consequences of production. If the listener interprets those consequences as the goals of production instead, then a change in the underlying representations has taken place (see, e.g., McMahon et al. 1994), and the phonetic pattern (excrescent glide) becomes phonologized (segmental glide) (Hyman 1976). An epenthetic pattern, like that in Malay (1), in which the segment is conditioned by phonetic context is taken as a possible output of the Phonologization mechanism.

(1)	/bantu + an/	→[bantuwan]	"relief"
	/udʒi + an/	→[udʒijan]	"test"
	/məŋ + gula + i/	→[məŋgulaʔi]	"cause to sweeten"

There is an additional mechanism, whereby the phonologized pattern becomes regularized to the point of context-independence. Incrementality within the model refers to the fact that Type 2 default epenthesis languages (such as Turkish (2)) can only arise from Type 1 systems (like Malay). Thus, they must pass through two stages: re-analysis of gradient information as categorical ('Phonologization'), and abstraction over a class of segments and words ('Generalization').

(2)	/iste + en/	→[istijen]	"one who wants"
	/anla + ɪʃ/	→[anlajɨʃ]	"understanding"
	/su + ɪ/	→[suju]	"his/her water"

3 The Experimental Design

The current work is intended as a contribution towards developing an explanatorily adequate model of both the mechanisms of ‘Phonologization’ and ‘Generalization’. This is done by investigating some of the characteristics of language learners via a series of artificial grammar learning experiments (Saffran & Thiessen 2003, Newport & Aslin 2004, Gerken 2006, Wilson 2006, Kapatsinski 2010, and many others). This paradigm allows for the learner’s input to be carefully controlled and manipulated. Additionally, the learning problem can be simplified by limiting exposure to small sets of words, and focusing the learner’s attention on the linguistic alternation of interest. Participants have been shown to learn novel words and patterns quickly within this paradigm, and to implicitly acquire grammatical rules of which they are consciously unaware.

There are three experiments described in this paper. The first two are focused on the ‘Phonologization’ mechanism in Fig. 1, and investigate the interface between phonetics and phonology. The third experiment is explicitly designed to probe the ‘Generalization’ stage of the model by manipulating the relative frequency with which certain morphemes occur, and reliability of their contextual predictors.

3.1 Procedure

A total of 101 participants were run in the three experiments of this paper (21 in Experiment 3, 20 in all other experimental conditions). All were undergraduates participating for course credit at the Ohio State University. Participants were told that they would be hearing words in a new language, and that they would later be asked questions about those words. What followed was a passive training stage in which participants listened to words over the headphones and looked at pictures on the computer monitor. The words occurred in doubles: the singular (e.g., ratu), followed by the plural (e.g., ratuwək). A picture of a single object (e.g. an apple) accompanied playback of the singular word, and a picture of two objects (e.g., two apples) accompanied playback of the plural word. See Fig. 2.

Training		Test	
[ˈratu]	[ˈratuwək]	[ˈdarum]	???

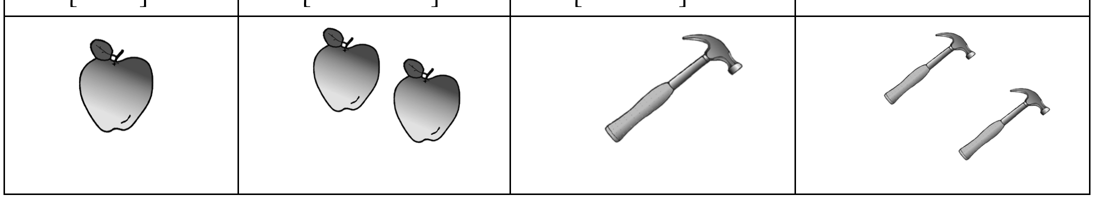

Figure 2: Schematic of training and test conditions for Experiments 1-3.

Participants were trained on 12 or 18 different singular-plural word pairs, depending on experimental condition. This set of words was repeated in 4 randomized blocks. A feedback stage occurred halfway through training, and again after training was complete. This stage provided a singular form (heard during training) and prompted the participant for the plural: "Now you say the plural…" Participants were instructed to speak their response into the microphone. After an interval of 5 seconds the words "The correct answer is…" appeared on the screen and the computer played the plural. The procedure was the same at test, except the correct answer was not provided; additionally participants were exposed to singular forms that they had never heard before. There were 18 test words, repeated twice, in randomized order, for a total of 36 test items. All participants in a given condition heard the same set of words, associated with the same pictures; the order of presentation, however, was randomized. The entire experiment took roughly half an hour.

3.2 Stimuli

The three experiments each utilized a subset of a common pool of auditory stimuli. All stimuli were spoken by a phonetically trained female speaker of American English. Singular and plural words were recorded separately. The full set of singular forms (stems) is given in Table 1. All plural forms were of the form singular + Xək. Depending on experiment and condition X was either a glide homorganic to the place of the preceding vowel, an anti-homorganic glide, or a pause/discontinuity (e.g., skibejək, skibewək, skibe.ək).

A	ˈratu	ˈrilo	fraˈbomu	tʃoˈræno	kro	vu
B	ˈhædi	ˈskibe	teˈlʌpi	gluˈdɛbe	fi	sme
C	ˈpishu	ˈhago	bəˈhʌʒu	ˈfædʒo	ˈzo	ˈgaidu
D	ˈvʌlki	ˈploke	ˈdʒimi	diˈʒare	ˈθuzi	ˈʃuvi
E	ˈdarum	keˈtɛlan	ˈhoʃɪn	ˈribæz	ˈprɛv	ˈbihɪl

Table 1: Full set of singular stimuli used in all experiments (training and test)

The last type of stimuli were necessary for distinguishing between a phonetic versus phonological hypothesis on the part of listeners (see discussion of Experiment 1). It is difficult, if not impossible, to avoid producing some sort of glide-like transition between adjacent vowels of a certain type. Conversely, distinguishing between glides in onset position in unstressed syllables and onsetless vowels is difficult. Productions of [ratuwək] and [ratuək] by the speaker of the stimuli were auditorily and spectrographically highly similar. For this reason, tokens were spliced to avoid emergent glides.

For example, the token [skibe.ək] was created by recording a single utterance with a long pause: [skibe ək], then splicing out the pause, as well as the very beginning of the final vowel, to minimize glottalization. The resulting stimuli exhibited a discontinuity in the spectrogram, eliminating the natural transitional period between adjacent vowels, and as a result were rather unnatural sounding; they were, however, clearly distinguishable from the full glide variants. Participants developed various strategies for reproducing these unnatural words: altering vowel quality, introducing glottal stops, or significantly drawing out their articulations. Furthermore, the experimental results show that such tokens were natural-enough sounding to be categorized by listeners as reflexes of underlying vowel-vowel sequences.

Participant responses were coded by a phonetically trained listener as either being consistent with a phonetically natural production, or not. That is, since finer measurements are needed to definitively establish whether listeners intended to produce /ratuwək/ or /ratuək/, the two are not differentiated for the purpose of the analyses presented here. For test items consisting of consonant-final stems, however, the presence versus absence of a glide is quite clear (e.g., /darumwək/ versus /darumək/).

4 Experiment 1: Phonetic Primacy

Experiment 1 consisted of three conditions: the Natural, Anti-Natural, and Bi-Modal. All stems heard in training were vowel final, either ending in o or u (back vowels), or i or e (front vowels): rows A and B of Table 1. The critical items at test were novel consonant-final stems: Row E. In the Natural condition, all plurals contained homorganic glides. In the Anti-Natural condition, all plurals contained anti-homorganic glides. The surface forms for both of these conditions can be analyzed as the result of an epenthetic operation. See (3a): Natural, and (3b): Anti-Natural. The linguistic information available to the participants, however, is impoverished (in the sense of Wilson 2006); since they never hear inflected consonant-final stems, the underlying form of the plural suffix is ambiguous.

(3a) skibe + ək→ skibejək ratu + ək→ratuwək

(3b) skibe + ək→ skibewək ratu + ək→ratujək

There are only two types of segments that occur at the end of the inflected stem; which one appears is completely predictable, and the conditioning context is part of the natural feature set (i.e. vowel height). These three factors are taken to define phonological naturalness in the experimental context, and predict that learners will interpret the glide as epenthetic for both conditions. There are, however, two types of 'naturalness' at play here; the names of the conditions refer to the phonetic naturalness of the glide in this context. The 'Natural' condition represents an assimilatory alternation, a phonetically natural reflex of coarticulation. The 'Anti-Natural' condition represents a dissimilatory alternation, one that could arise through hypo-articulation, something that would presumably only occur in contexts in which confusion due to neutralization might result.

Suffixation of consonant-final forms at test reveals which hypothesis participants have adopted. The phonological naturalness hypothesis (3) will lead to forms like (4) as responses for both conditions.

(4) darum + ək → darumək

Fig. 3 shows that this is true of the Natural condition in which participants overwhelmingly produced affixed forms which lacked glides. However, in the Anti-Natural condition participants produced the glide at test (e.g., [darumwək] or [darumjək]). This was in spite of the fact that they had never before heard any CC sequences (all training syllables were of the form CV).

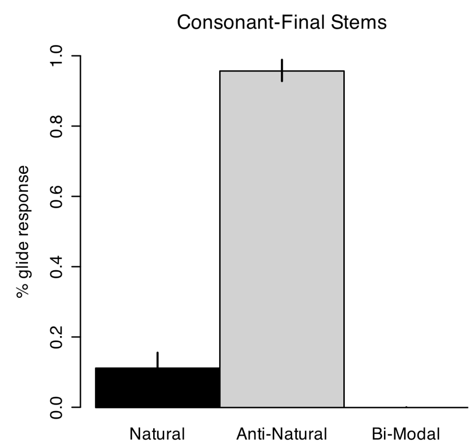

Figure 3: Results of Experiment 1 by Condition; Consonant-final stem test items only.

The contrasting results for the two conditions argues against the phonological naturalness hypothesis. The Anti-Natural condition results are, however, consistent with an account in which participants learned two predictable allomorphs for the plural suffix: -wək and -jək. The fact that those suffixes were applied to novel consonant-final stems indicates that implicit phonotactic information regarding preferred syllable shape was not enough to overcome participants' faithfulness to learned morphemic representations. If this analysis is correct, then why the two-allomorph representation did not occur for the Natural condition must be explained. The two conditions can be reconciled under the following hypothesis. Rather than posit an epenthetic segment, participants in the Natural condition both perceived and encoded a single plural suffix, attributing all segmental glide material to expected coarticulatory phonetics. This hypothesis was tested in the Bi-Modal condition.

The Bi-Modal condition mixed phonetically natural tokens (identical with Natural condition stimuli) with spliced tokens from which all traces of excrescent glides were removed, resulting in V-V sequences. Which tokens occurred with the spliced plurals depended on stem, and was not predictable from phonetic context. The reasoning was that the spliced tokens would, by comparison, force a segmental interpretation for the naturally produced tokens. This would result in participants' learning three (partially predictable) plural allomorphs: -wək, - jək, and -ək. A sample of representative training forms is given in (5).

(5) skibe/skibejək ratu/ratuwək
hædi/hædiək hago/hagoək

Fig. 1 shows that it was clearly not the case that participants learned three allomorphs in the Bi-Modal condition. Participants behaved in this condition as they had in the Natural condition – producing only consonant-final affixed forms lacking glides at test. Statistically, the two conditions do not differ from one another, but both differ significantly from the Anti-Natural condition ($p < .05$).

The full set of results from Experiment 1 is interpreted to indicate, firstly, a primacy of phonetic naturalness over phonological. While phonotactic information was implicitly available during training, listener response did not always conform to it. This is evident in the Anti-Natural condition, where participants seem to have learned two glide-initial allomorphs for the plural suffix. The phonetic naturalness of the training items, on the other hand, strongly predicted participant response. In both the Natural and the Bi-Modal conditions participants responded consistently as though they had learned a single vowel-initial form for the plural, one with phonetic variants[1].

The auditorily distinct tokens in the Bi-Modal condition, rather than inducing learners to encode three allomorphs, produced a phonologically unimodal response distribution. Listeners appear to treat ..Vək and ..VGək plurals as phonetic tokens of the same morpheme class. This is plausible under a model in which listeners have phonetic expectations based on their native language competence (e.g., Beddor & Krakow 1986), as well as a predilection for limiting the surface realizations of semantically identical morphemes. Specifically, the fact that these two acoustically quite distinct phone sequences could be classified together is attributable to speaker knowledge that carefully articulated vowel-vowel sequences result in an intrusive pause or glottal stop, whereas rapid or colloquial speech often produces phonetic gliding between two vowels. This interpretation accounts for the results of the Natural condition (in which listeners interpret all glides as phonetic). It is also consistent with the Anti-Natural condition, where even a bias towards a single-morpheme solution would not be

[1] A separate orthographic test supports the interpretation that participants treated glides heard in training for the Natural condition as the phonetic product of coarticulation.

enough when faced with two tokens that were impossible phonetic variants of the same target (..ujək and ..iwək).

5 Experiment 2: Morphological Economy

The results of Experiment 1 present a potential hurdle for the model of epenthesis emergence outlined in Section 2. For naturally produced glides to become epenthetic segments learners must be forced to interpret them as such; they must overcome an apparently quite strong bias towards their prior phonetic expectation. Experiment 2 (Consistent 'w') shows that a phonemic analysis of the expected glide can be induced via morphological decomposition.

In addition to half of the phonetically natural tokens of Experiment 1, participants heard consonant-final stems with a clear /w/; training consisted of rows A and E of Table 1. Front-vowel final stems were held out. Example training items are given in (6).

(6) darum/darumwək ratu/ratuwək

Under the phonetic naturalness hypothesis attributed to the results of Experiment 1, learners should interpret the glide in inflected back-vowel final stems as phonetic. This interpretation is not possible for the consonant-final stems. Therefore, a two-allomorph analysis is predicted: -wək, and -ək. Critical items at test are the novel front-vowel final words, where the presence versus absence of a [w] in the plural suffix could be clearly heard.

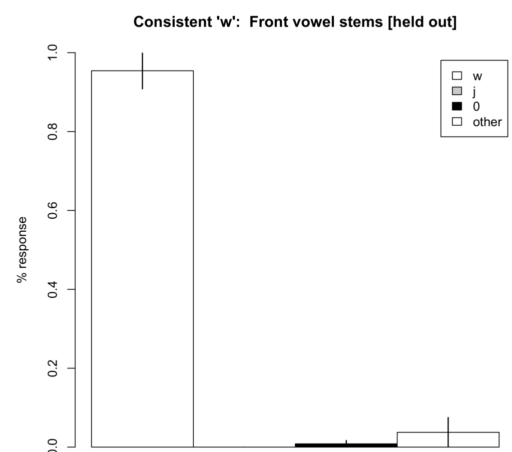

Figure 4: Results of Experiment 2 (Consistent 'w'), by response type.

The results in Fig. 4 show that participants responded almost categorically with the morpheme -wək, e.g., [plokewək]. Instead of an equiprobable two-allomorph distribution, participants clearly learned a single plural surface form. This outcome is taken to indicate the strong effect of a bias against multiple allomorphs. Additionally, this shows that the segmental interpretation is, in fact, accessible. Listeners have been trained, potentially both by general phonetic considerations, as well as their own language-specific experience, to automatically compensate for the glide-like acoustics that emerge predictably between vowel articulations of certain kinds. The bias towards phonetic naturalness, while strong enough to trump phonological and acoustic factors, can be suppressed. Morphological economy over-rides phonetic predictability. This is an important result as it provides a clear mechanism for initiating the model stage of Phonologization.

6 Experiment 3: Within-Category Variability

Historical sound changes are often modeled as diffusion of already existing variation (e.g., Kroch 1989, Niyogi 1997). Another potential source of variation, however, is that which can result directly from sound change: variation within a single speaker. Take the example of historic consonant loss in Maori. Bare verb stems underwent this change in pre-Polynesian, e.g., [hopuk] > [hopu]. Passive forms, however, were unaffected: [hopukia] > [hopukia] (Hale 1973). A sample of the surface alternations available to modern day Maori learners is given in Table 2.

verb	passive	gloss
awhi	awhitia	"to embrace"
hopu	hopukia	"to catch"
aru	arumia	"to follow"
tohu	tohuŋia	"to point out"
patu	patua	"to strike, kill"
kite	kitea	"to see, find"

Table 2: Stem/Passive alternations in Maori (Hale 1973)

These data resemble the impoverished alternations of Experiment 1. Since there are no consonant-final words in present-day Maori, learners are never exposed to those forms. Instead they see data that could support at least three distinct synchronic linguistic analyses: (7a) deletion of stem-final consonant, (7b) epenthesis of unpredictable segments, (7c) unpredictable allomorphy.

(7a) /hopuk/→[hopu] /hopuk + ia/ →[hopukia]

(7b) /hopu/ →[hopu] /hopu + k + ia/→[hopukia]

(7c) /hopu/ →[hopu] /hopu + kia/ →[hopukia]

Hale argues against analysis (7a) based on the intuition that learners would be unlikely to posit consonant-final underlying forms in a language environment in which no consonant-final words ever surface. Hypothesis (7b) may be similarly unlikely, based in part on the results of Experiment 1 which argue against epenthesis as a readily accessible learner hypothesis. That leaves hypothesis (7c), the type of allomorphy that was demonstrated, to a degree, in the Anti-Natural condition of Experiment 1. As can be seen from Table 2, however, the number of unpredictable allomorphs is quite large, posing a potential memory problem for the learner. In fact, Hale argues that uncommon verbs tend to lose their historic passive consonants, becoming 'regularized' to the same default suffix allomorph.

Within the field of creolization, it has been argued that learners will regularize their input in similar ways. Essentially, given a system with large amounts of variability (what amounts to widespread allomorphy, or suppletion), there is a threshold at which the cognitive system begins to introduce regularity. This could be attributed to a simple overloading of memory capacity. When this happens, the most frequent pattern or element will be 'boosted' such that it becomes the only pattern or element (see Singleton 1989, Ross and Newport 1996, Senghas and Coppola 2001, Hudson Cam & Newport 2005).

Experiment 3 was designed to test the effect of variability on learning, and to localize this proposed irregularity threshold. The example stimuli in (8) demonstrate the non-uniform distribution and partially predictable pattern of the training data.

(8)	2	skibe/skibewək	4	ratu/ratuək
	1	hædi/hædiək	2	hago/hagojək

Two different stem types (front vowel/back vowel) appeared with two different allomorphs each; one majority, one minority. Of the three possible allomorphs, one (-jək) appeared only with back-vowel stems; one (-wək) appeared only with front-vowel stems; and one (-ək) appeared with both: as the majority allomorph for back-vowel stems, and as the minority allomorph for front-vowel stems. The over-all ratio of types is given to the left of each alternation type in (8). Back-vowel final stems occurred twice as frequently as front-vowel final stems. Consonant-final stems were held out.

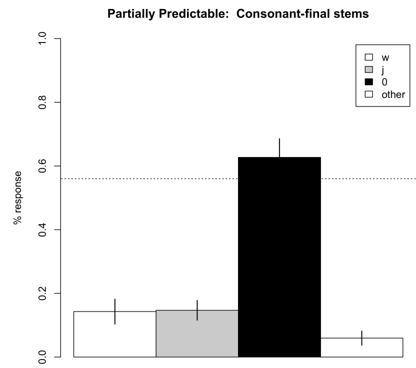

Figure 5a: Consonant-final stems; by response type.

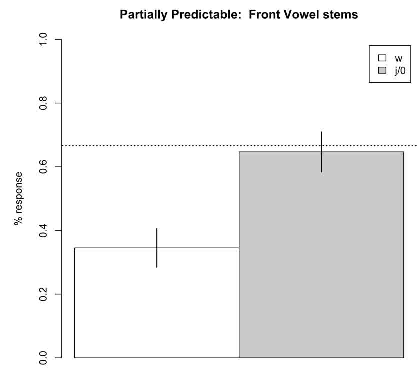

Figure 5b: Novel front-vowel final stems; by response type

On average, participants matched fairly closely the over-all frequency of the training items for the novel consonant-final stems heard at test. That is, the -wək and -jək allomorphs each accounted for a bit less than 22% of the responses; -ək for somewhat more than 56% (dotted line). On average, the proportion of -ək responses is about 12% higher than expected (excluding "other" responses). See Fig. 5a. Interestingly, however, for a stem class that was familiar, learners showed less faithfulness to their training distribution. The testing phase also consisted of a set of front-vowel stems that had not been heard during training (e.g., Row D in Table 1). Fig. 5b shows that for these stems, the number of -wək responses is about 30% less than expected (dotted line at 66%). This result seems to indicate an effect of the larger allomorph category which has been 'boosted' beyond its actual rate of occurrence[2]. A paired Wilcoxon sign test comparing the percent difference from expected for the two types of test item finds a significantly larger boosting effect for front-vowel final stems ($p<.05$ for a one-tailed test of the no difference null hypothesis).

On the one hand, it is not surprising that learners can, and will, keep track of a high degree of variability or detail (see, e.g., Goldinger 1996). Additionally, when there is predictability, e.g., across vowel category boundaries, one might expect representational stability. The interesting result, therefore, is exactly where and when a breakdown in memory and/or association can be induced. The -wək allomorph occurs more frequently with front vowel stems, but the -ək allomorph occurs more frequently over all. It is impossible to predict which of the front-vowel stems will take which allomorph. In the presence of the larger category of

[2] For the purposes of the measurements here, segmental and phonetic glides were grouped: "0/j" in Fig. 5b.

back-vowel stems, a type of 'attraction' to the more frequent allomorph seems to be at work. Listeners may actually think they have heard the novel stems before, and be 'misremembering' with a bias towards the larger category (cf. Wilson 2003). Consonant-final stems may be sufficiently novel, however, for listeners to treat them differently, and, in effect, more conservatively. This may result in part because the consonant/vowel distinction is much stronger cognitively than the front vowel/back vowel distinction. Additional experiments are needed to fully explore these results.

7 Discussion & Conclusions

The impoverished stimuli of the Natural condition of Experiment 1 are ecologically valid reflexes of natural coarticulation in a language like present day Maori that lacks consonant-final words. The results of the three experiments presented here, however, suggest that such initial conditions may be unlikely to lead to a synchronic epenthesis grammar. Learners consistently interpret what they hear such that the original underlying forms are preserved. Something further is needed to initiate the perceptual shift from phonetics to phonology. Experiment 2 provides evidence for one type of information that can strongly affect categorization on the part of the learner – the cohort of inflected forms. Additionally, Experiment 3 provides some indication of the role variability may play in introducing regularity to the non-conservative products of sound change.

If generalization based on phonological or morphological information must contend with a large phonetic bias, then certain limits can be placed on the likelihood of particular outcomes. The trajectory shown in Fig. 1 of a natural Type 2 epenthesis pattern resulting from a Type 1 pattern, for example, is predicted to have low probability (*a la* Experiment 1). This contrasts with the generalization that might result due to a regular pattern that was not phonetically natural, exemplified by the stimuli of Experiment 3, the Maori passive alternations, as well as other systems arising from historical consonant deletion (Blevins 2008).

This work provides a bridge from the apparent limitlessness of an emergentist, diachronic perspective to the apparently highly constrained space of possible synchronic grammars. The present research methodology discovers non-UG, and non-substantive, restrictions on language trajectories. In the first place, the experiments identify certain cognitive predispositions (Moreton (2008)'s 'analytic bias'): the weight given to phonetic expectation and morphological economy, and the permeability of certain linguistic categories in the face of variability. Additionally, one can identify a third factor that is properly attributable neither to the transmission channel, nor to the language acquirer. This might be termed an *environmental bias*: the types of words and morphological paradigms learners must be exposed to as a function of their current synchronic state. This factor has the potential to create quite narrow bottlenecks in 'grammaticalization' state space, and is thus a critical consideration in the generativist/emergentist debate.

Further experiments currently underway will pin down further the learning conditions which are most determinative in producing one synchronic state over another. Simultaneously, this will allow us to identify more narrowly the types of initial-state grammars that are potential precursor epenthesis systems (see Morley *under review)*. The entire research program provides a finely-grained predictive tool, both typologically and theoretically. This applies not only to the particular case study of consonant epenthesis, but to trajectories of any synchronic phonological pattern for which a plausible historic source can be constructed. As such, the methodology is a promising one for the investigation of linguistic universals.

Acknowledgments

This work was supported by the Targeted Investment in Excellence from the OSU College of Humanities. I would like to thank Paul Smolensky and Colin Wilson for their valuable insights on this project. Thank you also to Christina Heaton, MarDez Desmond, Lark Hovey and Sara Pennington for their work running participants. And thank you to Mary Beckman and the rest of the Phonies group at OSU for their input and advice.

References

Beddor, P. S., R. A. Krakow, et al. 1986. Perceptual constraints and phonological change: a study of nasal vowel height. In *Phonology Yearbook 3*, ed. By C. Ewen & J. Anderson, 197-217. Cambridge: Cambridge University Press.

Blevins, J. 2004. *Evolutionary Phonology: the emergence of sound patterns*. New York: Cambridge University Press.

Blevins, J. 2008. Consonant Epenthesis: Natural and Unnatural Histories. In *Linguistic Universals and Language Change*, ed. By J. Good, 79-106. Oxford University Press.

Browman, C. P. & L. M. Goldstein. 1986. Towards an articulatory phonology. *Phonology* 3. 219-252.

Byrd, D. & E. Saltzman. 1998. Intragestural dynamics of multiple prosodic boundaries. *Journal of Phonetics*. 26. 173-199.

Casali, R. F. 1997. Vowel elision in hiatus contexts: which vowel goes? *Language* 73(3). 493-533.

de Lacy, P. 2006. *Markedness: Reduction and Preservation in Phonology.* Cambridge University Press.

Gerken, L. 2006. Decisions, decisions: infant language learning when multiple generalizations are possible. *Cognition.* 98(3). B67-B74.

Goldinger, S. D. 1996. Words and voices: episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 22. 1166-1183.

Hale, K. 1973. Deep-surface canonical disparities in relation to analysis and change: an Australian example. *Diachronic, Areal and Typological Linguistics.* 401-458.

Hudson Kam, C. L. & E. L. Newport. 2005. Regularizing unpredictable variation: the roles of adult and child learners in language formation and change. *Language Learning and Development.* 1(2). 151-195.

Hyman, L. M. 1976. Phonologization. In *Linguistic Studies Offered to Joseph Greenberg: Vol. 2 Phonology*, ed. By A. Juilland, 407-418.

Kapatsinski, V. 2010. Velar palatalization in Russian and artificial grammar: Constraints on models of morphophonology. *Laboratory Phonology*. 1(2). 361-393.

Kroch, A. 1989. Reflexes of grammar in patterns of language change. *Language Variation and Change*. 1. 199-244.

Lombardi, L. 2002. Coronal epenthesis and markedness. *Phonology.* 19. 219-251.

McMahon, A., P. Foulkes, *et al*. 1994. Gestural Representation and Lexical Phonology. *Phonology.* 11(2). 277-316.

Moreton, E. 2008. Analytic bias and phonological typology. *Phonology.* 25. 83-127.

Morley, R. L. (under review). The Emergence of Epenthesis: An Incremental Model of Grammar Change.

Newport, E. L. & R. N. Aslin. 2004. Learning at a Distance I. Statistical Learning of Non-adjacent Dependencies. *Cognitive Psychology.* 48. 127-162.

Niyogi, P. & R. C. Berwick. 1997. A dynamical systems model for language change. *Complex Systems.* 11. 161–204.

Ross, D. S. & E. L. Newport. 1996. The development of language from non-native linguistic input. In *Proceedings of the 20th annual Boston University Conference on Language Development Boston*, vol 2, 634-635.

Saffran, J. R. & E. D. Thiessen. 2003. Pattern induction by infant language learners. *Developmental Psychology*. 39(3). 484-494.

Senghas, A. & M. Coppola. 2001. The creation of Nicaraguan Sign Language by children: Language genesis as language acquisition. *Psychological Science.* 12. 323–328.

Singleton, J. L. 1989. Restructuring of language from impoverished input: Evidence for linguistic compensation. Doctoral Dissertation. University of Illinois at Urbana-Champaign.

Wilson, C. 2003. Experimental Investigation of Phonological Naturalness. In *Proceedings of the 22nd West Coast Conference on Formal Linguistics San Diego*, 101-114.

Wilson, C. 2006. Learning Phonology with Substantive Bias: An Experimental and Computational Study of Velar Palatalization. *Cognitive Science.* 30. 945-982.

Zsiga, E. C. 2000. Phonetic alignment constraints: consonant overlap and palatalization in English and Russian. *Journal of Phonetics*. 28. 69-102.