

Squibs and replies

*On measuring phonetic precursor robustness : a response to Moreton**

Alan C. L. Yu

University of Chicago

Much debate in recent years has focused on the relative contribution of analytic and channel biases in shaping the typology of sound. Moreton (2008) argues forcefully for the strength of analytic bias, such as Universal Grammar and other non-modality-specific cognitive biases that facilitate the learning of some phonological patterns and inhibit that of others, in creating typological asymmetries on its own, unassisted by the robustness of phonetic precursors. This article focuses on the assessment of phonetic precursor robustness. The main goal of this article is two-fold: (i) to establish the inadequacy of Moreton's method of evaluating relative phonetic precursor robustness and to offer an alternative to his approach; (ii) to report the results of a cross-linguistic study comparing the nature of vowel-to-vowel coarticulation and the interaction between obstruent voicing and vowel height with the same languages – no previous studies have directly compared these two phonetic precursors.

1 Introduction

What factors shape the synchronic typology of sound patterns and how should they be detected? To the extent that the synchronic typology of sound patterns follows from the results of language change, it is commonly, if not implicitly, assumed that analytic and channel biases are two major factors involved in shaping phonological typology (Wilson 2006, Zuraw 2007, Moreton 2008, 2009). ANALYTIC BIASES are limitations in computation or markedness relations and constraints imposed by Universal

* I have benefited tremendously from discussion with Sharon Peperkamp and Morgan Sonderegger and comments from four anonymous reviewers, an associate editor and the audiences at the Manchester Phonology Meeting in 2009, the Workshop on Computational Modeling of Sound Pattern Acquisition at the University of Alberta, University of Pittsburgh, Michigan State University and Stony Brook University. Naturally, all errors are my own. This material is based upon work partially supported by the National Science Foundation under grant #0949754. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

Grammar. An analytic bias might render certain patterns difficult to acquire even from perfect learning data. CHANNEL BIAS is the relative likelihood of a phonetic precursor to sound change becoming phonologised into full-fledged sound patterns (e.g. Hyman 1976, Ohala 1993, Lindblom *et al.* 1995, Hume & Johnson 2001a, Blevins 2004). This view of phonological typology is motivated by the commonly held assumption that sound patterns and sound changes that recur across unrelated languages originate in properties of human articulatory, perceptual and/or auditory mechanisms (Ohala 1983, 1993, Beddor *et al.* 2007). Context-induced phonetic variation in speech production and perception is taken to be the phonetic precursor to listener misperception-based sound changes. PHONOLOGISATION refers to gradient phonetic variation (i.e. intrinsic allophones) becoming entrenched and developing into categorical phonological patterns (i.e. the emergence of extrinsic allophones). Rare or uncommon sound patterns are due to the low probabilities of the corresponding phonetic effects being phonologised through sound change.

Much debate in recent years has focused on the relative contribution of analytic and channel biases in shaping the typology of sound patterns (Hale & Reiss 2000, Hyman 2001, Blevins 2004, 2006, Seidl & Buckley 2005, Kiparsky 2006, 2008, Wilson 2006, Zuraw 2007, Moreton 2008, 2009). The main issues are whether analytic and channel biases are in fact distinct, and, if so, how they should be teased apart. Moreton (2008), in a recent thought-provoking paper, argues forcefully for the strength of analytic bias, such as Universal Grammar and other non-modality-specific cognitive biases that facilitate the learning of some phonological patterns and inhibit that of others, in creating typological asymmetries on its own, unassisted by the robustness of phonetic precursors. His approach is three-pronged. Using as a case study the asymmetric rate of attestation between patterns involving vowel-to-vowel height dependencies (HH patterns) and those involving vowel height to consonantal voicing dependencies (HV patterns), he first conducts a brute-force library survey of grammars for the attestation of vowel-height harmony and consonantal voicing–vowel-height interaction. He then surveys previous acoustic studies of vowel-to-vowel coarticulation and the effect of consonantal voicing on vowel formants, in order to assess the strengths of these two coarticulatory effects. He concludes (2008: 93) that the skewed attestation of HH patterns over HV patterns cannot be explained by the relative robustness of their phonetic precursors, since the relative magnitudes of coarticulatory effects are comparable (i.e. the HV precursor being more robust than the HH precursor). He then tests for potential analytic biases against HV patterns, using an artificial grammar learning paradigm, and shows that English-speaking subjects are less able to learn HV patterns than HH patterns, concluding that there must be a preference for patterns involving the co-occurrence of a single feature over patterns involving the concurrence of multiple features.

This article focuses on the assessment of phonetic precursor robustness. The main goal of this article is twofold: (i) to establish the inadequacy of

Moreton's method of evaluating relative phonetic precursor robustness and to offer an alternative to his approach; (ii) to report the results of a cross-linguistic study comparing the nature of HH and HV interactions with the same languages – no previous studies have directly compared these two phonetic precursors. The article is organised as follows: I begin by first establishing that Moreton's method of evaluating relative phonetic precursor robustness is inadequate, as it measures inter-contextual, rather than intra-contextual, variation (§2). Using a rational optimal listener to model speech perception, I propose a method that more directly and objectively measures phonetic precursor robustness (§3). The effectiveness of this measure is demonstrated using results of a cross-linguistic acoustic investigation of HH and HV effects in English and Turkish (§4). The results of this investigation suggest that Moreton's claim about channel bias having no place in explaining the skewed typology of HH and HV patterns is premature. This paper concludes with a discussion about the significance of the present study and the complexity of isolating analytic bias from channel bias (§5).

2 Variation and channel bias

Moreton (2008) tests the relevance of channel bias in explaining the asymmetric rate of attestation between HH patterns and HV patterns by carrying out a survey of existing phonetic studies 'to assess the effect on target vowel F1 of the phonological height of a neighbouring vowel, and compare it with the effect of phonological voicing, aspiration or fortis/lenis status of an immediately following consonant' (2008: 93). He proposes to quantify the relative magnitude of the phonetic precursors (i.e. the magnitude of the coarticulatory effects) by comparing the target-vowel F1 in contexts where it is likeliest to raise or lower. In Moreton's own words, 'for HH studies, the Raising context consisted of high vowels, and the Lowering context consisted of low vowels. For HV studies, the Raising context was voiced, unaspirated or lenis obstruents, and the Lowering context was voiceless, aspirated or fortis obstruents' (2008: 93). The effect of context, which I refer to as the L/R ratio, is therefore defined as the target-vowel F1 in the Lowering context divided by the target-vowel F1 in the Raising context. This L/R ratio thus estimates the degree of inter-contextual variation; the further the ratio deviates from 1, the wider the range of F1 variation is going to be. To illustrate this concretely, consider the data in Table I. Table Ia summarises the average F1 values and the corresponding L/R ratios at the offset of /e/ and /a/ preceding /i/ or /a/ reported in Manuel's (1990) study of vowel-to-vowel coarticulation in three Bantu languages. Table Ib shows the average F1 values at the vowel offset of /a/ or /æ/ preceding voiced and voiceless obstruents /b p v f/ reported in Summers' (1987) study of the effect of final consonant voicing on vowel production in English.

		Raising	Lowering	L/R ratio
(a)	Manuel (1990)	__ Ci	__ Ca	__ Ca/ __ Ci
	Sotho	400·67	443·17	1·11
	Shona	410·67	471·00	1·14
	Ndebele	396·67	442·50	1·12
(b)	Summers (1987)	__ voiced C	__ voiceless C	L/R ratio
	English	578·67	691·67	1·20

Table I

Vocalic F1 values in Hz in different contexts and languages.

Drawing from measurements reported in a set of studies similar to Manuel's and Summer's, Moreton calculated L/R ratios for the HH and HV precursors in different languages and found the HH precursor to be less robust than the HV precursor. This conclusion is surprising if the likelihood of phonologisation is the primary driver behind the frequency typology of sound patterns, since Moreton's own typological survey establishes that HH interactions are more widely attested than HV interactions. He thus concludes that the overattestation of HH patterns relative to HV patterns must stem from some sort of analytic bias that favours the learning of intra-tier feature interaction over inter-tier featural co-occurrence restrictions.

While Moreton's L/R ratio uses existing phonetic studies to establish a phonetic typology, its effectiveness as a measure of phonetic precursor robustness is dubious. To begin with, it is limited to the comparison of precursors that influence segments along the same phonetic dimension. For example, the L/R ratio would have difficulty comparing the robustness of the precursor to velar palatalisation with that of the precursor to HV interaction, since the main phonetic cues to velar palatalisation are peak spectral frequency and formant transitions (Guion 1996), while the primary cue to HH and HV interaction is the trajectory of the first formant. More problematic is how the L/R ratio relates to the robustness of phonetic precursors. The L/R ratio is, at its core, a measure of intra-category variability induced by a phonetic precursor. That is, it provides a rough measure of the spread of the category distribution along certain phonetic dimensions; the numerator and denominator are the means of a category's variants at the extremes of a category's distribution. To illustrate this point more concretely, consider Fig. 1, which shows the distribution of category *a* within the context of a phonetic precursor *k* along some phonetic dimension *X*. The grey lines delineate the distributions of *a* in the Lowering and Raising contexts of precursor *k*. From the perspective of regressive HH interaction, for example, where *a* is a vowel,

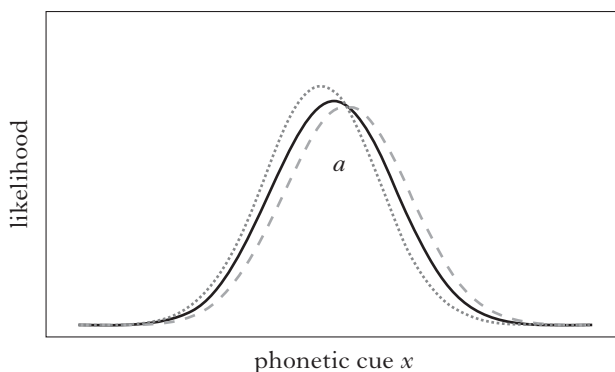


Figure 1

The distribution of category *a* along some phonetic cue dimension *x*. The grey lines delineate distributions of category *a* in the Lowering (dashed line) and Raising (dotted line) contexts.

the phonetic cue *X* is F1 and the precursor *k* is the presence of a following vowel, the Lowering context (dashed line) is *a* in the context of a low vowel, while the Raising context (dotted line) is *a* in the context of a high vowel. The black solid line shows the overall distribution of *a* in the presence of precursor *k*. What this illustration shows is that the further apart the category distribution in the Lowering and Raising contexts is (i.e. the further the L/R ratio deviates from 1), the wider range of variation the precursor induces.

While variation is a necessary condition for change, it is not a sufficient one. A category that has a wide variance may nonetheless show great stability if it overlaps minimally with other contrastive categories in the perceptual space. On the other hand, a narrowly distributed category might be highly prone to sound change if it has a high degree of overlap with other contrastive categories in the system. Consider, for example, Fig. 2, which shows the distributions of three categories, *a*, *b* and *c*, in the presence of precursor *k*. While these categories exhibit identical distributions with respect to cue *X* (i.e. they have the same L/R ratio and variance; they only differ in their means), they nonetheless have very different phonologisation profiles: categories *a* and *b* are prone to undergo sound change in the presence of precursor *k*, while category *c* is less likely to do so. The logic behind these conclusions follows from the way these distributions overlap with each other. From the perspective of a listener-misperception view of sound change, the likelihood of sound change is determined by the likelihood of misperception prompted by the presence of some phonetic precursor. The more overlap there is between two category distributions, the more uncertain and less accurate the recognition of the two categories are going to be. From this point of view, then, the likelihood of sound change amounts to the degree of overlap between a set of category distributions along some perceptually relevant phonetic dimensions.

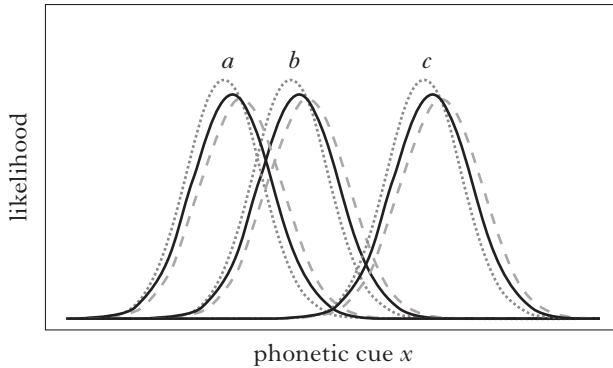


Figure 2

Distributions of categories *a*, *b* and *c* along some phonetic cue dimension *x*. All share the same L/R ratio and mean variance.

The robustness of a phonetic precursor is therefore a measure of the degree of confusion induced by the presence of that precursor. That is, the mean and variance of a category might shift as a result of the presence of a phonetic precursor. Such a shift may result in an increase or decrease of overlap with other sound categories along the same phonetic dimension. The larger the precursor-induced shift toward the direction of greater overlap, the more robust this precursor is going to be. Thus, given that categories *a* and *b* show much overlap under precursor *k*, we expect these categories would more likely undergo sound change in the presence of *k* than category *c*, which has only a small overlap with categories *a* and *b*.

In sum, an objective measure of phonetic precursor robustness should ideally be feature- or cue-independent, and should offer a way to capture the degree of confusion between sounds in the presence of a phonetic precursor. One way to evaluate the relative robustness of a phonetic precursor is to ascertain the likelihood of one sound being confused with another in a given context. That is, a means to evaluate intra-contextual variation is needed. For example, to assess the magnitude of anticipatory coarticulation which /i/ has on /a/, it is imperative to determine the likelihood of /a/ being confused with a higher vowel, say /e/, in the same context. Likewise, to evaluate the robustness of /i/-lowering due to the presence of a following /a/, it is useful to know how likely /i/ is confused with /e/ in that context. The most direct way of arriving at this estimate is to conduct a perceptual study. For example, Beddor *et al.* (2007) asked English and Shona speakers to classify an /a/ ~ /e/ continuum embedded in the contexts __C/a/ and __C/i/. The magnitude of the effect of vocalic context on the classification of /e/ or /a/ can be evaluated by examining the resulting identification function, the canonical S curve. In particular, the relative strengths of different contexts have on sound categorisation can be assessed by comparing the slope at the cross-over point across contexts (e.g. the point where the probability of a signal is maximally ambiguous

between categories). Clayards *et al.* (2008), for example, show that listeners are sensitive to the entire probability distribution of acoustic phonetic cues for a sound category, and that the precision or amount of certainty about a category that a particular cue provides is inversely proportional to the variance of that cue for that category. That is, the more overlap in the probability distributions of two categories, the more uncertainty the listeners will have about which category they are hearing.

To be sure, cross-linguistic perception studies are few. Perceptual studies of vowel identification in the context of different following obstruents with varying voicing specification are even fewer, if any exist at all. Thus, a meaningful cross-linguistic phonetic survey based on existing perceptual studies is not yet feasible. However, recent rational models of human cognition have shown that human perceptual behaviours can be insightfully modelled as ‘ideal observers’. Building on this line of reasoning, I propose in the next section a method for evaluating phonetic precursor robustness even in the absence of perceptual data; human perceptual responses may be estimated, given certain assumptions about human perceptual behaviour and the probability distributions of the cues of the target categories.

3 Phonetic precursor robustness as degree of uncertainty

Optimal rational listener models of speech perception (Feldman & Griffiths 2007, Clayards 2008, Clayards *et al.* 2008, Flemming 2010) assume that decisions about perceptual information are guided by certain basic principles that optimise signal identification. To illustrate what a rational analysis might look like, let us consider a scenario where the listener has to decide whether some signal S in context k belongs to some category, say c_1 or c_2 , as in a canonical two-alternative forced choice task. Adopting Sonderegger & Yu’s (2010) model of perceptual compensation for coarticulation, itself an elaboration of Feldman & Griffiths (2007) and Feldman *et al.* (2009), listeners are assumed to perceive a noisy stimulus S that is normally distributed around a target pronunciation T , as listeners cannot recover the target pronunciation T directly. The target pronunciation T is itself normally distributed around a category mean, as in (1).

$$(1) T|c_i, k \sim N(\mu_{c_i, k}, \sigma_c), \quad S|T, c_i \sim N(T, \sigma_s)$$

where $\mu_{c_i, k}$ is the mean of category c_i in context k , σ_c^2 is the variance in T around the category mean and σ_s^2 is the variance in S around T . For the sake of simplicity, the variance of the category σ_c^2 and the variance of the signal σ_s^2 are assumed to be the same for categories 1 and 2.

The ideal optimal listener categorises based on the likelihood of S being an instance of the speaker producing an example from c_i in context k , with

target T . The probability of S coming from category c_1 can be calculated with Bayes' rule, as in (2).

$$(2) \quad P(c_1|S, k) = \frac{P(c_1|k)P(S|c_1, k)}{P(c_2|k)P(S|c_2, k) + P(c_1|k)P(S|c_1, k)}$$

As shown in Equation (2), $P(c_i|k)$ is the probability of category i occurring in context k , i.e. in the lexicon as a whole. $P(S|c_i, k)$ is calculated by integrating over all possible targets T , giving the equation in (3).¹

$$(3) \quad P(c_1|S, k) = \left(1 + \frac{f_2}{f_1} e^{b-Sg}\right)^{-1}$$

where $f_i = P(c_i|k)$ is the frequency of category i in context k and the variables b and g are defined as in (4).

$$(4) \quad b = \frac{1}{2} \frac{\mu_{c_1,k}^2 - \mu_{c_2,k}^2}{\sigma_S^2 + \sigma_c^2}, \quad g = \frac{\mu_{c_1,k} - \mu_{c_2,k}}{\sigma_S^2 + \sigma_c^2}$$

Given that the confusability between two sounds is characterised by the nature of the identification function of these sounds, an objective measure of the nature of the identification function between two sounds is also a

¹ Perceptual compensation for coarticulation is treated in Sonderegger & Yu (2010) as a consequence of the variability in the differences between the means, variances and/or frequencies of two categories in a given context. For example, if the mean F1 shifts downwards for both categories in context k_i compared to context k_j , then the perceptual boundary between the two categories in context k_i would also shift downwards relative to the perceptual boundary in context k_j . A reviewer notes that such a model of perceptual compensation for coarticulation might not be compatible with the approach to sound change which Moreton implicitly assumes, that is, new sound variants arise when unsophisticated listeners, such as children and second language learners, fail to properly perceptually compensate for coarticulation (Ohala 1993). To fully address this matter would take the present discussion too far afield. However, there are reasons to think that this is not an insurmountable problem. To begin with, the idea that only unsophisticated listeners might participate in sound change is empirically dubious; recent experimental evidence have suggested that the sound system might change throughout an individual's life time (Sancier & Fowler 1997, Harrington *et al.* 2000, Sankoff 2004, Harrington 2006, Evans & Iverson 2007). Even if unsophisticated listeners were contributors to change, there are reasons to believe that differences in perceptual compensation responses by unsophisticated listeners such as children may be experientially driven. Children have been shown to adjust the range of cues they attend to as they develop (Nittrouer 2002, Nittrouer & Lowenstein 2009); children initially weight dynamic spectral cues (i.e. formant transitions) more than adults, and weight stable spectral cues (e.g. noise spectra) less (Nittrouer & Miller 1997, Watson 1997, Mayo *et al.* 2003). Finally, recent studies have suggested that the role of perceptual compensation for coarticulation in sound change might be better conceptualised in terms of across-individual differences than in terms of within-individual variation. That is, the introduction of stable new variants into a language might be better conceptualised as resulting from systematic differences in perceptual and production norms of individuals with different perceptual compensation strategies, rather than from occasional misperception by individuals (Yu 2010, to appear).

measure of the confusability of these sounds. With the above-mentioned rational model in mind, I propose to quantify context-induced confusability in terms of the normalised slope of the identification function in (5).²

(5) *Slope of Precursor Robustness* (SPROB)

$$\frac{(\mu_{c_1,k} - \mu_{c_2,k})^2}{4(\sigma_S^2 + \sigma_c^2)}$$

Thus the magnitude of perceptual confusion in the identification of category *a* relative to category *b* under the influence of Precursor 1 and Precursor 2 (i.e. the relative robustness of Precursor 1 and Precursor 2) is determined by the respective SPROB scores of the *a*~*b* identification functions in Context 1 and Context 2. The lower the SPROB score a precursor has, the more uncertainty the precursor introduces to the discrimination between *a* and *b*. The more uncertainty there is between *a* and *b* in the presence of the precursor, the more robust the precursor is.

To further clarify the relationship between variance and slope and their effects on the interpretation of phonetic precursor robustness, consider Fig. 3. Figure 3a shows hypothetical probability distributions for two vowel categories along some phonetic dimension *X* under the influence of two phonetic precursors, 1 and 2. The means of the distributions are the same distance apart, but the variances differ. The dashed lines correspond to some sound categories, such as vowels, which are produced consistently and thus have narrow distributions under the influence of Precursor 1; the solid lines correspond to the same sound categories under the influence of Precursor 2, where the sounds are produced less consistently, with wide distributions. The optimal solution for each pair of distributions in Fig. 3a is illustrated in Fig. 3b. Even though for both solutions the category boundary (the point where the categorisation of *a* and *b* is at chance) is in the same place along the *x*-axis, distributions with greater degree of overlaps (solid lines) have higher degree of uncertainty in category identification, as reflected in the shallowness of slope of the categorisation function.

Here, I propose to measure the strength of a phonetic precursor (i.e. the likelihood of a phonetic precursor resulting in sound change) in terms of the SPROB in (5), which measures the degree of uncertainty (quantified in terms of the normalised slope (5)) engendered by the intra-context variation inherent in the signal. The lower the SPROB, the more robust the phonetic precursor, since a lower SPROB score corresponds to a higher degree of uncertainty. The main idea motivating this approach is

² Normalisation is accomplished by multiplying the actual slope of the identification function with the differences between the means of the two categories ($\mu_{c_1,k} - \mu_{c_2,k}$). This normalisation feature allows the SPROB score to be dimensionless. Thanks to Morgan Sonderegger for suggesting this method of normalisation.

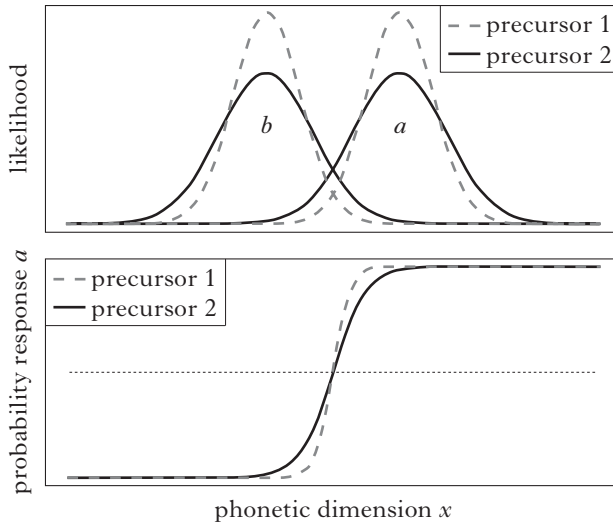


Figure 3

The top panel shows the probability distributions of categories *a* and *b* along some phonetic dimension *x* in the contexts of two phonetic precursors: the variance of the cue distributions under Precursor 1 is narrow (dashed lines); the variance under Precursor 2 is wide (solid lines). The bottom panel shows the optimal response curves calculated from the probability distribution using equation (3) for the narrow condition (dashed line) and the wide condition (solid line). The horizontal line in the bottom panel indicates the crossover point; SPROB measures the normalised slope at that crossover point.

the hypothesis that an increase in uncertainty in perception leads to an increase in the likelihood of listener misperception-driven sound change (i.e. the strength of the channel bias).

To better illustrate how this estimate is assessed and its application to understanding relative phonetic precursor robustness, I conducted a cross-linguistic production studies contrasting the effects of anticipatory vowel-to-vowel height coarticulation and anticipatory interaction between vowel height and consonantal voicing.³ The relative robustness of these phonetic precursors is evaluated in terms of the uncertainties of /i/ and /e/, on the one hand, and /ɑ/ and /e/ on the other, in the context of a following

³ While it would have been instructive to apply the present model to the phonetic studies reviewed in Moreton's phonetic survey, this is unfortunately not possible, given that only English has been studied for both vowel-to-vowel coarticulation and height-voice interaction, and, in addition, the relevant studies did not report all the parameters (e.g. standard deviations) needed by the present model to calculate the SPROB scores.

obstruent that is either voiced or voiceless (the HV precursors) and a following vowel that is either /a/ or /i/ (see (6)).

(6) *Summary of the contexts and vowel targets investigated in the production experiment.*

Vowel pairs	HH context	HV context
/i/ ~ /e/	__ Ca	__ [+voice] or __ [-voice]
/a/ ~ /e/	__ Ci	__ [+voice] or __ [-voice]

The results of the production study serve as the basis for estimating the SPROB scores for the HH and HV phonetic precursors in English and Turkish. Given that previous studies have shown that obstruent voicing has the effect of lowering the F1 value of a preceding vowel, /e/ is expected to encroach on the vowel space of /i/, and /a/ on the vowel space of /e/. Given the F1-raising effect of /a/ on the preceding vowel, /i/ is expected to encroach on the vowel space of /e/, and /e/ on the vowel space of /a/. Likewise, the F1 values of /a/ and /e/ should be lower in the context of /i/. Thus, if the robustness of HH is stronger than HV, the SPROB scores for both pairs of vowels in the HH contexts should be lower than the same pairs of vowels in the HV contexts. Simply put, vowels in HH contexts should behave more like the wide variance distributions (solid lines) in Fig. 3, while the same vowels in HV context should resemble the narrow variance distributions (dashed lines).

4 A production study

The following study examined the effect of anticipatory vowel-to-vowel height coarticulation (HH) against the effect of consonantal voicing and vowel-height interaction (HV). Previous studies on vowel-to-vowel coarticulation in English reported greater carry-over effects, particularly for unstressed vowels. For example, using CVCV reiterant speech, Majors (1998) reported that unstressed vowels are subject to greater V-to-V influence in English than stressed ones. This stress asymmetry is subject to vowel-specific variation; the stress asymmetry is significant for /i/ but not for /o/. Focusing on the production of /bV'bVbV/ strings, where V = /i e a o u/, Beddor *et al.* (2002) also found more extensive carry-over V-to-V coarticulation than anticipatory effects. While coarticulation is strongest at vowel edges, statistically significant anticipatory coarticulation is also found in English when the vowel is unstressed. With respect to V-to-V coarticulation in Turkish, Beddor & Yavuz (1995), who focused on coarticulation of vowel frontness, found that anticipatory coarticulation is more consistent than carry-over coarticulation in Turkish. In particular, they found anticipatory and carry-over fronting of /a/ with neighbouring /i/, anticipatory backing of /i/ with following /a/, and slight dissimilation of /i/ with respect to preceding /a/. Beddor & Yavuz (1995) speculate that the reason why the anticipatory coarticulation pattern is more consistent

than the carry-over effect might have to do with the word-final stress pattern of Turkish, in view of the fact that stressed vowels tend to exert a greater coarticulatory influence on neighbouring vowels than unstressed ones do. Inkelas *et al.* (2001), who expanded on Beddor & Yavuz's study by including stress as a factor of vowel-to-vowel coarticulation in Turkish, replicated Beddor & Yavuz's results, but found no stress effects. To the best of my knowledge, there is no existing study that has investigated HV effects in Turkish. Previous studies on HV effects in English have focused on the effect of final obstruent voicing in CVC strings. For example, Summers (1987), who focused on /bVC/ syllables, where V = /a æ/ and the final consonant = /b p v f/, reported higher F1 frequencies throughout the vowel (i.e. during the onset, steady state and offset) when the final consonant was voiceless than when it was voiced. Hillenbrand *et al.* (1984) found that CVC stimuli which were identified at least 75% of the time as ending in a voiceless consonant tended to have higher F1 offset frequencies than those which were identified at least 75% of the time as ending in a voiced consonant. This effect was much larger for the vowels with a higher first formant frequency; F1 offset frequency differences were only 17 Hz and 23 Hz when the vocalic contexts were /i/ and /u/ respectively, but was 76 Hz for /a/, and ranged from 169 to 258 Hz for /ɛ/. The nature and extent of coarticulation is language-specific. The range of phonemic contrasts within a language, for example, restricts the magnitude of variability induced by coarticulation (Manuel 1990, Beddor *et al.* 2007); the greater the number of phonemic contrasts along an acoustic-perceptual dimension, the less coarticulatory variability is observed. Thus, given the smaller range of height contrast in Turkish than in English, HH and HV coarticulatory effects in Turkish are expected to be more pronounced than in English.

4.1 Methods

The stimuli consisted of nonsense trisyllabic strings of the form /'adV₂CV₃/, where C was /p b/ and V₂ and V₃ /i e a/.⁴ The three vowels in the stimulus materials were selected as the closest counterparts in the two languages: phonetically [i e a] in Turkish and [i e a] in English. A total of eighteen stimuli were constructed. In order to avoid the influence of stress on coarticulation (Beddor *et al.* 2002), subjects were instructed to accentuate the first syllable of each target word. The medial vowel, V₂, was the target of this analysis.

⁴ One reviewer pointed out that the use of potentially phonotactically illegal strings here might present difficulties for the native English subjects, especially since vowels are often reduced in unstressed positions. This concern is mitigated by the fact that this is a cross-linguistic study. Given that the main focus of this research is in uncovering the nature of phonetic precursors across languages, to the extent that coarticulatory effects are observed even in phonotactically illegal strings in the native language of the speaker, it only strengthens the universal basis of such effects.

Six native speakers of North American English (2 female and 4 male) and four native speakers of Turkish (1 female and 3 male) were recruited to take part in the experiment. The English speakers were linguistics graduate students at the University of Chicago, with some phonetic training; the Turkish speakers were graduate students at the University of Chicago, with no linguistic training. Following Beddor *et al.* (2002), phonetically trained native English speakers were chosen in order to avoid difficulties with English orthography and excessive vowel reduction and diphthongisation of the mid vowel /e/. Subjects were paid a nominal fee for participation. Speakers were recorded reading a randomised list of test stimuli ten times. The target trisyllables (given in IPA transcription in brackets []) were embedded in a sentence context, English *Give me an __ tomorrow* and Turkish *açıkça __ söyle* ‘clearly __ say’. The test stimuli were displayed one at a time on a computer screen, using a Java script. Subjects were instructed to read the sentence aloud, at a normal speed. English subjects were additionally instructed to read the target trisyllables as closely to their IPA values as possible. Turkish subjects were instructed to read the target trisyllables as they would sound in Turkish. The utterances were recorded in a soundproofed room and were digitised directly at 48 kHz, using an Ediral digital recorder. F1 values were measured using an LPC algorithm, implemented in Praat.⁵ Measurements were taken at eleven equidistant points of the target vowel, starting from the onset of the vowel, measuring at 10% increments of the total duration of the vowel. Measurements at the eleventh measure point were taken at the beginning of stop closure. Only the values of points 2 to 10 were considered in the statistical analysis, to avoid errors introduced by closure boundaries. Outliers, defined as F1 values more than two standard deviations from the means of their condition, were discarded from further analysis. The linear formant frequency scale was converted to a Bark scale (Traunmüller 1990). Statistical analyses were performed on the Bark-transformed formant frequencies.

4.2 Results

Figures 4 and 5 summarise the results of the production study (see the Appendix for a detailed summary). The effect of vowel-to-vowel coarticulation on F1 is given in Fig. 4, while the nature of the plosive voicing–F1 interaction is presented in Fig. 5. For ease of reference, the F1 values are given in Hz rather than in a Bark-transformed scale. Visual inspection confirms that the subjects indeed articulated three vowels with distinct vowel heights, suggesting that the subjects were able to accomplish the task as instructed.

⁵ Perceived vowel height is influenced by multiple acoustic phonetic factors (e.g. F0 and duration). However, to ensure maximal comparability with Moreton’s original study, the present study follows his lead in focusing on just F1 values. It is worth noting, though, that the same SPROB calculation can be conducted for all relevant acoustic phonetic cues.

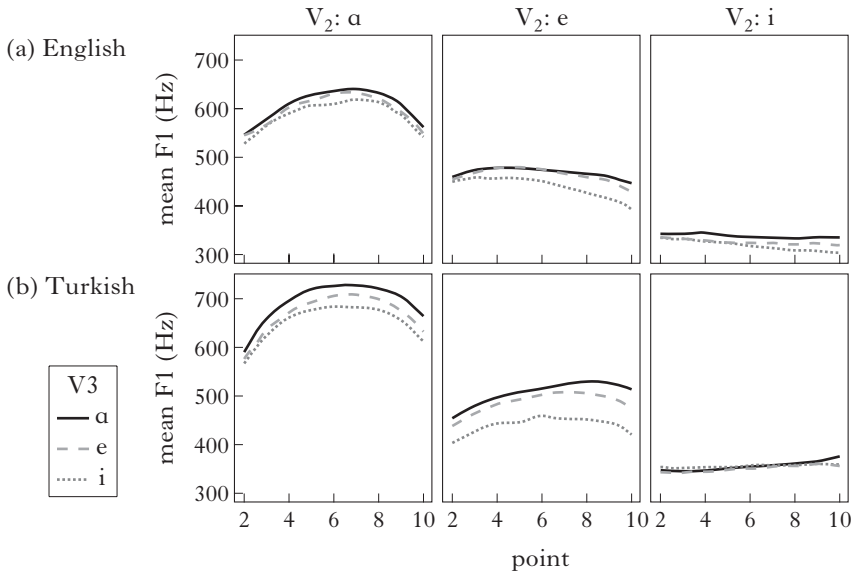


Figure 4
 F1 values of /a e i/ preceding /a/, /e/ or /i/
 from speakers of English and Turkish.

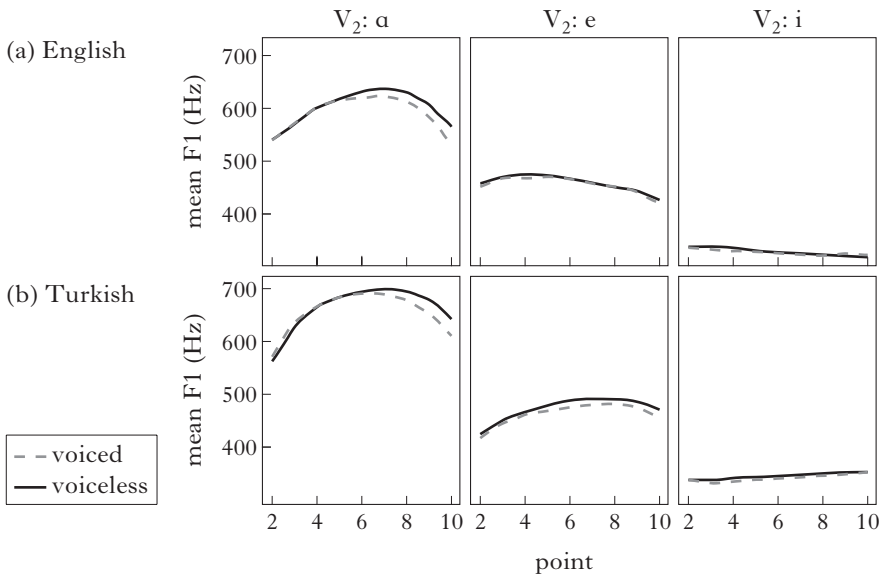


Figure 5
 F1 values of /a e i/ preceding [±voice] obstruents
 from speakers of English and Turkish.

	df	English		Turkish	
		F value	<i>p</i>	F value	<i>p</i>
V ₂	2	26240.028	< 0.001	9737.023	< 0.001
V ₃	2	208.851	< 0.001	133.910	< 0.001
Voicing	1	10.301	0.001	14.457	< 0.001
Point	8	91.830	< 0.001	56.112	< 0.001
V ₂ × V ₃	4	10.429	< 0.001	43.038	< 0.001
V ₂ × Voicing	2	8.364	< 0.001	0.622	0.537
V ₃ × Voicing	2	3.863	0.021	3.852	0.021
V ₂ × Point	16	55.420	< 0.001	15.002	< 0.001
V ₃ × Point	16	2.887	0.000	1.080	0.368
Voicing × Point	8	0.699	0.693	0.731	0.665
V ₂ × V ₃ × Voicing	4	4.989	0.001	0.939	0.440
V ₂ × V ₃ × Point	32	1.166	0.238	0.088	1.000
V ₂ × Voicing × Point	16	2.846	< 0.001	0.695	0.802
V ₃ × Voicing × Point	16	0.423	0.978	0.092	1.000
V ₂ × V ₃ × Voicing × Point	32	0.504	0.991	0.158	1.000

Table II

Results of ANOVA for first formant values (based on 9518 responses from six English speakers and 6050 measurements from four Turkish speakers).

Figure 4 shows that the F1 values of the target vowel is highest when the following vowel is /a/ (solid line) and that the F1 values are low when the following vowel is /i/ (dotted line). As illustrated in Fig. 5, the effect of consonantal voicing on the F1 values of a preceding vowel is less pronounced. The F1-lowering effect of a voiced stop is most apparent when the target vowel is low. Visual inspection also reveals a difference in the temporal dynamics of HH and HV coarticulations. While HH coarticulation is evidenced across the entire target vowel, the HV effect, if any, is evident only in the latter half of the target vowel (i.e. the portion closest to the consonantal trigger).

To examine the statistical significance of these differences, a series of four-way repeated measures ANOVA with V₂ (/i e a/), V₃ (/i e a/), Voicing (voiced *vs.* voiceless) and Measurement Point (Point 2 to Point 10) as fixed factors and Subject as the error term was performed for each language group. The statistical results are summarised in Table II. As expected, main effects of V₃ and Voicing, as well as V₂ and Measuring Point, are observed in both the English and Turkish data. Significant interaction between V₂ and V₃ is observed in both English and Turkish, suggesting that the degree of anticipatory coarticulation varies depending on the quality of the target vowel. Post hoc analyses (α level adjusted to 0.017 for

three comparisons) shows that, with the sole exception of $V_2 = /i/$ in Turkish, where the V_3 effect does not reach the adjusted α level ($F(2, 2032) = 3.14, p = 0.04$),⁶ the effect of V_3 on all other target vowels in both languages are extremely robust ($p < 0.001$). Turkish shows greater degree of coarticulation than English, at least with respect to the non-high V_2 . This finding is consistent with the idea that a language with a smaller number of phonemic contrast along an acoustic-perceptual dimension should be more susceptible to coarticulation (Manuel 1990, Beddor *et al.* 2007). The lack of a strong coarticulatory effect when V_2 is $/i/$ in Turkish might be attributed to the effect of the preceding vowel; previous studies show that $/i/$ dissimilates (i.e. $/i/$ is higher and fronter rather than lower and backer) when the preceding vowel is $/a/$ (Beddor & Yavuz 1995, Inkelas *et al.* 2001). The cause of this dissimilatory effect is not known, although Beddor & Yavuz (1995) speculate that 'dissimilation' may actually be lack of centralisation when $V_1 \neq V_2$. In the case of English, significant interactions between V_2 and Voicing, and between V_2 , Voicing and Measuring Point, suggest that the degree of F1 raising from obstruent voicing depends not only on the target vowel but also varies in time. Post hoc analyses (α level adjusted to 0.017 for three comparisons) show the V_2 -dependent voicing effect to be significant only when $V_2 = /a/$ ($F(1, 3113) = 12.23, p < 0.001$). It is noteworthy that previous studies of HV effects have mainly focused on low or non-high vowels in English. The fact that the present study shows that no significant HV effect is observed with $/e/$ and $/i/$ further highlights the specificity of the HV effect, suggesting that this might be an additional obstacle to HV effects being phonologised. One of the more intriguing findings is a significant, albeit weak, interaction between V_3 and Voicing in both English and Turkish. Figure 6 shows that F1 values generally trend lower preceding a voiced plosive when a non-low vowel follows the target vowel. However, this V_3 -dependent voicing effect did not reach statistical significance in post hoc analyses.

4.3 Evaluating relative phonetic precursor robustness

4.3.1 *Predictions of the SPROB approach.* The main goal of the cross-linguistic production study is to assess the relative robustness of HH and HV precursors. To quantify the relative phonetic precursor robustness, the SPROB scores for $/i \sim e/$ and $/a \sim e/$ were calculated using equation (5). Specifically, four contextual effects were considered: $/i \sim e/$ in a lowering context ($V_3 = /a/$), $/a \sim e/$ in a raising context ($V_3 = /i/$) and $/i \sim e/$ and $/a \sim e/$ followed by a plosive that is voiced (a vowel-raising context) or voiceless (a vowel-lowering context). With respect to V_1 and V_2 in (5), in the $/i \sim e/$ comparison, $V_1 = /e/$ and $V_2 = /i/$, while in the $/a \sim e/$ comparison $V_1 = /a/$ and $V_2 = /e/$. Formant values were based on the midpoint of the target vowel in the specific context (see the Appendix for actual values).

⁶ The V_3 effect on $/i/$ is trending in the right direction.

	i ~ e			a ~ e		
	__Ca	__ [+voi]	__ [-voi]	__Ci	__ [+voi]	__ [-voi]
English	1.13	1.33	1.25	0.58	0.60	0.66
Turkish	1.32	0.91	0.93	0.50	0.64	0.69

Table III

SPROB scores for /i ~ e/ and /a ~ e/ in four vowel raising and lowering contexts.

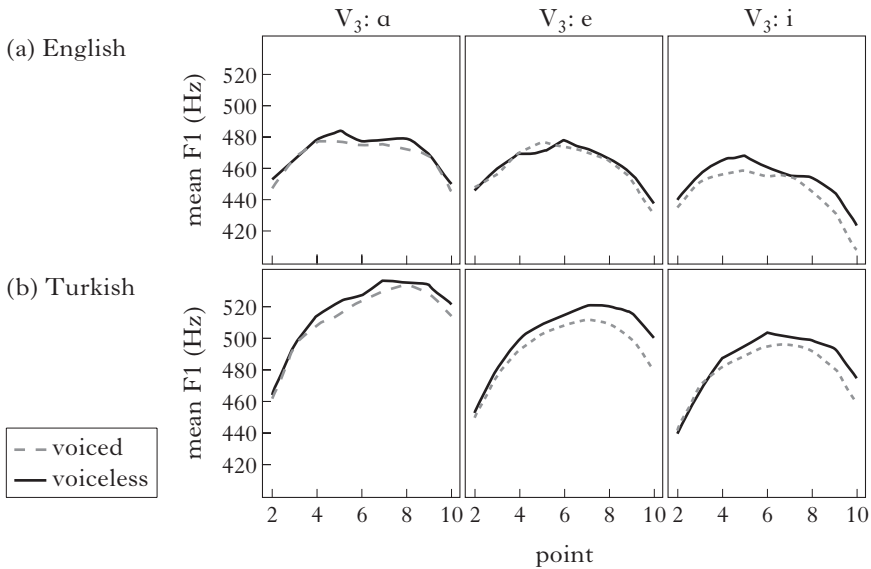


Figure 6

F1 values of V_2 under the influence of V_3 and the voicing of the following obstruent from speakers of English and Turkish.

Following Sonderegger & Yu (2010), variances for the target vowels (the $\sigma_S^2 + \sigma_c^2$ term in (5)) were taken to be the mean of the variances for the target vowels in a given context. For example, for /i ~ e/ in a lowering context ($V_3 = /a/$), the variances would be the mean of the variances for the /'adiba/ and /'adeba/ stimuli. Table III summarises the SPROB results.

Recall that a larger SPROB score corresponds to a steeper category boundary. Thus, the shallower the slope, the more categorisation uncertainty exists. Table III shows that, all else being equal, listeners' classification of vowels is more uncertain when the target signal is influenced by an HH precursor than by an HV precursor. This result suggests that the HH precursor is more robust than the HV precursor; it is more likely to

result in sound change than the HV precursor. The only exception to this generalisation is /i ~ e/ in Turkish, which shows HV environments yielding lower SPROB scores than the vowel-lowering __Ca context. This exceptional behaviour is to be expected in light of the results of the acoustic study and the general dissimilatory behaviour of /i/ when the preceding vowel is /a/, as reported in earlier studies (Beddor & Yavuz 1995, Inkelas *et al.* 2001). Recall that the effect of V₃ on /i/ is weak in Turkish. Thus while /e/ is drastically lowered by the following /a/, little lowering is evidenced in /i/. This leads to an increased separation between /i/ and /e/ before /a/. Thus, all else being equal, a wider separation between distributions of two vowel categories translates into a steeper category boundary slope (i.e. a higher SPROB score). The behaviour of /i ~ e/ in Turkish is thus the exception that proves the rule. Exceptions of this sort highlight the importance of not treating channel bias as a monolithic factor that applies to all languages indiscriminately; the nature and magnitude of coarticulatory effects, and intra-context variation in general, is largely language-specific.

The robustness of the HH precursor compared to that of the HV precursor extends beyond the SPROB comparison. While the production study shows that anticipatory vowel-to-vowel coarticulation is evidenced in both English and Turkish, the extent of HV interaction is mediated temporally as well as by target vowel quality, at least in the case of English. That is, while anticipatory HV coarticulation dissipates the further away the vocalic gesture is from the following obstruent, the anticipatory HH effect shows no such temporal dependency. The locality of the HV effect may further weaken its effectiveness as a phonetic precursor to sound change. The temporal restrictiveness of the HV precursor suggests two potential venues for further investigation. To begin with, the phonological effect of HV interaction might also be localised in the part of the vowel closest to the triggering consonant. For example, voicing in a preceding obstruent may only affect the first portion of a following vowel's height. Such a pattern is observed in the history of the Austronesian language Haroi (Lee 1977, Denning 1989): following a voiced obstruent, *a > /ia/, *au > /iau/, *əi > /ii/; voicelessness in a preceding obstruent, on the other hand, conditioned lowering of part or all of a high vowel (after a voiceless obstruent, *u > /o/, *i > /ej ε/).⁷ Another potential venue for exploring HV effects might be in its interaction with other precursors. Recall, for example, that there is an interaction effect between Voicing and V₃. Such an interaction may manifest itself in the form of passive resistance in natural language. For example, in a language with height harmony, voiced

⁷ The opposing effect of obstruent voicing on vowel height is consistent with the fact that obstruent voicing lowers F1, and hence raises vowel height, while obstruent voicelessness raises F1, thus lowering vowel height. That post-voiced obstruent raising only affects low vowels might be due to the fact that high vowels already have a low F1 to start with. Likewise, the fact that lowering after a voiceless obstruent only targets high vowels is presumably because low vowels already have a high F1.

		Raising	Lowering	L/R ratio
(a)	V-to-V	__ Ci	__ Ca	__ Ca/ __ Ci
	English	540·1	560·2	1·04
	Turkish	612·3	664·5	1·09
(b)	Voicing	__ voiced C	__ voiceless C	L/R ratio
	English	533·5	566·2	1·06
	Turkish	621·6	654·8	1·05

Table IV

Vocalic F1 values in Hz measured at the target vowel offset (Point 10) in different contexts and languages.

obstruents may serve as blockers. In the Buchan dialect of Scots, for example, there exists a partial height harmony where an unaccented high vowel in a suffix agrees in height with the preceding stressed root vowel unless the intervening consonant is either a voiced consonant or [l m n ŋ] followed by voiceless obstruents (Dieth 1932, Fitzgerald 2002, Paster 2004). Thus, *messy* is pronounced [mese] and *mealie* as [mili], but *doggie* as [dogi] and *bendy* as [bendi] (Paster 2004: 365–366). Assuming that partial height harmony in Buchan Scots originated from vowel-to-vowel coarticulation, the local F1-lowering effect of obstruent voicing might have been sufficient to counteract the stressed non-high vowel's effect of F1 raising on the unstressed high vowel of the suffix. This type of interaction between the effect of voicing on F1 and vowel-to-vowel coarticulation is readily evident in the production data presented above. Further investigation is needed to substantiate this type of synergistic interaction between consonantal voicing and vowel height.

4.3.2 *Comparison with the L/R ratio approach.* It is instructive at this point to return to Moreton's original proposal and to consider what predictions the L/R ratio analysis might make, given the present dataset. Following Moreton, the L/R ratios are calculated here on the basis of measurements taken at the target vowel offset. Table IV shows the L/R ratios for vowels before /a i/ and before voiced and voiceless stops in English and Turkish.

The results show that, at least in English, the L/R ratio is higher in the HV context than in the HH context. The opposite pattern, however, is observed in the Turkish case. Even though the L/R ratios based on the English data appear to show that the HH precursor is less robust than the HV precursor, while the Turkish results point to a more robust HH precursor than an HV precursor, the magnitude of difference in L/R ratios across precursor contexts is very small (the average L/R ratios for HH and HV precursors are 1·061 and 1·057 respectively), especially compared to

the differences observed in Moreton's survey (the average L/R ratios for HH and HV precursors in Moreton's study are 1.065 and 1.207 respectively). The results of the L/R ratio analysis would lead to the conclusion that the precursors of HH and HV interactions are comparable in robustness, and thus do not support the hypothesis that relative precursor robustness predicts the typological frequency of phonologisation. This conclusion is contradicted by the overall prediction of the SPROB analysis. This difference is to be expected, since Moreton's approach only measures the variation of a category across contexts. If a category varies by approximately the same amount in both HH and HV contexts, the L/R ratio analysis would suggest no differences in phonetic precursor robustness. The SPROB approach, on the other hand, emphasises the effects that context-induced variation has on between-category confusion. Thus, even though the HH and HV precursors might induce the same amount of variation with respect to one category, their effects might not be the same with respect to another category. The amount of potential confusion generated by a phonetic precursor might therefore not be identical across sound categories.

5 Conclusion

This paper began with a discussion of the inadequacy of Moreton's L/R ratio as a measure of phonetic precursor robustness. A method of evaluating the robustness of phonetic precursors, the SPROB, was proposed. This parameter, derived from a rational model of speech perception and production, offers a means to estimate listeners' perceptual response to intra-contextual variation, even when actual perceptual data from human listeners is absent. Based on the results of a cross-linguistic production experiment and the corresponding SPROB comparison, I have shown that the HH precursor is more robust than the HV precursor. As such, a channel-bias explanation for the underphonologisation of HV precursor remains very plausible.

To be sure, the viability of a channel bias account of the observed underphonologisation between the HH and HV precursors does not obviate the possible contribution of analytic bias in shaping this typology. In fact, I would submit that the type of rational model proposed in this and other work can be viewed as a type of analytic bias, since rational models assume humans perform optimally and rationally in cognitive tasks. Such models are also not modular-specific, in that they have been invoked to account for problems in domains of human cognition beyond language. In a very real sense, the model advocated in this paper provides a natural framework for bridging the divide between a purely experience-based approach to the emergence of sound pattern and a formal/cognitive approach, since the parameters which the rational model depends on are experience-driven. The outcome of categorisation is just as much affected by the model of optimisation selected as it is by the distribution of cues assumed by the speaker-listener who performs the categorisation task.

Naturally, there remains the issue of whether analytic bias can be diagnosed independently from channel bias. As Moreton painstakingly explains, the contribution of analytic bias is best revealed when the confound of channel bias is ruled out (this is the basic logic behind all under-phonologisation research). The best evidence for the independence of analytic bias from channel bias would come from a case that allows double dissociation between these two biases. The present study argues that the HH/HV asymmetry does not rise to the challenge of a double-dissociation test.⁸ Some might point to the results of the artificial grammar learning experiment as evidence for hints of an analytic bias at work. There are, however, reasons to be cautious, at least with respect to the significance of Moreton's experimental results. Moreton constructed an artificial language where the V_1 and C_2 of $C_1V_1C_2V_2$ words were either high and voiced or non-high and voiceless. Subjects were exposed to samples of this artificial language, and were tested subsequently for awareness of the co-occurrence restriction. Participants were found to have a hard time learning the HV patterns. Moreton attributes this difficulty to the learning of cross-featural tier dependencies (i.e. it might be more difficult to learn the dependencies between voicing features and height features than the dependencies among height features). The relative difficulty of acquiring feature dependencies notwithstanding, there is a potential confound which might obscure the interpretation of the experimental findings. As is well known, certain varieties of English possess a phonological rule, generally known as Canadian Raising, whereby a low vowel is raised when preceding a voiceless stop. Despite the name of this pattern, this raising phenomenon is not limited to Canada. Moreton & Thomas (2007), for example, noted that the /ai ~ li/ alternation has been reported in other parts of North America, including Virginia and adjacent parts of Maryland and North Carolina, as well as coastal South Carolina and Georgia (see Moreton & Thomas 2007 for references). As such, it cannot be discounted that students at the University of North Carolina at Chapel Hill, the subject population of Moreton's artificial grammar learning experiment, might either possess Canadian Raising in their phonologies or at least be aware of such an alternation in other varieties of English. If correct, the fact that subjects in Moreton's experiment exhibit difficulties in learning HV dependencies might be attributed to the fact that they were asked to learn a pattern that conflicts with either the phonology of their own language or the phonology of some neighbouring English varieties they are aware of. This contradiction might be sufficient to hinder the learning of the HV patterns to such a degree that the subjects achieved a lower success rate in learning the HV patterns than in learning the HH patterns.

⁸ Double dissociation, commonly used in neuropsychological testing, refers to a demonstration of two experimental manipulations having different effects on two dependent variables; one manipulation affects the first variable but not the second, and the other manipulation affects the second but not the first.

A reviewer of this article suggests that the strength of the HH precursor itself might have originated from the effect of analytic bias. That is, coarticulatory effects might be seen as the application of phonetic rules in the language rather than as mechanical side-effects that fall out as a consequence of the coproduction of various articulators (Whalen 1990, Kingston & Diehl 1994). Such an interpretation is particularly appropriate in models of the phonetics–phonology interface that eschew any differences between phonetics and phonology and propose instead to account for both the categorical and gradient patterns within the same grammatical system (Flemming 1995, 1996, 2001, 2002, Ní Chiosáin & Padgett 2001, Padgett 2003). The idea that vowel-to-vowel coarticulation and the effects of voicing on F0 and F1 are part of the phonetic knowledge of the speaker and that they are, if only partially, under the control of the speaker is certainly not new (Whalen 1990, Kingston & Diehl 1994, Scarborough 2004). However, if phonetic knowledge is to be classified under the scope of analytic bias, then it is not clear to what extent it is meaningful to differentiate channel bias from analytic bias in the first place. To what extent, for example, should the type of rational model advocated in this and other work be conceptualised in terms of channel or analytic bias, particularly in light of the applicability of Bayesian rational models to many domains of human cognition beyond language, as alluded to earlier? Given that this paper aims at addressing the more narrow issue of the proper measurement of the robustness of phonetic precursor as conceptualised within Moreton’s framework, I shall take refuge in the fact that Moreton himself took vowel-to-vowel coarticulation and the interaction between voicing and F1 as part of the channel bias, and leave the ontological question of where to draw the channel bias and analytic bias divide to another occasion.

Finally, it is worth mentioning that the SPROB approach, as described here, has limitations, since it only provides an expression of phonetic precursor robustness in the context of structure-preserving changes (i.e. the reassignment of ambiguous phonetic signals to *existing* phonological categories), as it requires the comparisons of cues between two already established sound categories. In the case of structure-building changes (i.e. the emergence of new allophones or phonological categories based on some ambiguous phonetic signals), the question of phonetic precursor robustness is no longer about confusion and reassignment between categories, but about inducing a distinct allophonic or phonemic category from a set of context-specific cues. Without taking the discussion too far afield, it is worth noting that there has been a host of recent work applying computational techniques of clustering and pattern recognition to the problem of phonological category induction (e.g. de Boer & Kuhl 2003, Lin 2005, Vallabha *et al.* 2007, Feldman *et al.* 2009, McMurray *et al.* 2009, Kirby 2010, to appear). Kirby, for example, illustrates a promising method of phonetic category induction using a model-based clustering approach. He models phonetic categories as Gaussian mixtures of cue distributions and shows that, given a set of phonetic cues, a two-category

solution (i.e. the positing of two allophones or two phonological categories) may be justified under certain combination of cues; other cue combinations would lead to complete merger. To be sure, further research is needed to ascertain whether such a model-based clustering approach is adequate for inducing novel allophones in the context of a structure-building phonetic precursor. The study of the robustness of phonetic precursors is only in its early years. As formal models of speech perception, such as the rational model promoted here, are developed and refined, better and more sophisticated methods of measuring precursor robustness will no doubt emerge.

In sum, channel bias, i.e. the relative robustness of phonetic precursors, remains a strong contender in explaining the underphonologisation of HV patterns as compared to HH patterns. To be sure, the ideal of isolating analytic bias by controlling for potential channel bias is a tall order. Moreton is to be saluted for laying out an interesting approach to this thorny issue and for attempting to offer concrete evidence for the role of analytic bias in shaping the typology of sound patterns. Nevertheless, at least in the case of the underphonologisation of HV over HH patterns, irrefutable evidence for analytic bias remains elusive.

AppendixSummary of F1 values measured at the mid-point of V₂ in /'adV₂CV₃/ sequences.

V ₂	V ₃		English		Turkish	
			Mean (Hz)	SD	Mean (Hz)	SD
a	a	voiced	629	104	733	85
		voiceless	645	94	718	103
		<i>total</i>	637	99	725	94
	e	voiced	630	104	693	116
		voiceless	637	94	718	80
		<i>total</i>	634	99	705	100
	i	voiced	603	106	684	139
		voiceless	618	124	683	123
		<i>total</i>	610	115	684	130
	<i>total</i>	voiced	620	105	702	117
		voiceless	633	103	706	105
		<i>total</i>	627	104	704	111
e	a	voiced	479	73	514	81
		voiceless	468	71	519	79
		<i>total</i>	473	72	516	80
	e	voiced	474	68	495	64
		voiceless	481	78	509	60
		<i>total</i>	478	73	502	62
	i	voiced	449	58	450	104
		voiceless	454	63	467	117
		<i>total</i>	451	60	458	110
	<i>total</i>	voiced	467	68	486	88
		voiceless	468	71	499	89
		<i>total</i>	468	69	493	89
i	a	voiced	336	43	351	37
		voiceless	336	42	356	37
		<i>total</i>	336	42	353	37
	e	voiced	328	37	348	31
		voiceless	322	37	355	39
		<i>total</i>	325	38	351	35
	i	voiced	315	33	355	41
		voiceless	323	36	358	44
		<i>total</i>	319	35	356	42
	<i>total</i>	voiced	326	39	351	37
		voiceless	327	39	356	39
		<i>total</i>	327	39	354	38

REFERENCES

- Beddor, Patrice Speeter, Anthony Brasher & Chandan Narayan (2007). Applying perceptual methods to the study of phonetic variation and sound change. In Maria-Josep Solé, Patrice Speeter Beddor & Manjari Ohala (eds.) *Experimental approaches to phonology*. Oxford: Oxford University Press. 127–143.
- Beddor, Patrice Speeter & Handan Yavuz (1995). The relation between vowel-to-vowel coarticulation and vowel harmony in Turkish. In Kjell Elenius & Peter Branderud (eds.) *Proceedings of the 13th International Congress of the Phonetic Sciences*. Vol. 2. Stockholm: KTH & Stockholm University. 44–51.
- Beddor, Patrice Speeter, James D. Harnsberger & Stephanie Lindemann (2002). Language-specific patterns of vowel-to-vowel coarticulation: acoustic structures and their perceptual correlates. *JPh* **30**. 591–627.
- Blevins, Juliette (2004). *Evolutionary Phonology: the emergence of sound patterns*. Cambridge: Cambridge University Press.
- Blevins, Juliette (2006). A theoretical synopsis of Evolutionary Phonology. *Theoretical Linguistics* **32**. 117–166.
- Boer, Bart de & Patricia K. Kuhl (2003). Investigating the role of infant-directed speech with a computer model. *Acoustic Research Letters Online* **4**. 129–134.
- Clayards, Meghan (2008). *The ideal listener: making optimal use of acoustic-phonetic cues for word recognition*. PhD thesis, University of Rochester.
- Clayards, Meghan, Michael K. Tanenhaus, Richard N. Aslin & Robert A. Jacobs (2008). Perception of speech reflects optimal use of probabilistic speech cues. *Cognition* **108**. 804–809.
- Denning, Keith (1989). *The diachronic development of phonological voice quality, with special reference to Dinka and the other Nilotic languages*. PhD dissertation, Stanford University.
- Dieth, Eugen (1932). *A grammar of the Buchan dialect (Aberdeenshire): descriptive and historical*. Vol. 1: *Phonology – Accidence*. Cambridge: Heffer.
- Evans, Bronwen G. & Paul Iverson (2007). Plasticity in vowel perception and production: a study of accent change in young adults. *JASA* **121**. 3814–3826.
- Feldman, Naomi H. & Thomas L. Griffiths (2007). A rational account of the perceptual magnet effect. In Danielle S. McNamara & J. Gregory Trafton (eds.) *Proceedings of the 29th Annual Cognitive Science Society*. Austin: Cognitive Science Society. 257–262.
- Feldman, Naomi H., Thomas L. Griffiths & James L. Morgan (2009). The influence of categories on perception: explaining the perceptual magnet effect as optimal statistical inference. *Psychological Review* **116**. 752–782.
- Fitzgerald, Colleen M. (2002). Vowel harmony in Buchan Scots English. *English Language and Linguistics* **6**. 61–79.
- Flemming, Edward (1995). *Auditory representations in phonology*. PhD dissertation, University of California, Los Angeles.
- Flemming, Edward (1996). Evidence for constraints on contrast: the dispersion theory of contrast. *UCLA Working Papers in Phonology* **1**. 86–106.
- Flemming, Edward (2001). Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* **18**. 7–44.
- Flemming, Edward (2002). *Auditory representations in phonology*. New York: Routledge.
- Flemming, Edward (2010). Modeling listeners: comments on Pluymaekers *et al.* and Scarborough. In Cécile Fougeron, Barbara Kühnert, Mariapaola D’Imperio & Nathalie Vallée (eds.) *Laboratory phonology 10*. Berlin & New York: De Gruyter Mouton. 587–605.
- Guion, Susan (1996). *Velar palatalization: coarticulation, perception and sound change*. PhD dissertation, University of Texas at Austin.

- Hale, Mark & Charles Reiss (2000). Phonology as cognition. In Noel Burton-Roberts, Philip Carr & Gerard Docherty (eds.) *Phonological knowledge: conceptual and empirical issues*. Oxford: Oxford University Press. 161–184.
- Harrington, Jonathan (2006). An acoustic analysis of ‘happy-tensing’ in the Queen’s Christmas broadcasts. *JPh* **34**. 439–457.
- Harrington, Jonathan, Sallyanne Palethorpe & Catherine Watson (2000). Monophthongal vowel changes in Received Pronunciation: an acoustic analysis of the Queen’s Christmas broadcasts. *Journal of the International Phonetic Association* **30**. 63–78.
- Hillenbrand, James, Dennis R. Ingrisano, Bruce L. Smith & James E. Flege (1984). Perception of the voiced–voiceless contrast in syllable-final stops. *JASA* **76**. 18–27.
- Hume, Elizabeth & Keith Johnson (2001a). A model of the interplay of speech perception and phonology. In Hume & Johnson (2001b). 3–26.
- Hume, Elizabeth & Keith Johnson (eds.) (2001b). *The role of speech perception in phonology*. San Diego: Academic Press.
- Hyman, Larry M. (1976). Phonologization. In Alphonse Juilland (ed.) *Linguistic studies offered to Joseph Greenberg on the occasion of his sixtieth birthday*. Vol. 2: *Phonology*. Saratoga: Anna Libri. 407–418.
- Hyman, Larry (2001). On the limits of phonetic determinism in phonology: *NC revisited. In Hume & Johnson (2001b). 141–185.
- Inkelas, Sharon, Jonathan Barnes, Andrew Dolbey, Jeff Good, Gunnar Hansson, Darya Kavitskaya, Orhan Orgun, Ronald Sprouse & Alan Yu (2001). The effects of stress on vowel-to-vowel coarticulation in Turkish. Paper presented at the 10th International Conference on Turkish Linguistics, Boğaziçi University, Istanbul.
- Kingston, John & Randy L. Diehl (1994). Phonetic knowledge. *Lg* **70**. 419–454.
- Kiparsky, Paul (2006). The amphichronic program *vs.* Evolutionary Phonology. *Theoretical Linguistics* **32**. 217–236.
- Kiparsky, Paul (2008). Universals constrain change; change results in typological generalizations. In Jeff Good (ed.) *Linguistic universals and language change*. Oxford: Oxford University Press. 23–53.
- Kirby, James (2010). *Cue selection and category restructuring in sound change*. PhD dissertation, University of Chicago, Chicago.
- Kirby, James (to appear). Acquisition of covert contrasts: an unsupervised learning approach. *CLS* **46**.
- Lee, Ernst W. (1977). Devoicing, aspiration, and vowel split in Haroi: evidence for register (contrastive tongue-root position). In David Thomas, Ernst W. Lee & Nguyen Dang Liem (eds.) *Papers in South East Asian linguistics No. 4: Chamic studies*. Canberra: Australia National University. 87–104.
- Lin, Ying (2005). *Learning features and segments from waveforms: a statistical model of early phonological acquisition*. PhD thesis, University of California, Los Angeles.
- Lindblom, Björn, Susan Guion, Susan Hura, Seung-Jae Moon & Raquel Willerman (1995). Is sound change adaptive? *Rivista di Linguistica* **7**. 5–36.
- McMurray, Bob, Richard N. Aslin & Joseph C. Toscano (2009). Statistical learning of phonetic categories: insights from a computational approach. *Developmental Science* **12**. 369–378.
- Majors, Tivoli (1998). *Stress-dependent harmony: phonetic origins and phonological analysis*. PhD dissertation, University of Texas, Austin.
- Manuel, Sharon Y. (1990). The role of contrast in limiting vowel-to-vowel coarticulation in different languages. *JASA* **88**. 1286–1298.
- Mayo, Catherine, James M. Scobbie, Nigel Hewlett & Daphne Waters (2003). The influence of phonemic awareness development on acoustic cue weighting strategies in children’s speech perception. *Journal of Speech, Language, and Hearing Research* **46**. 1184–1196.

- Moreton, Elliott (2008). Analytic bias and phonological typology. *Phonology* **25**, 83–127.
- Moreton, Elliott (2009). Underphonologization and modularity bias. In Steve Parker (ed.) *Phonological argumentation: essays on evidence and motivation*. London: Equinox. 79–101.
- Moreton, Elliott & Erik R. Thomas (2007). Origins of Canadian Raising in voiceless-coda effects: a case study in phonologization. In Jennifer Cole & José Ignacio Hualde (eds.) *Laboratory Phonology 9*. Berlin & New York: Mouton de Gruyter. 37–63.
- Ní Chiosáin, Máire & Jaye Padgett (2001). Markedness, segment realization, and locality in spreading. In Linda Lombardi (ed.) *Segmental phonology in Optimality Theory: constraints and representations*. Cambridge: Cambridge University Press. 118–156.
- Nittrouer, Susan (2002). Learning to perceive speech: how fricative perception changes, and how it stays the same. *JASA* **112**, 711–719.
- Nittrouer, Susan & Joanna H. Lowenstein (2009). Does harmonicity explain children's cue weighting of fricative-vowel syllables? *JASA* **125**, 1679–1692.
- Nittrouer, Susan & Marnie E. Miller (1997). Predicting developmental shifts in perceptual weighting schemes. *JASA* **101**, 2253–2266.
- Ohala, John J. (1983). The origin of sound patterns in vocal tract constraints. In Peter F. MacNeilage (ed.) *The production of speech*. New York: Springer. 189–216.
- Ohala, John J. (1993). The phonetics of sound change. In Charles Jones (ed.) *Historical linguistics: problems and perspectives*. London: Longman. 237–278.
- Padgett, Jaye (2003). Contrast and post-velar fronting in Russian. *NLLT* **21**, 39–87.
- Paster, Mary (2004). Vowel height harmony and blocking in Buchan Scots. *Phonology* **21**, 359–407.
- Sancier, Michele L. & Carol A. Fowler (1997). Gestural drift in a bilingual speaker of Brazilian Portuguese and English. *JPh* **25**, 421–436.
- Sankoff, Gillian (2004). Adolescents, young adults and the critical period: two case studies from 'Seven Up'. In Carmen Fought (ed.) *Sociolinguistic variation: critical reflections*. Oxford: Oxford University Press. 121–139.
- Scarborough, Rebecca A. (2004). *Coarticulation and the structure of the lexicon*. PhD thesis, University of California, Los Angeles.
- Seidl, Amanda & Eugene Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development* **1**, 289–316.
- Sonderegger, Morgan & Alan C. L. Yu (2010). A rational account of perceptual compensation for coarticulation. In Stellan Ohlsson & Richard Catrambone (eds.) *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Austin: Cognitive Science Society. 375–380.
- Summers, W. Van (1987). Effects of stress and final-consonant voicing on vowel production: articulatory and acoustic analysis. *JASA* **82**, 847–863.
- Trautmüller, Hartmut (1990). Analytic expressions for the tonotopic frequency scale. *JASA* **88**, 97–100.
- Vallabha, Gautam K., James L. McClelland, Ferran Pons, Janet F. Werker & Shigeaki Amano (2007). Unsupervised learning of vowel categories from infant-directed speech. *Proceedings of the National Academy of Sciences* **104**, 13273–13278.
- Watson, Jocelyne M. M. (1997). *Sibilant-vowel coarticulation in the perception of speech by children with phonological disorders*. PhD dissertation, Queen Margaret College, Edinburgh.
- Whalen, D. H. (1990). Coarticulation is largely planned. *JPh* **18**, 3–35.
- Wilson, Colin (2006). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science* **30**, 945–982.
- Yu, Alan C. L. (2010). Perceptual compensation is correlated with individuals' 'autistic' traits: implications for models of sound change. *PLoS One* **5:8**, e11950.

Available at <http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0011950>.

Yu, Alan C. L. (to appear). Individual differences in socio-cognitive processing and the actuation of sound change. In Alan C. L. Yu (ed.) *Origins of sound change: approaches to phonologization*. Oxford: Oxford University Press.

Zuraw, Kie (2007). The role of phonetic knowledge in phonological patterning: corpus and survey evidence from Tagalog infixation. *Lg* **83**. 277–316.