Phonology **25** (2008) 83–127. © 2008 Cambridge University Press
doi:10.1017/S0952675708001413 Printed in the United Kingdom

*Analytic bias and phonological typology**

Elliott Moreton
University of North Carolina, Chapel Hill

Two factors have been proposed as the main determinants of phonological typology: channel bias, phonetically systematic errors in transmission, and analytic bias, cognitive predispositions making learners more receptive to some patterns than others. Much of typology can be explained equally well by either factor, making them hard to distinguish empirically. This study presents evidence that analytic bias is strong enough to create typological asymmetries in a case where channel bias is controlled. I show that (i) phonological dependencies between the height of two vowels are typologically more common than dependencies between vowel height and consonant voicing, (ii) the phonetic precursors of the height-height and height-voice patterns are equally robust and (iii) in two experiments, English speakers learned a height-height pattern and a voice-voice pattern better than a height-voice pattern. I conclude that both factors contribute to typology, and discuss hypotheses about their interaction.

1 Introduction

Some phonological patterns are common across unrelated languages, while others are rare or non-existent. It must be the case that the common patterns either are innovated more often, or survive better from generation to generation. This paper addresses the two leading proposals as to the factors which determine innovation and survival rates. One is CHANNEL BIAS, phonetically systematic errors in transmission between speaker and hearer, caused largely by subtle phonetic interactions which serve as precursors for phonologisation (Ohala 1993, 2005, Hale & Reiss 2000, Barnes 2002, Blevins 2004). The other is ANALYTIC BIAS, cognitive

* I am indebted to many people for discussion of the ideas and facts presented here, including Adam Albright, Dani Byrd, Andries Coetzee, Paul de Lacy, Abigail Kaun, Shigeto Kawahara, John McCarthy, Steve Parker, Joe Pater, Amanda Seidl, Jennifer L. Smith, Paul Smolensky, Donca Steriade, Anne-Michelle Tessier, Rachel Walker, audiences at the University of North Carolina, the University of Southern California, Rutgers University, NELS 37 at the University of Illinois, WCCFL 29 at Berkeley, four anonymous reviewers and the Associate Editor. Thanks are also due to Chris Wiesen of UNC's Odum Institute for statistical advice. This research was supported in part by a grant from the University Research Council (UNC). Remaining errors are mine.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

biases which facilitate the learning of some phonological patterns and inhibit that of others. One hypothetical type of analytic bias, Universal Grammar, forms the basis for typological explanation in generative phonology.

Channel bias and analytic bias are often treated as mutually exclusive, either passively, by neglecting one factor, or actively, by arguing for the primacy of the other (see review in §2). I believe that this is a mistake, and that an adequate theory of typology will have to take both into account (Hyman 2001, Myers 2002, Kiparsky 2006). Towards that end, this paper presents new empirical evidence that selective pattern learning shapes typology in ways that cannot be explained by channel bias. Specifically, the study shows that phonological patterns relating the height of two or more neighbouring vowels outnumber patterns relating vowel height to consonant voicing (§3); that the phonetic precursor of the height-height patterns is not larger than that of the height-voice patterns (§4); and that a height-height pattern is learned better in a laboratory situation than a height-voice pattern (§5). A complementary question – whether *every* analytic bias corresponds to a typological asymmetry – is addressed with a second learning experiment in which a long-range voice-voice dependency is learned better than a height-voice pattern (§6). Three alternative explanations for the experimental results, based on the lexical statistics of English, are considered and rejected in §7. Concluding discussion is in §8, where the evidence of this study is used to argue that a two-factor theory of typology is necessary and feasible. Hypotheses as to how analytic and channel bias interact are proposed and discussed.

The principal novelty of this study is that it connects a specific typological asymmetry to a demonstrated analytic bias, while excluding channel bias as a cause. Previous laboratory studies of analytic bias have concentrated on analytic biases which mimic channel biases (phonetically 'natural' analytic biases), and so could not unambiguously identify the source of the typological bias. Previous studies which eliminated channel bias inferred an analytic bias, but did not demonstrate it directly in the laboratory.

2 Theoretical context

Phonology is acquired by a learner from a corpus of phonological representations received from other speakers. CHANNEL BIAS refers to systematic errors which cause the phonological representation received by the learner to differ from the one intended by the speaker. ANALYTIC BIAS refers to systematic predispositions in what a learner infers from the received representations (Wilson 2003b). The sources of bias are shown schematically in Fig. 1. In principle, either factor could lead to systematic drift in phonological systems as they are passed from one generation to the next, favouring the innovation or survival of particular patterns. Both

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

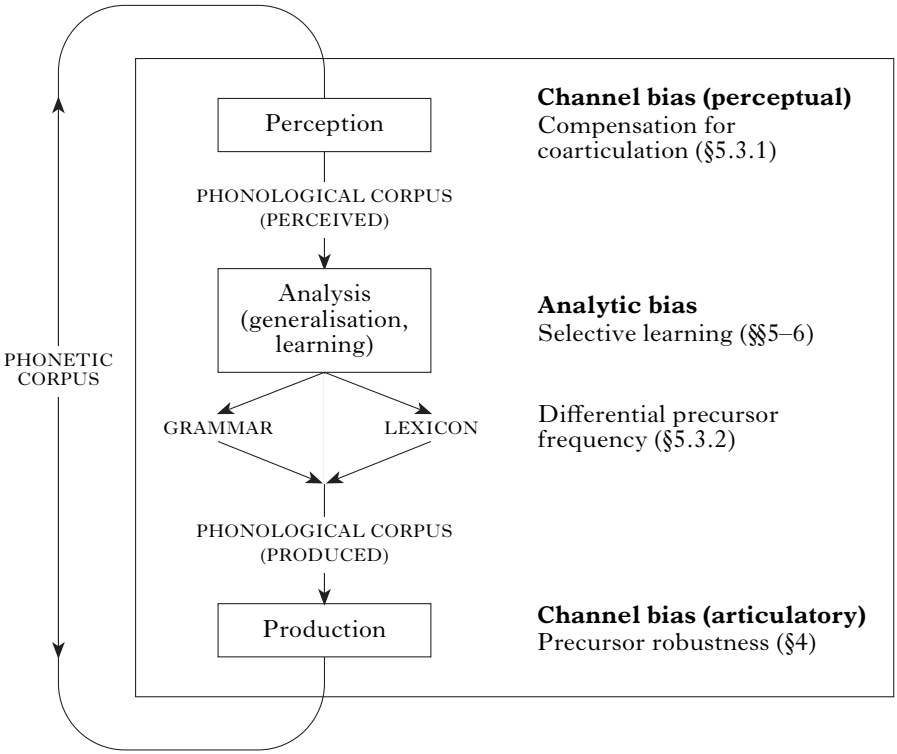


Figure 1

Factors influencing the innovation and survival of phonological patterns in intergenerational transmission. The large box encloses an individual speaker.

have been proposed as general, causal explanations for phonological typology.[1]

2.1 Deriving typology

In typological theories based on analytic bias, asymmetries between attested and unattested phonologies are attributed to cognitive predispositions which admit some phonological patterns and exclude others. For example, vowel-height harmony is common, while consonant-continuancy harmony is non-existent or nearly so (Hansson 2001: 137–149, Rose & Walker 2004). A typical analytic-bias account might run like this (adapted from Baković 2000: 4–6): Universal Grammar provides a constraint AGREE[high] against adjacent vowels that disagree in height, but no

[1] The two approaches are sometimes referred to as 'synchronic' (analytic bias) and 'diachronic' (channel bias), but this is misleading. On the one hand, analytic bias and channel bias both exist synchronically; on the other hand, the only way that any factor can affect typology is diachronically, through its impact on the innovation and retention rates.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

corresponding AGREE[cont] for consonants. The universal constraint set can therefore be ranked so as to enforce height harmony, but not continuancy harmony. Given training data which instantiated both patterns equally well, a learner would find continuancy harmony entirely unlearnable, or would acquire it slowly or imperfectly via the mechanisms used for idiosyncratic patterns. As a result, continuancy harmony would be less likely to be innovated, and more likely to be lost, than height harmony, leading to lower typological frequency.

Most proposals which use analytic bias to explain typology take that bias to be Universal Grammar (Chomsky & Halle 1968: 4, 251, 296–297, Archangeli & Pulleyblank 1994: 391–395, Clements & Hume 1995: 245, Steriade 2001b: 235–237, Davidson *et al*. 2004, Hayes & Steriade 2004: 1–2, 6). However, in other proposals, typologically effective analytic bias may also emerge from the interaction between Universal Grammar and a learning mechanism (Boersma 2004), or from cognitive biases which are not specifically linguistic (Saffran 2002, 2003, Newport & Aslin 2004). Thus, Universal Grammar is a kind of analytic bias, but there may be analytic biases other than Universal Grammar.

At the other end of the spectrum are approaches which aim to minimise the role of analytic bias by shifting the burden of typological explanation to properties of the communication channel between the speaker and hearer. In this view, Universal Grammar provides a cognitive framework that can represent a much larger range of phonological patterns than is found in nature. It may supply a universal set of representational units, or regularise phonetic variability, but does not otherwise favour one phonological pattern over another (Ohala 1990, 2005, Haspelmath 1999: 204–205, Hale & Reiss 2000, Hume & Johnson 2001a, Blevins 2004: 19–21, 41, 281–285). Instead, phonological typology is caused principally by systematic errors occurring in the transmission of phonological representations between the mind of a speaker and that of a learner (who induces a grammar from the erroneously received forms). Such an explanation for the rarity of continuancy harmony compared to height harmony might go as follows (after Ohala 1994b, Beddor *et al*. 2001, Blevins 2004: 142–144, Przezdziecki 2005): vowel-to-vowel height coarticulation is normally 'compensated' by perceptual mechanisms which allow the hearer to recover the intended phonological height, but sometimes compensation for coarticulation fails. When this happens, the listener perceives one vowel as having been phonologically assimilated to the other, and may use this perception as evidence to acquire a height-harmony process ('phonologisation'; Hyman 1976, Ohala 1993, Beddor *et al*. 2001). There is no such phonetic precursor for continuancy harmony, so continuancy harmony is rarely innovated. A learner exposed to equally good instantiations of both patterns would, one assumes, acquire them equally well.

Coarticulation and other patterns of phonetic covariation are hypothesised to be a major source of channel bias. Asymmetries in phonetic precursors introduce biases into the data available to learners, leading to more

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

frequent innovation of some sound patterns than others, and hence to asymmetries in phonological typology. This hypothesis is most often invoked to explain why a pattern which has a phonetic precursor is more frequent than another pattern which has none, but it also applies to patterns whose precursors differ in size: the more robust the precursor, the more opportunities arise for phonologisation, and hence the more frequent is the phonological pattern (Ohala 1994a, Hale & Reiss 2000, Barnes 2002: 151–159, Kavitskaya 2002: 123–133, Blevins 2004: 108–109). Other proposed sources of channel bias include differences in perceptual similarity between sounds (Ohala 1993), differences in auditory robustness of acoustic cues (Chang *et al.* 2001), and cognitive biases, specific to language, in how acoustic cues are parsed into phonological representations (Blevins 2004: 151–153).

2.2 Evidence and arguments

If typology can be explained by analytic bias, then analytic bias, properly understood, should fit snugly around typology (Przezdziecki 2005: 7–20). The main arguments against a general analytic-bias account of typology are based on typological data showing that no model of Universal Grammar can achieve this fit. One argument comes from 'crazy rules', the other from the 'too-many-solutions problem'. 'Crazy' (i.e. phonetically bizarre) rules are attested in nature as the result of a succession of phonetically transparent sound changes (Bach & Harms 1972, Anderson 1981). A theory of Universal Grammar which is liberal enough to admit crazy rules must also admit so many unattested processes that it can no longer make useful typological predictions. The 'too-many-solutions' problem occurs when optimality-theoretic factorial typology overpredicts the number of ways in which a markedness constraint can be satisfied (Steriade 2001a). Revisions to the theory of Universal Grammar have has some success (see Blumenfeld 2006 for a review); however, some of the missing processes have to date been explainable only by lack of a phonetic precursor. For example, the configuration (nasal) + (voiceless obstruent) is resolved in many ways, but never by epenthesis (Pater 2004). This fact has resisted UG-based explanation for ten years, but Myers (2002) has pointed out that the process lacks a robust phonetic precursor. Thus, current UG-based analytic-bias theories both overpredict and underpredict in ways that can be explained by channel bias.

A parsimony argument is also advanced against the UG-based theories. The most salient of all typological facts is that phonological patterns tend to be 'phonetically natural', in the sense that they resemble exaggerated or stylised expressions of some phonetic fact. UG-based theories rely on 'phonetically grounded' constraints to explain this typological asymmetry. Thus, in order to explain typology, facts already immanent in the phonetics have to be stated a second time in the characterisation of Universal Grammar, often in a way that implicitly describes a channel

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

bias. If the phonetics is admitted to cause a channel bias which can account for the observed typology, it is argued to be 'extravagant' (Ohala 2005) to hypothesise a phonetically informed Universal Grammar to do the same job a second time (Hale & Reiss 2000: 160, 162, Barnes 2002: 364–365, Blevins 2004: 81–85, 237).

In response, two arguments have been put forth in favour of analytic bias as a typological factor. The first is that, parsimoniously or not, analytic bias exists and resembles typology. In pattern-processing experiments, systematic 'naturalness' biases have been reported in what learners acquire and what they overlook (Schane *et al.* 1974, Saffran & Thiessen 2003, Wilson 2003a, b), how they generalise what they do acquire (Chambers *et al.* 2006, Wilson 2006), and what predispositions they have without training (Pertz & Bever 1975, Davidson *et al.* 2004, Mintz & Walker 2006, Berent *et al.* 2007, Moreton *et al.*, in press). While these findings defuse the parsimony argument, they do not remove the confound between analytic and channel bias, and so do not show a causal role in typology for analytic bias.

The second argument in favour of analytic bias in typology is that channel bias alone does not predict typology correctly. There are several ways in which this is true. First, there exist 'diachronic conspiracies', in which otherwise common sound changes fail to occur when the resulting grammar would violate a language universal. For instance, a language with final-obstruent voicing could in principle arise from intervocalic voicing followed by final-vowel deletion, but in fact never does. Sound change is blocked by some other factor, presumably analytic bias (Kiparsky 1995, 2008, Bermúdez-Otero 2006). A related point is that channel bias observed in perceptual experiments does not always predict the relative frequencies of sound changes occurring in nature, again suggesting that some sound changes are resisted or facilitated by analytic bias (Steriade 2001b).

Finally, some phonetic precursors seem to undergo phonologisation less often than others of similar magnitude ('underphonologisation'; Moreton, in press). Two cases have been described to date. (i) Vowel F0 is affected to about the same extent by the height of the vowel and by the voicing or aspiration of a preceding consonant, but phonological height-tone patterns are hard to find compared to voice-tone patterns (Hombert *et al.* 1979: 51–53). (ii) The effect on vowel F0 of consonant voicing and aspiration is about the same size as that of tone-to-tone coarticulation, but phonological voice-tone patterns are significantly rarer than tone sandhi affecting tone height (Moreton, in press). This, too, suggests that analytic bias may facilitate the learning of some phonetically 'natural' sound patterns over others.

This study asks whether analytic bias is strong enough to create typological asymmetries on its own, unassisted by precursor robustness. The point of departure is a new case of underphonologisation. In the next section, phonological 'height-height' (HH) patterns, defined as dependencies between the height of neighbouring vowels, are shown to be

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

more common than 'height-voice' (HV) patterns, defined as dependencies between the height of a vowel and the voicing, aspiration or fortis/lenis status of an immediately following consonant. Subsequent sections of the paper investigate the contributions of each of the factors identified in Fig. 1.

3 Typological asymmetry: height-height outnumbers height-voice

A pilot survey, encompassing a wide range of phonological and phonetic variables, was conducted to locate cases of underphonologisation. The pilot results suggested that HH patterns are typologically more frequent than HV patterns. An intensive survey was carried out to test this hypothesis. The survey consisted of a brute-force search of the descriptive grammars and secondary phonological literature available at Johns Hopkins University, the University of Southern California and the University of North Carolina at Chapel Hill, supplemented by a query on the LINGUIST e-mail list (2002). Only sources written in Germanic and Romance languages were accessible to the author.

In order to qualify for the survey, a language has to provide the opportunity for both HH and HV patterns to occur. Specifically, it had to have both a height contrast and a postvocalic voicing, aspiration or 'fortis-lenis' contrast. The language must have been described while still alive; reconstructions were excluded. For the purposes of the survey, an HH pattern was defined as a static phonotactic restriction or morphophonemic alternation in which the height of one vowel was predictable from that of another vowel across at least one intervening consonant. To be sure that the pattern involved height, rather than just being an idiosyncratic property of a particular phoneme, the pattern was required to involve at least two different vowels of the same height. An HV pattern was defined as an analogous dependency between the height of a vowel and the voicing, aspiration or 'fortis-lenis' status of an immediately following consonant. Allophonic (non-neutralising) patterns were excluded, since there was no way to distinguish them from especially robust phonetic precursors. The existence of lexical exceptions was construed as evidence that a pattern was *not* allophonic. Alternations limited to, or triggered by, a single affix did not qualify.

As a crude precaution against double-counting instances of shared inheritance, the survey counted language families rather than individual languages. 'Family', for the purpose of this survey, was defined as 'top-level category in Ethnologue' (Gordon 2005). The assumption is that in counting the language families in which living languages instantiate the HH and HV patterns, we are counting *surviving independent innovations* of those patterns, and thus approximating an answer to the question of whether HH patterns are likelier than HV patterns to be innovated or retained in the face of language change.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Survey results were divided into two tiers. The 'strict' tier consisted of those cases which fit the survey criteria perfectly. Cases which were partially defective in one of the survey criteria were relegated to the 'lax' tier. The results are shown in (1) and (2). For each family, the strongest case is cited; others are noted briefly if known to me.

(1) HH *patterns*

a. *Strict tier: 7 families*

Afro-Asiatic: Awngi
The nucleus of the last syllable of a nominal or verbal stem alternates between [e] and [i] depending on the following suffix. Nuclei of earlier syllables alternate between [e] and [i], or between [o] and [u], to match. Voicing contrast (Palmer 1959, Hetzron 1969: 8, 1997: 484–485). Height harmony is also found in Kera, but the voicing contrast may be redundant with tone (Ebert 1976, 1979: 14–18, Pearce 2003, 2005, personal communication).

Altaic: Udihe
In roots, non-high vowels agree in height and rounding. Suffixes harmonise to root. Voicing contrast (Nikolaeva & Tolskaya 2001: 50–51, 72–76).

Basque: Basque
In many dialects, /a/ raises to [e] after a syllable containing a high vowel. Voicing contrast (Hualde 1991: 10, 23–31).

Indo-European: Buchan Scots
Unstressed suffixal high vowels become non-high when preceded by a stressed non-high vowel. Certain consonants are blockers. Voicing contrast (Paster 2004). Numerous other height-assimilation patterns occur in the Romance languages (for reviews see Hualde 1989, Parkinson 1996, Walker 2005).

Niger-Congo: C'Lela
High and non-high vowels do not co-occur in roots. Suffixes alternate. Voicing contrast (Dettweiler 2000). Height harmony is very widespread in the Bantu branch of this family (Parkinson 1996, Hyman 1998).

Oto-Manguean: Maltinaltepec Tlapaneca
/a/ unrestricted, but vowels of non-final syllable are mid or high depending on whether the final vowel is mid. Voicing contrast (Suárez 1983: 7–9, 12–16, 20–22, 48–49).

Sino-Tibetan: Lhasa Tibetan
Non-high vowels become high in the presence of a high vowel. Aspiration contrast in stops (Dawson 1980: 3, 11–12, 63–80).

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

b. *Lax tier: 8 families*

Austronesian: Woleaian
/a/, the only low vowel, becomes [e] before a syllable containing [a], and also becomes [e] between two syllables containing high vowels. Voicing contrast marginal (only /ʂ/ *vs.* /ʐ/) (Sohn 1971, 1975).

Chukotko-Kamchatkan: Chukchee
/i u e/ lower to [e o a] when in same morphological constituent as /e o/ or some kinds of laterals. Voicing contrast marginal: /k/ *vs.* /g/ only (Bogoras 1922). Later authors describe the /g/ as [ɣ], making voicing redundant with continuancy (Kämpfe & Volodin 1995).

Dravidian: Spoken Tamil
/i u/ in a word-initial syllable do not occur before a singleton consonant followed by /a/ or /ai/; /e o/ occur instead. Voicing contrast marginal, only in loans (Asher 1985: 211–214, 229, Schiffman 1999: 19). Co-occurrence facts have been questioned on the basis of phonetic measurements (Keane 2001: ch. 4).

Gulf: Tunica
Mid vowels do not co-occur in underived lexical items. /e o/ lower to [ɛ ɔ] before /a/ in same morpheme. Voicing contrast marginal; mostly in loans (Haas 1946, Wiswall 1991: 82–125).

Hokan: Washo
The final vowels of certain prefixes are realised as [a] or [e] depending on whether the first stem vowel is /a o/ or /i ɨ u e/. The same rule determines the vowel in an epenthesis process. Voicing contrast. Since [e] in this language is phonetically higher than [o], it is not clear that the pattern is conditioned by height (Jacobsen 1964: 52–54, 300–302, 305–306).

Korean
In ideophones, 'dark' /e y ə u/ do not co-occur with 'light' /æ œ a o/. Numerous dark/light pairs with identical consonants, but vowels differing by one height step exist. Since they have augmentative/diminutive meanings, it is unclear that this would pass the single-affix test (McCarthy 1983, Sohn 1986).

Nilo-Saharan: Murle
/ɛ o ɔ/ raise to [e u o] before a voiced stop followed by /i u/, or in some cases /e o/. Voicing contrast. Productivity and phonological status doubtful (Arensen 1982: 19, 134, and examples *passim*).

Penutian: Wintu
In a 'very large' number of verb roots, /e o/ raise to [i u] before a singleton consonant followed by /a/ (Pitkin 1984: 43–45). Voicing and aspiration contrast. Productivity uncertain.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

(2) HV *patterns*

a. *Strict tier: 0 families*

b. *Lax tier: 3 families*

Indo-European

(i) Polish

/ɔ/ raises to [o] before underlyingly voiced non-nasal coda. Productivity is doubtful (Sanders 2003).

(ii) Canadian English

[ʌɪ] and [aɪ] contrast before [ɾ], but in other environments [ʌɪ] is found only before voiceless obstruents and [aɪ] is found only elsewhere. Contrast is marginal (Chambers 1973).

Nilo-Saharan: Murle

See (1b) above.

Sino-Tibetan: Lungtu Fujien Chinese

Stops contrast for aspiration in onset. In codas, voiced stops occur after non-low vowels, voiceless stops after low vowels. Coda voiced/voiceless redundant with preglottalised/glottalised, and not phonemically contrastive (Egerod 1956: 27–51).

HH patterns outnumbered HV patterns by 7 to 0 in the strict survey, and 15 to 3 in the lax one. If their true frequencies were the same, half of the cases found should have been HH and half HV. This null hypothesis was tested using a two-sided exact binomial test, and was decisively rejected for both the strict and the lax survey ($p < 0{\cdot}016$ and $p < 0{\cdot}008$ respectively using the *binom.test* function of the *stats* package of the R statistical software, R Development Core Team 2005).[2] We have thus identified a previously unremarked typological asymmetry: vowel height interacts more often with vowel height than with consonant voicing.[3]

4 Channel bias does not favour height-height patterns

Given that the typological asymmetry exists, the question arises of whether channel bias provides an explanation. The high typological frequency of vowel harmony has been ascribed to channel bias caused by its phonetic precursor, vowel-to-vowel height coarticulation (Ohala 1994b,

[2] I am indebted to Chris Wiesen, of the Odum Institute for Research in the Social Sciences at the University of North Carolina, Chapel Hill, for suggesting this analysis.

[3] By counting families rather than languages, we have if anything *under*stated the extent of the asymmetry. Expanding the 'strict' survey to include multiple representatives of each family can only increase the HH count, since no HV cases at all were found. Naturally, this procedure is only a heuristic, adapted for convenience to deal efficiently with a large amount of data. It is in the end no substitute for the careful historical scholarship required to establish e.g. which of the Bantu lowering rules are in fact independent innovations.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Blevins 2004: 143, Przezdziecki 2005), and it has been proposed in various contexts that weaker precursors lead to less phonologisation (Ohala 1994a, Barnes 2002: 151–159, Kavitskaya 2002: 123–133, Myers 2002, Blevins 2004: 108–109, Moreton & Thomas 2007). If that explanation is correct, then the phonetic precursor of the HV pattern should be smaller than that of the HH pattern.

4.1 Survey: HV precursor is not smaller than HH precursor

To test whether this is so, we must first identify the phonetic precursor of the HV pattern, and then compare its magnitude to that of height coarticulation. The HV patterns appear to have two sources. One is the tendency for vocalic articulations to be exaggerated before voiceless obstruents (Thomas 2000, Moreton 2004, Moreton & Thomas 2007); the other is the pharyngeal cavity expansion which occurs during the production of voiced obstruents (for a review, see Thomas 2000). Both of these phonetic interactions lead to a slightly lower vowel F1 before a voiceless obstruent than before a voiced one. A survey was carried out 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.

The survey proceeded as follows. Studies were found in which vowel F1 was measured in the relevant contexts. Among the contexts used in the study, two were identified which were deemed likeliest to raise or lower target-vowel F1. 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. The effect of context was defined to be the target-vowel F1 in the Raising context divided by the target-vowel F1 in the Lowering context. This procedure automatically normalises away interspeaker differences in vocal tract length (Thomas 2000). Some studies reported measurements at different points in the target. Where that was the case, the point closest to the context was used. For example, if the study measured F1 at the target vowel's onset and offset, then the onset measurement was used when estimating the effect of preceding /i/ *vs.* /a/ context, and the offset measurement was used when estimating the effect of following /i/ *vs.* /a/ context. Survey results are plotted in Fig. 2, and given in detail in Tables I and II. A ratio of 1 (solid horizontal line) indicates no effect of context, while values greater than 1 signify a higher F1 (lower vowel) in the Lowering context.

The smaller precursor hypothesis is not confirmed: there is no evidence that the HH precursor is larger than the HV precursor; if anything, the reverse is true. This finding adds a third case of underphonologisation to the two that are already known. In all three cases, differences in phonological typology exist without corresponding differences in precursor robustness. Hence, it is not in general true that precursor robustness

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

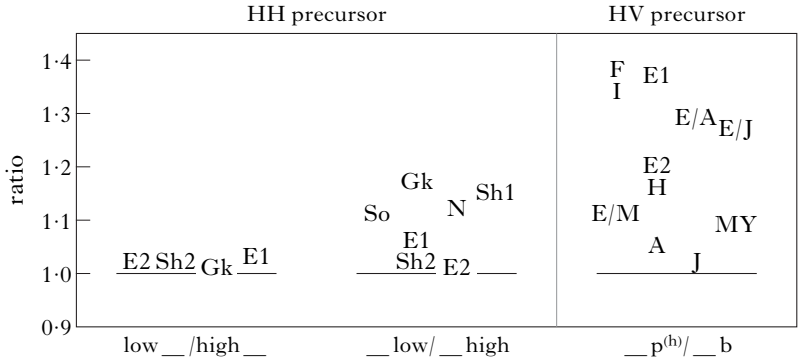


Figure 2

Ratio of F1 in the Raising context to F1 in the Lowering context. A value of 1 corresponds to no effect (see text for explanation). The points in each group have been randomly dispersed on the horizontal axis to avoid overlapping.

predicts typological frequency (*contra* the suggestion of Archangeli & Pulleyblank 1994: 178–179). Since precursor robustness is the only kind of channel bias that is relevant to these cases, it follows that channel bias does not in general predict typological frequency (this claim is further defended below, in §5.3).

5 Experiment 1: height-height *vs*. height-voice

The previous section showed that channel bias is not a plausible explanation for the typological preponderance of HH over HV patterns. Can analytic bias do better? In particular, is the HH pattern easier to learn?

Patterns of segmental occurrence and co-occurrence can be acquired by learners in laboratory experiments. In a typical such experiment, participants are familiarised with a set of stimuli that conform to a particular pattern, then tested on novel stimuli which may or may not conform. In adults, pattern conformity affects phoneme restoration (Ohala & Feder 1994), speech errors (Dell *et al*. 2000, Goldrick 2004), speeded-repetition latency (Onishi *et al*. 2002, Chambers *et al*. 2006, Koo & Cole 2006) and segmentation of continuous speech (Newport & Aslin 2004, Bonatti *et al*. 2005), as well as allomorph selection in an artificial language (Schane *et al*. 1974, Pycha *et al*. 2003, Wilson 2003a, b) and language-game responses (Wilson 2006). In infants, pattern-conformity effects are found in preferential listening paradigms (Chambers *et al*. 2003, Saffran & Thiessen 2003, Seidl & Buckley 2005).

Experiment 1 used a learning paradigm to compare learning of HH and HV patterns. Participants were familiarised with an instantiation of one or the other pattern by practising pronouncing 'words' of an artificial 'language' instantiating the pattern, and were then asked to distinguish

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

code	study	ratio
E1	English (Beddor *et al*. 2002): 5 speakers Stressed /i e a o u/ measured at target offset: [__ Ca] *vs*. [__ Ci] measured at vowel onset: [aC __] *vs*. [iC __]	 1·06 1·03
E2	English (Koening & Okalidou 2003): 3 speakers Stressed /i e ɑ ɔ u/, measured at steady state [__ Ca] *vs*. [__ Ci] [aC __] *vs*. [iC __]	 1·01 1·02
Gk	Greek (Koening & Okalidou 2003): 3 speakers Stressed /i ɛ a ɔ u/, measured at steady state [__ Ca] *vs*. [__ Ci] [aC __] *vs*. [iC __]	 1·17 1·01
N	Ndebele (Manuel 1990): 3 speakers /e a/, measured at target offset [__ Ca] *vs*. [__ Ci]	 1·12
Sh1	Shona (Manuel 1990): 3 speakers /e a/, measured at target offset [__ Ca] *vs*. [__ Ci]	 1·15
Sh2	Shona (Beddor *et al*. 2002): 7 speakers Stressed /i e a o u/ measured at target offset: [__ Ca] *vs*. [__ Ci] measured at target onset: [aC __] *vs*. [iC __]	 1·02 1·02
So	Sotho (Manuel 1990): 3 speakers /e a/, measured at target offset [__ Ca] *vs*. [__ Ci]	 1·11

Table I
Phonetic effect of context vowel height on target vowel F1.

new 'words' from non-'word' foils, i.e. to distinguish stimuli conforming to the pattern from stimuli violating it. If analytic bias favours the HH pattern over the HV pattern in nature, then we might expect participants to show better performance in the HH Condition in the lab. On the other hand, if participants' performance is better in the HV Condition, that would be evidence against the hypothesis that analytic bias favours the HH pattern in nature.

5.1 Method

5.1.1 *Design*. The 'words' used in the artificial 'languages' had the phonological structure $C_1V_1C_2V_2$, where C_1 and C_2 were drawn from the

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

code	study	ratio
A	Arabic (de Jong & Zawaydeh 2002) Stressed /a/, measured at mid point [__t] *vs.* [__d]	1·05
E1	English (Wolf 1978): 2 speakers /æ/, average F1 in last 30 ms [__p t k] *vs.* [__b d g]	1·37
E2	English (Summers 1987): 3 speakers /ɔ ɑ/, measured at vowel offset [__p f] *vs.* [__b v]	1·20
E/A	L2 English (L1 Arabic) (Crowther & Mann 1992): 10 speakers /ɑ/, measured at vowel offset [__t] *vs.* [__d]	1·29
E/J	L2 English (L1 Japanese) (Crowther & Mann 1992): 10 speakers /ɑ/, measured at vowel offset [__t] *vs.* [__d]	1·27
E/M	L2 English (L1 Mandarin) (Crowther & Mann 1992): 10 speakers /ɑ/, measured at vowel offset [__t] *vs.* [__d]	1·11
F	French (Fischer-Jørgensen 1972): 1 speaker /a/, measured just before closure [__p t k] *vs.* [__b d g]	1·38
H	Hindi (Lampp & Reklis 2004): 5 speakers /ɔ/, measured just before closure [__k] *vs.* [__g]	1·16
I	Italian (Vagges *et al.* 1978): 10 speakers /a/, measured at closure [__p t k f s t͡s ʧ] *vs.* [__b d g v z d͡z ʤ]	1·34
J	Japanese (Kawahara 2005): 3 speakers /e a o/, measured just before closure [__p t k] *vs.* [__b d g]	1·02
MY	Mòbà Yoruba (Przezdziecki 2005): 1 speaker /i/, measured at midpoint [__t k] *vs.* [__d g]	1·09

Table II

Phonetic effect of context consonant voicing on target vowel F1.

set /t d k g/, and V_1 and V_2 from the set /i u æ ɔ/. The CVCV shape was chosen with an eye to future experiments, because it is the smallest unit within which nucleus-to-onset, nucleus-to-nucleus and onset-to-onset dependencies could occur. Within these limits, 256 'words' were possible.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

A 'word' was defined as HH-CONFORMING if V_1 and V_2 were both phonologically high (/i u/) or both phonologically non-high (/æ ɔ/). It was HV-CONFORMING if V_1 and C_2 were respectively high and voiced, or non-high and voiceless. Consequently, there were 64 'words' that were both HH- and HV-conforming, 64 that were HH- but not HV-conforming, 64 that were HV- but not HH-conforming and 64 that were neither HH- nor HV-conforming.

For each participant, a unique set of 32 HH-conforming 'words' was randomly chosen for use in the FAMILIARISATION PHASE of the HH Condition, subject to the constraint that each of the eight permitted V_1V_2 combinations occur in four 'words', and each of the 16 permitted V_1C_2 combinations occur in two 'words'. Another set of 32 HH-conforming words, disjoint from the first one, was randomly chosen for use as positive test items in the TEST PHASE of the HH Condition. An analogous procedure was followed to choose 32 familiarisation stimuli and 32 positive test items for the HV Condition, with each of the 16 permitted V_1V_2 combinations occurring in two 'words' and each of the eight permitted V_1V_2 combinations occurring in four 'words' in each of the two lists. Finally, the 64 'words' that were neither HH- nor HV-conforming were randomly assigned to the HH and HV Conditions as negative test items, subject to the requirement that the eight permitted V_1V_2 combinations and eight permitted V_1C_2 combinations occur in four 'words' each. No 'word' occurred in both Phases of the same Condition, or in both Conditions. All familiarisation items in a given Condition conformed to the relevant pattern, and were 50% likely to conform to the other pattern; the same was true for the positive test items. The negative test items in both Conditions were HH- and HV-non-conforming, to make the Conditions as similar as possible. All participants were familiarised and tested in both Conditions, with even-numbered participants receiving the HH Condition first.

5.1.2 *Stimuli.* Stimuli were synthesised using the MBROLA diphone concatenative synthesiser (Dutoit *et al.* 1996), using 'US 3' voice (a male speaker of American English). Each 'word' was synthesised individually. The nominal duration parameters for both consonants were set to 100 ms, while those for both vowels were set to 225 ms, with 150 ms of silence initially and finally. Intonation was left at the default monotone of 123 Hz. In order not to perturb the natural intensity difference between high and low vowels, no amplitude normalisation was applied.

5.1.3 *Procedure.* All participants were tested individually in a double-walled soundproof chamber (Ray Proof Corporation, Norwalk, Connecticut, Model AS-200) using a Macintosh iBook G4 laptop computer (Apple Computer Corporation) under the control of software written for this experiment in Java 2, Version 1.4.2_09 (Sun Microsystems). Participants received oral instructions from the experimenter, recapitulated by detailed written instructions on the computer screen. These

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

instructions are reproduced in the Appendix. The instructions stated that the experiment was 'about learning to recognize words in an artificial language', and that it would consist of a 'study phase' (i.e. familiarisation) in which they practised pronouncing individual 'words' of the language, followed by a 'test phase' in which they would be tested on how well they could recognise them. No indication was given at any time as to whether or not words from the familiarisation phase would recur in the test phase. At the beginning of each phase, a message box appeared to remind participants of the procedure for that phase. The experimenter stayed with the participant through the first 5–10 familiarisation trails, then left the soundproof chamber, and was not present during the rest of the experiment.

On each familiarisation trial, the computer played a single 'word' to the participant through binaural mono headphones, which the participant was to repeat back into a head-mount microphone attached to the headphones (Altec Lansing). Participants were instructed to 'match the pronunciation as closely as possible', and told that their pronunciations would be recorded. A large button labelled 'Next' was permanently visible on the screen; mouse-clicking it after the end of one trial started the next. Presentation rate was thus under participant control, and no instructions were given as to speed. One familiarisation block consisted of one trial for each of the 32 familiarisation stimuli, in random order. The familiarisation phase contained four such blocks.

On each test trial, participants heard one positive and one negative test item, separated by 450 ms (i.e. the 150 ms of MBROLA-synthesised silence after the offset of the first test item, followed by a 150 ms pause, followed by the 150 ms of MBROLA-synthesised silence preceding the onset of the second test item). Buttons labelled '1' and '2' were permanently visible on the screen, and participants were instructed to mouse-click 'the one that you think was in the language you studied … . If you can't tell, make your best guess'. The buttons remained inactive until the second stimulus had finished playing; thereafter, clicking either button initiated the next trial.

When the test phase of the first condition finished, a message box on the computer screen told the participant that it was time for a break, and 2–3 minutes of instrumental music was played over the headphones. When the music ended, the break continued until the participant was ready to proceed with the familiarisation phase of the second condition.

5.1.4 *Participants*. Twenty-five participants were recruited from the community at the University of North Carolina at Chapel Hill. Their average age was 20·7 years (SD = 3·2). All reported English as their first language and normal hearing; one reported a speech condition (stuttering). Three were natively bilingual (Estonian, Korean, Kru). All had studied a foreign language (Spanish 17; French 7; Latin 5; Italian 3; Arabic 2; Chinese, Japanese, Luganda, Portuguese and Swahili 1 each). Participants were paid US$7 for the experiment, which lasted about half

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

		0 (HV Condition)		1 (HH Condition)	
		HH-non-conformity		HV-non-conformity	
Same-Vowel	Order	0 (vowels agree in height)	1 (vowels disagree in height)	0 (V1 high iff C2 voiced)	1 (V1 high iff C2 voiceless)
0 (V1≠V2)	0 (1st half)	[tidu] (8) 50·0	[tidæ] (16) 55·7	[tidu] (8) 67·7	[titu] (8) 63·5
	1 (2nd half)	53·1	55·7	70·8	57·3
1 (V1=V2)	0 (1st half)	[tidi] (8) 57·3	*impossible*	[tidi] (8) 75·0	[titi] (8) 56·3
	1 (2nd half)	50·0		51·0	63·5

Table III

Design and results of Experiment 1. A typical positive test item is shown in each cell, along with the raw percentage of correct responses. Parenthesised numbers show how many positive test items were in that cell.

an hour. Data from one participant was lost due to equipment failure, leaving 24 valid participants.

5.2 Results and discussion

Subject responses were analysed using a mixed-effects logistic-regression model in which the dependent variable was the probability of choosing the test item that was consistent with the familiarisation pattern.

All of the independent variables (representing the factors whose effects were to be tested) were binary. Condition was 0 for test trials in the HV Condition, and 1 for test trials in the HH Conditions, thus making the HV Condition the reference category. The reason for this choice was that Experiment 2 also had an HV Condition, but no HH Condition. HH-non-conformity was 0 for test trials in which the positive test item was HH-conforming (i.e. those in which the vowels agreed in height), and 1 for those in which the positive test item was HH-non-conforming. Likewise, HV-non-conformity was 0 when the positive test item was HV-conforming (i.e. when the first vowel was high and the second consonant voiced, or when the first vowel was low and the second consonant voiceless). The negative test item was in every instance both HH- and HV-non-conforming. Since positive test items in the HH Condition were always HH-conforming, and those in the HV Condition were HV-conforming, this meant that HH- and HV-conformity were nested within Condition. The variable Same-Vowel was 1 when the positive test item had the exact same vowel twice (e.g. [tugu]), and 0 when the two vowels

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

differed (e.g. [tugi]). Since only HH-conforming items could have the same vowel twice, this variable was nested within HH-non-conformity. Negative test items, being HH-non-conforming, never had the same vowel twice. Finally, Order was 0 for test trials which occurred in the first half of the experiment (before the musical break), and 1 for those which occurred in the second half. The complete fourteen-cell design, together with typical positive test items and the raw percentage of correct responses in each cell, is shown in Table III.

The statistical analysis proceeded by stepwise reduction from an initial saturated model, which was guaranteed to fit the data perfectly. The initial model included fixed-effect terms for the main effect of each of the independent variables, as well as all possible interactions up to redundancy (since it was not known in advance which ones would matter). There were a total of fourteen fixed-effects terms, saturating the fourteen cells of the design. A random effect was included for subject intercepts to absorb within-subject variability. The model was fit by maximum likelihood using the *lmer* function in the *Matrix* library of the statistical software package R (R Development Core Team 2005). Parameter estimates for the fixed effects are shown along with their standard errors and significance levels in Table IV.

The model was reduced by stepwise deletion of non-significant terms, beginning with the highest-order interactions and, within the interactions, with the numerically smallest coefficients, subject to the restriction that a lower-order term could only be deleted after the deletion of all higher-order terms in which it occurred. Each reduced model was compared to the initial saturated model, using analysis of variance. Reduction stopped when the next reduced model would have differed significantly from the saturated model, using a criterion of $p < 0{\cdot}25$ to err on the side of retaining rather than eliminating terms. This procedure yielded a reduced model, shown in Table V, with ten terms. The reduced model did not differ significantly from the saturated model by an analysis-of-variance test ($\chi^2 = 4{\cdot}2514$ on 11 degrees of freedom, $p = 0{\cdot}3730$).

The intercept term in the final model was small, but greater than zero, indicating that the positive test item was chosen with greater than chance frequency by participants in the baseline HV Condition (non-significant; $p = 0{\cdot}301$). A numerically larger and highly significant main effect of Condition meant that the probability of choosing the positive test item was greater in the HH Condition than in the HV Condition ($p < 0{\cdot}002$; also significant in the original saturated model at $p = 0{\cdot}018$).

Two interactions reached the usual statistical-significance criterion of $p \leq 0{\cdot}05$. One, Order × Same-Vowel, reflected the fact that when the HV Condition came second, participants were less likely to choose positive test items in which the same vowel occurred twice – perhaps because they had heard many such items as familiarisation stimuli in the HH Condition, and associated them with 'the other language'. The second, Order × HV-non-conformity × Same-Vowel, cancels out both the Order × Same-Vowel interaction just mentioned and the sizeable but non-significant

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

variable	coefficient	SE	z	P(>|z|)
(Intercept)	−0·000	0·215	0·000	1·000
Condition	0·744	0·313	2·375	0·018 *
HH-non-conformity	0·231	0·251	0·921	0·357
HV-non-conformity	−0·186	0·305	−0·609	0·543
Same-Vowel	0·295	0·291	1·015	0·310
Order	0·126	0·304	0·415	0·678
Condition × Same-Vowel	0·064	0·434	0·148	0·882
HV-non-conformity × Same-Vowel	−0·665	0·437	−1·521	0·128
Order × Condition	0·021	0·465	0·046	0·963
Order × HH-non-conformity	−0·126	0·356	−0·345	0·723
Order × HV-non-conformity	−0·411	0·432	−0·952	0·341
Order × Same-Vowel	−0·421	0·411	−1·026	0·305
Order × Condition × Same-Vowel	−0·789	0·604	−1·305	0·192
Order × HV-non-conformity × Same-Vowel	1·779	0·610	2·917	0·004 **

Table IV
Experiment 1: initial saturated model.[4]

HV-non-conformity × Same-Vowel term when those test items were also HV-conforming, i.e. when the items which had one vowel twice also fit the pattern of the HV Condition (the one that the participants had just been familiarised on).

[4] As mentioned above, this is a logistic-regression model, in which the coefficients represent effect magnitudes in terms of logarithms of odds ratios (natural logarithm of the effect of that factor on the odds of a correct response). Here is an example of how it works. Suppose the cell we are interested in is the one in which the participant has been familiarised in the HH Condition (Condition = 1) during the first half of the experiment (Order = 0), with the positive test item being HV-conforming (HV-non-conformity = 1) and having vowels which agree in height (HH-non-conformity = 0) and are identical (Same-Vowel = 1) – the [titi] cell in Table III. To predict the probability that the participant chooses the positive rather than the negative test item in such a case, we first add together the coefficients for each term of the model for which all factors are equal to 1: Condition (0·744), HV-non-conformity (−0·186), Same-Vowel (0·295), Condition × Same-Vowel (0·064) and HV-non-conformity × Same-Vowel (−0·665), plus the Intercept term (in this case 0), which is included in all cells. This yields 0·252, which is the model's predicted log-odds of the probability of a correct response in that cell, corresponding to a predicted probability of 56·3%. The actual probability of a correct response is shown in Table III; it is 56·3%. (The predicted and actual probabilities are identical because the model is saturated.) For reasons why logistic regression is superior to older techniques such as analysis of variance (with or without the arcsine transformation), see Macmillan & Creelman (1990).

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

variable	coefficient	SE	z	P(>|z|)
(Intercept)	0·135	0·131	1·033	0·301
Condition	0·675	0·216	3·127	0·002 **
HV-non-conformity	−0·252	0·281	−0·894	0·371
Same-Vowel	0·219	0·192	1·143	0·253
Order	0·125	0·184	0·676	0·499
HV-non-conformity × Same-Vowel	−0·525	0·353	−1·488	0·137
Order × Condition	−0·259	0·328	−0·791	0·429
Order × HV-non-conformity	−0·129	0·392	−0·328	0·743
Order × Same-Vowel	−0·665	0·267	−2·491	0·013 *
Order × HV-non-conformity × Same-Vowel	1·234	0·497	2·483	0·013 *

Table V
Experiment 1: final reduced model.

These results – poor performance in the HV Condition and superior performance in the HH Condition – are consistent with the hypothesis that the HH pattern is learned more readily than the HV one, and thus provide support for the position that a cognitive bias is responsible for the underphonologisation of the HV pattern relative to the HH pattern in natural language. However, there are other possible interpretations which we must deal with first.

Since participants in the HH Condition heard only HH-conforming positive test items, while participants in the HV Condition heard a mix of HH-conforming and HH-non-conforming test items, the superior performance in the HH Condition might have had nothing to do with learning in the experiment, being due instead to a pre-existing preference for HH-conforming stimuli. If that had been the case, participants in the HV Condition would have been more likely to choose the positive test item when it was HH-conforming, and the statistical analysis would have found a negative effect of HH-non-conformity. However, no such effect was found. The coefficients in the saturated model associated with HH-non-conformity and its interaction with Order did not survive the elimination process, and in any case had the wrong sign.

A second alternative has to do with the fact that in half of the Familiarisation and positive Test stimuli in the HH Condition, the same vowel occurred twice (e.g. in the [titi] and [tidi] cells in Table III). In the HV Condition, only one-quarter of the Familiarisation and positive Test stimuli had two identical vowels (the [tidi] cell). Perhaps participants in the HH Condition did not learn to recognise stimuli whose vowels agreed

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

in height, but merely those whose vowels were identical. If that were true, however, we would have found an interaction of Condition × Same-Vowel (i.e. better performance on same-vowel stimuli when familiarised and tested in the HH Condition) instead of a main effect of Condition. Likewise, if performance had been better in the HH Condition because of a pre-existing preference for repeated vowels or rhyming syllables, we would have found a main effect of Same-Vowel rather than one of Condition.

A third alternative possibility is that participants did not detect the HV pattern because they misperceived the intended voicing of the medial consonant. The vowels used in this experiment were longer, more intense and acoustically more stable than the consonants, with the result that the HH pattern may have been supported by better-quality acoustic cues than the HV pattern. Previous research shows that this scenario is not impossible: in a study of CVC confusions in multi-talker babble noise, it was found that about 60–65 % of the information carried by vowel height was transmitted at all signal-to-noise ratios (0 dB, 8 dB and 16 dB). A similar proportion of the information carried by consonant voicing was transmitted at high SNR, but for initial consonants it fell to about 40 % at an SNR of 0 dB (Cutler *et al*. 2004).

To assess how accurately consonant voicing was perceived, the audio productions of participants from the familiarisation phase were examined. Each of the 24 speakers produced four repetitions each of 32 familiarisation stimuli in each of two pattern conditions, for a total of 6144 utterances. A subset of 500 recorded trials was selected randomly, assigned unique but meaningless identifying codes, and put in random order. The experimenter examined each one by ear and as an oscillogram and spectrogram using the Praat software (Boersma & Weenink 2005), and transcribed as much of the utterance as possible. Of the 500 recordings, 364 contained an entire C_2 (the other 136 consisted mainly of cases in which the participant had clicked the 'Next' button before finishing the utterance, or in which a faulty microphone had recorded no signal or an insufficient signal). The experimenter's transcriptions were then compared with the stimuli played to the participants. The two disagreed in voicing in four cases out of 364 (1·1 %), and some of these cases may have been due to the experimenter misperceiving the participant's utterance, rather than the participant misperceiving the stimulus. In no case was a non-high stimulus vowel (in either vowel position) produced as high, and in only one case out of 375 was a high stimulus vowel produced as non-high (/i/ produced as /æ/).

Perception of voicing may therefore have been slightly worse than that of height, but both features were perceived with high accuracy. Moreover, it has been found that phonotactic learning effects can persist in the face of small amounts of contrary evidence. Chambers *et al*. (2006) used a simultaneous train-and-test design in which conforming and non-conforming test items were interspersed amongst (conforming) training items. Although 10 of the 35 items (28·6 %) in some blocks of

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

their experiment violated the experimental phonotactic pattern, there was no difference in performance between blocks which contained test items and blocks which did not. Hence, it is not likely that the differences between the HH and HV Conditions in the present experiment were due to relatively worse perception of voicing than height in the stimuli.

5.3 Alternative explanations for underphonologisation

In sum, participants' superior performance in the HH Condition shows better learning of the HH experimental pattern than the HV pattern. The results are particularly striking in light of evidence from other sources that dependencies between phonetically adjacent segments are more salient than more remote relationships (Moreton & Amano 1999, Creel *et al*. 2004, Newport & Aslin 2004). If the same bias operates in natural language acquisition, it could produce the observed typological skew in favour of the HH pattern in natural language. It is tempting to conclude that this is indeed what happened, and hence that cognitive biases can shape typology. Before we can take this step, there are two alternative hypotheses that must be dealt with.

5.3.1 *Perceptual distortion of precursors*. Acoustic measurements of the precursors may not accurately reflect their perceptual magnitudes. In the HV precursor, the two coarticulated segments are adjacent, whereas in the HH precursor some time passes between them. Suppose that compensation for coarticulation takes place within a shorter time window than does coarticulation itself. Then compensation would be less reliable for the HH precursor, leading to a higher rate of phonologisation. The suspicion that this is indeed what is happening is bolstered by the observation that in the case of effects on tone, typological frequency seems to increase with distance: tone-tone interactions (between neighbouring vowels) are more common than voice-tone interactions (between a vowel and an adjacent consonant), which are more common than interactions between a tone and the height of the vowel on which it is realised. The hypothesis has not yet been tested directly; however, there are two indirect arguments that it is not correct.

The first has to do with the nature of compensation for coarticulation. Compensation occurs when the perception of a feature on a potential target of coarticulation is influenced by potential triggers of coarticulation; e.g. when a phonetically nasalised vowel is perceived as less nasal in the environment of a nasal consonant. The perceptual influence appears to have two sources. One source is linguistic, and is sensitive to the coarticulatory patterns of the perceiver's native language (Beddor & Krakow 1999, Beddor *et al*. 2002, Darcy *et al*., in press). The other is auditory, not specific to humans or to speech, and sensitive to spectral similarity between trigger and target (for a review, see Lotto & Holt 2006). Spectral contrast can have long-range effects; e.g. categorisation of

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

a syllable as [ga] or [da] can be influenced by a 70 ms sine-wave tone occurring 1·3 seconds previously (Holt 2005). Since vowels are maximally similar to other vowels but maximally different from obstruent consonants, it is likely that vowel-to-vowel coarticulation is compensated for in both ways, but vowel-'voice' interaction in only the first. This should, if anything, lead to superior compensation for the HH precursor.

The second argument is typological. Suppose that compensation for coarticulation does, in fact, have a shorter range than coarticulation itself. Then most of the uncompensated coarticulation should occur at an intermediate distance from the coarticulatory trigger, in the zone between the (narrow) limits of compensation and the (wide) limits of coarticulation. Phonologisation of the uncompensated coarticulation would result in bizarre patterns. Coarticulation of lip rounding, for example, may anticipate the phonologically rounded segment by up to half a second (Lubker & Gay 1982). If coarticulatory rounding is removed from the closest segments by compensation, but remains uncompensated on the more distant neighbours, phonologisation could create a process that spreads rounding but skips over the vowel nearest the source, e.g. /ɯhɯ+ku/→[uhɯku]. Similarly, in a V_1CV_2 sequence, where V_1 is coarticulated with V_2, compensation should be best for that portion of V_1 which is closest to V_2. Phonologisation of the uncompensated coarticulation would lead to a diphthongising vowel-harmony pattern in which only the initial portion of V_1 changed to match V_2, e.g. /eː+hi/→[iehi]. Since these patterns are not (to my knowledge) found in nature, the hypothesis is unlikely to be true.

5.3.2 *Differential within-language precursor frequency*. The statistical properties of individual natural languages may afford speakers more opportunities to observe one precursor than the other, making its phonologisation more likely. The HV precursor can only be observed in sequences consisting of a vowel and an obstruent, whereas the HH precursor can only be observed when two vowels of different height occur in adjacent syllables. Is the HH context more frequent than the HV context across languages?

A definitive answer would require a database of corpus (token) frequencies in a large genetically and geographically balanced sample of languages, something which does not now exist. However, a database of lexical (type) frequencies in a small genetically and geographically balanced sample does exist, in the form of the UCLA Lexical and Syllabic Inventory Database (ULSID), and an approximate answer to our question can be constructed on the basis of the analysis of Rousset (2004).

The languages used by Rousset are a subset of those in ULSID: Agar, Finnish, Kannada, Kanuri, Kwakw'ala, Navajo, Ngizim, Nyah Kur, Quechua, Sora, Thai, Wa, Yup'ik and !Xóõ, plus French and Swedish. All of them have a voicing contrast in obstruents (either stops or fricatives, but not necessarily both). The data underlying the study is in the form of

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

	proportion of syllable types in lexicon					
	open syllables			closed syllables		
language	CV	CCV	V	VC	CVC	CCVC
Afar	64	0	3	5	29	0
Finnish	58	0	6	3	33	0
French	54	11	8	2	18	3
Kannada	76	0	3	2	18	0
Kanuri	60	1	3	1	33	1
Kwakw'ala	65	0	2	1	26	0
Navajo	59	0	0	0	40	0
Ngizim	73	1	3	1	22	0
Nyah Kur	23	1	0	0	57	18
Quechua	58	0	3	3	35	0
Sora	43	0	7	4	45	0
Swedish	33	4	6	8	34	6
Thai	28	3	0	0	64	5
Wa	19	3	1	3	61	14
Yup'ik	43	0	9	6	41	0
!Xóõ	81	0	5	0	14	0
mean	52	2	4	2	36	3

Table VI

Occurrence (%) of the six most common syllable types in Rousset's sample (retabulated from Rousset 2004: 115, Table III.8).

syllabified lexica, with recent loanwords excluded. Syllabification was based on either editorial judgements in published lexica, or on native speaker judgements collected by the database compliers (Rousset 2004: 53). On tabulating the lexical frequencies of different syllable types, Rousset found that, on average, 99% of them fell into the six categories CV, CCV, V, VC, CVC and CCVC. The data is shown in Table VI.

First, we estimate p_{HV}, the probability that a vowel will be followed by an ordinary voiced or voiceless obstruent of the type surveyed in Table II (i.e. not ejective, implosive, prenasalised, etc., not [h] or [ʔ], and not a sonorant). We adopt certain simplifying assumptions: we ignore the 13 syllable types which account for the remaining 1% of the lexica (CCCVC, VCCC, etc.), and we estimate the discourse (corpus, token) frequency of a syllable type by the lexical frequency of that syllable type. Also, we assume that the discourse is long enough that we can ignore the complication of the final syllable, and calculate as if every syllable were followed by another syllable.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

language	obstruents	all consonants	proportion
Afar	11	17	65
Finnish	10	17	59
French	13	21	62
Kannada	16	27	60
Kanuri	12	22	55
Kwakw'ala	21	43	49
Navajo	19	38	50
Ngizim	20	37	54
Nyah Kur	13	30	43
Quechua	19	33	58
Sora	18	51	35
Thai	13	22	59
Wa	17	37	46
Yup'ik	27	40	68
mean			55

Table VII

Proportion (%) of ordinary obstruents in the inventories of the languages in the sample (data from Rousset 2004: 58–71).

Under these assumptions, 58% of syllables are open. Of these, 4% are followed by V, while the rest are followed by an onset consonant. Thus 56% of vowels are followed by an onset consonant, 2% by a vowel and the remaining 42% by a coda consonant. According to Rousset (2004: 127), on average 96% of the consonants in a language's inventory can appear in the onset, so we will assume for simplicity that the discourse frequency of ordinary voiced or voiceless obstruents in onset position is equal to their proportion of the inventory. Inventory statistics, given in Table VII, indicate that about 55% of onset consonants are ordinary obstruents.[5]

In the languages of Table VII, only about 68% of all inventory consonants appeared in coda position (Rousset 2004: 127, Table III.12), and no information is given about which ones are codas in which languages. We are told only that [p t k ʔ s m n ŋ l] are by far the most frequent (Rousset 2004: 128). If we assume that the proportion of all codas which are ordinary obstruents is the same as the proportion of ordinary obstruents in that set, we arrive at an estimate of 44%.

[5] Of the 129 consonants of !Xóõ, 83 are clicks, and 13 of the rest have double releases. The majority of the 145 vowels are accompanied by secondary laryngeal, pharyngeal or glottal articulations. Since nothing is known about phonetic height-height or height-voicing interactions under these circumstances, !Xóõ was omitted from the calculations. Rousset (2004) gives no data for Swedish.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

	Vowels in inventory, by height					
language	H	HM	M	ML	L	p_{HH}
Afar	14		4		2	64
Finnish	6		4		2	61
French	3	3		6	2	70
Kannada	6		5		4	66
Kanuri	2		2		1	64
Kwakw'ala	2		2		1	64
Navajo	5		8		4	64
Ngizim	4		4		2	64
Nyah Kur	6	6		4	3	73
Quechua	2		2		1	64
Sora	6	5	4		2	72
Thai	6		4		6	66
Wa	3	3		2	1	72
Yup'ik	4				2	44
mean						65

Table VIII

Proportion (%) of vowel heights in the inventories of the languages in the sample (data from Rousset 2004: 58–71). Schwa and diphthongs are excluded.

Putting the pieces together, we find that the proportion of vowels which are followed by an obstruent is $(0{\cdot}56 \times 0{\cdot}55) + (0{\cdot}42 \times 0{\cdot}44) = 49\,\%$. This is our estimate for p_{HV}. As for p_{HH}, the probability that a vowel will be followed by a vowel of a different height, we assume that all vowels in a language's inventory occur with equal frequency. The relevant inventory statistics, given in Table VIII, yield an estimate of 65 % for p_{HH}.

Since the foregoing analysis depended on the questionable assumption that the segments making up a language's consonant or vowel inventory are all equally frequent, p_{HV} and p_{HH} were estimated a second time, using an opportunistic sample of 15 languages for which within-language phoneme-frequency counts were available. Here, the simplifying assumption was that all syllables are CV; i.e. no attempt was made to distinguish between coda and onset inventories. The results are shown in Table IX Averaged across the entire sample, p_{HV} is 54 %, and p_{HH} is 66 %.

Both approximations, arrived at using different data and assumptions, agree that in a long utterance of n syllables, the HV precursor can be expected to occur about $0{\cdot}50\ n$ times, and the HH precursor about $0{\cdot}65\ n$ times. There is indeed a difference in favour of the HH precursor, but it is not a large one. Considering the extreme nature of the typological skew in favour of HH, it is quite unlikely that within-language difference

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

language	corpus type	p_{HV}	p_{HH}	source
Austronesian				
Chamorro	lexicon	56	72	Seiden (1960)
Indonesian	text	50	66	Altmann & Lehfeld (1980: 165)
Samoan	text	38	61	Sigurd (1968)
Sea Dyak	text	54	62	Altmann & Lehfeld (1980: 202)
Indo-European				
Bengali*	text	56	65	Sigurd (1968)
Czech	text	74	64	Altmann & Lehfeld (1980: 139)
English	text	47	73	Sigurd (1968)
Swedish*	text	52	72	Sigurd (1968)
Niger-Congo				
Ewe*	speech, text	72	64	Bole-Richard (1983: 90)
Swahili	text	47	62	Gakuru *et al.* (n.d.)
Amharic	text	44	63	Bender (1974)
Finnish	speech	52	66	Vainio (1996)
Georgian	text	55	67	Altmann & Lehfeld (1980: 124)
Japanese	text	57	64	Tamaoka & Makioka (2004)
Kaiwa	text	53	67	Sigurd (1968)
mean		54	66	

Table IX

Estimated probability of occurrence (%) of HH and HV precursors in CVCV… utterances, based on within-language phoneme frequencies. All languages were analysed by the cited authors as having three degrees of vowel height, except those marked with *, which have four. All have an obstruent-voicing contrast.

in precursor frequency are the sole cause, or even the main cause, of HV/HH underphonologisation, though it may be a contributing factor.

6 Experiment 2: height-height *vs.* voice-voice

Our results so far have shown that there is a typological asymmetry favouring HH over HV patterns, that this asymmetry does not reflect a difference in the robustness of the phonetic precursors and that the HH pattern is learned more readily in a laboratory situation. These results clearly favour analytic bias over precursor robustness as an explanation for the underphonologisation of HV patterns relative to HH ones.

If the fit between Universal Grammar and natural language typology is very snug, then the set of easily learned patterns should be the same as the set of typologically common patterns. This is the situation we would expect if Universal Grammar is the *only* important factor shaping typology.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

The results of Experiment I could then be explained as a consequence of UG's support for vowel-height harmony, as discussed in §2. In that case, we would expect that an experimental VV pattern, in which the two consonants of the disyllabic stimulus agreed in voicing, would enjoy no learning advantage over the HV experimental pattern, since the VV pattern, like the HV pattern, is typologically very rare (Hansson 2004, Rose & Walker 2004).

On the other hand, analytic bias may favour the HH pattern over the HV one in some more general way. It could be that patterns taking place on a single autosegmental tier are easier to learn than those involving two tiers (Newport & Aslin 2004), or that patterns involving a single feature are easier than those involving multiple features (Chomsky & Halle 1968: 334–335, Clements & Hume 1995, Gordon 2004, Moreton, in press). In these cases, a VV experimental pattern should be learned better than the HV pattern, and some other factor would have to be responsible for the rarity of naturally occurring VV patterns.

6.1 Method

This experiment followed the same procedure as Experiment 1 in all respects except the construction of the artificial 'languages', where voicing agreement between the two consonants replaced height agreement between the two vowels. Twenty-seven volunteers participated (average age 20·4). One was natively bilingual (Korean); three others had some early-childhood foreign language exposure (German, Indonasian, Spanish). All had studied a foreign language (French 12; Spanish 12; Latin 6; Mandarin Chinese 3; Ancient Greek, German and Italian 2 each; Hebrew, Japanese, Portuguese and Russian 1 each). Results from three participants were discarded: in two cases, the software crashed after the musical break; in the third, the participant consciously noticed the HH pattern and responded exactly backwards, choosing the HH-disharmonic item on every trial.

6.2 Results and discussion

The same analysis procedure was followed as for Experiment 1. The design and raw response probabilities are shown in Table X, the initial (saturated) model in Table XI and the reduced model in Table XII. The final model did not differ significantly in fit from the saturated initial model by an analysis-of-deviance test ($\chi^2 = 1{\cdot}587$ on 4 degrees of freedom, $p = 0{\cdot}811$).

The intercept term in the final model was significantly greater than zero by the conventional 5% criterion, indicating that participants in the HV Condition chose the HV-conforming test item with greater than chance probability. This contrasts with the results of Experiment 1, in which the intercept had a smaller magnitude and missed significance. The effect of Condition was positive and significant, indicating better performance in

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

		0 (HV Condition)		1 (VV Condition)	
		VV-non-conformity		HV-non-conformity	
Same-Consonant	Order	0 (consonants agree in voicing)	1 (consonants disagree in voicing)	0 (V1 high iff C2 voiced)	1 (V1 high iff C2 voiceless)
0 (C1, C2)	0 (1st half)	[gidi] (8) 54·1	[kidi] (16) 57·3	[kidi] (8) 65·8	[kiti] (8) 59·1
	1 (2nd half)	55·0	53·8	51·8	43·6
1 (C1=C2)	0 (1st half)	[didi] (8) 50·0	*impossible*	[didi] (8) 55·6	[titi] (8) 68·1
	1 (2nd half)	43·3		57·0	51·5

Table X

Design and results of Experiment 2. A typical positive test item is shown in each cell, along with the raw percentage of correct responses. Parenthesised numbers show how many positive test items were in each cell.

variable	coefficient	SE	z	P(>|z|)
(Intercept)	0·167	0·184	0·908	0·363
Condition	0·489	0·266	1·835	0·066
VV-non-conformity	0·127	0·234	0·541	0·588
HV-non-conformity	−0·285	0·267	−1·066	0·287
Same-Consonant	−0·167	0·299	−0·560	0·576
Order	0·033	0·273	0·121	0·904
Condition × Same-Consonant	−0·266	0·427	−0·622	0·534
HV-non-conformity × Same-Consonant	0·818	0·438	1·869	0·062
Order × Condition	−0·616	0·386	−1·597	0·110
Order × VV-non-conformity	−0·177	0·347	−0·510	0·610
Order × HV-non-conformity	−0·044	0·381	−0·115	0·909
Order × Same-Consonant	−0·302	0·444	−0·680	0·497
Order × Condition × Same-Consonant	0·540	0·623	0·868	0·386
Order × HV-non-conformity × Same-Consonant	−0·308	0·622	−0·495	0·621

Table XI

Experiment 2: initial saturated model.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

variable	coefficient	SE	z	P(>|z|)
(Intercept)	0·260	0·107	2·429	0·015 *
Condition	0·393	0·169	2·325	0·020 *
HV-non-conformity	−0·112	0·153	−0·732	0·464
Same-Consonant	−0·309	0·179	−1·722	0·085
Order	−0·326	0·145	−2·246	0·025 *
Order × Condition	−0·460	0·217	−2·119	0·034 *
HV-non-conformity × Same-Consonant	0·675	0·264	2·556	0·011 *

Table XII
Experiment 2: final reduced model.

the VV Condition than the HV Condition; however, the coefficient was somewhat smaller, and its significance level much lower, than had been found for the HH Condition in Experiment 1. A negative main effect of Order shows that performance in the second half of the experiment was worse than that in the first, and the Order × Condition interaction means that this effect was especially pronounced when the second half of the experiment was the VV Condition. Finally, the HV-non-conformity × Same-Consonant interaction shows a strong tendency in both Conditions to choose positive test items which had a low vowel between identical voiced consonants, or a high vowel between identical voiceless ones.

Although these results are not as clear-cut as those of Experiment 1, they do suggest that there is an analytic bias favouring the VV experimental pattern over the HV experimental pattern, although both are typologically scarce. It follows that analytic bias need not entail a typological asymmetry, or, to put it a different way, that analytic bias is not the only important factor determining typological frequency.

I do not know why VV patterns are typologically rarer than HH ones. The difference may be due to analytic bias, since a direct experimental comparison found a VV pattern somewhat harder to learn than an HH pattern (Moreton, unpublished data). However, the VV pattern also seems to lack a robust phonetic precursor. The only positive report of which I know is that of Beardsley & Cullinan (1987), who found that five-year-old English-learning children have longer positive VOTs (i.e. less voicing) for initial /p/ in *pick* than in *pig* (by about 6 % in isolation and 16 % in a frame sentence) and longer negative VOTs (i.e. more voicing) for initial /b/ in the nonsense /boʊk/ than in /boʊg/ (by about 19 % in isolation). On the other hand, Weismer (1979), in a study of English CVC monosyllables produced by adults, found long-distance voicing *dissimilation*: the VOT of an initial voiceless stop was about 7 % longer when the final consonant was a voiced stop than when it was a voiceless one. In another study of English-speaking adults, Port & Rotunno (1979) found that VOT of initial /p t k/

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

was shorter by about 13–20% when the syllable was /CVpt/ than when it was /CVn/; however, it is not clear that the effect was due to the change in voicing rather than syllabic or morphological complexity, for example. Finally, an adult-English study by Port (1981) measured the closure duration (typically longer in voiceless stops) of the initial /d/ in /dVC/, /dVCV/ and /dVCVCV/ words. It was not significantly affected by the voicing of the next consonant (numerically, it was 1% longer in the voiceless context, averaged across all six conditions of the experiment). I know of no work on this topic outside of English. There is some evidence as well that the HH pattern is learned faster than the VV one in the lab (Moreton, in preparation), suggesting that analytic bias may play a role as well.

7 Could English phonology explain the experimental results?

One further alternative hypothesis remains for us to deal with: if the effects found in Experiments 1 and 2 are caused by experience of English, they are irrelevant to typology, and the argument collapses. The most direct way for English phonology to contaminate the results would be if participants came to the experiment predisposed to choose HH- and VV-conforming test items; i.e. they were trained by exposure to English rather than to the familiarisation items. The experiments were designed to test for that possibility by looking for effects of HH- and VV-conformity in the HV Condition. None were found (see §§5.2 and 6.2 above).

However, English could also have had an indirect effect, by facilitating *learning* of the HH and VV patterns in the familiarisation phase. The experiments did not test this possibility, but we can check its plausibility by asking whether there is anything in the corpus statistics of English to make HH and VV patterns easier to learn than HV ones. This cannot be done without a concrete hypothesis about the learner to tell us the right way to count. Three such hypotheses were tested.

The first is that the English learner acquires a gradient phonotactic constraint which prefigures the absolute constraint of the experimental pattern. It has been proposed that such gradient constraints are acquired when natural classes co-occur more or less often than would be expected if they were independent (Frisch *et al*. 2004: 215–216). The relevant co-occurrence statistics were extracted from the CELEX database of British English (Baayen *et al*. 1995a). Words with zero corpus frequency were excluded. Different inflected forms of the same stem were counted as different words. Obstruents were classified as voiced ([b d g v ð z ʒ ʤ]) or voiceless ([p t k f θ ʃ x ʧ]), vowels as high ([iː uː ɪ ʊ ɪə ʊə eɪ əʊ]) or non-high ([ɛ ʌ ə ɜː ɔː ɛə æ ɒ ɑː ɐ̃ ɑ̃ æ̃ː ɔ̃]), on the basis of the CELEX transcriptions (Baayen *et al*. 1995b: 4.25–4.26). The diphthongs [aɪ aʊ ɔɪ] were omitted, as their height was ambiguous. For each of the three patterns (HH, HV, VV), both conforming and non-conforming instances were counted.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

	corpus type			
	combined written and spoken		spoken only	
pattern	equally weighted	frequency-weighted	equally weighted	frequency-weighted
HH	0·99	1·02	0·96	1·01
HV	1·05	0·97	1·05	1·00
VV	1·02	1·04	1·02	1·08

Table XIII
Ratio of observed to expected frequency of pattern-conforming instances in the English lexicon.

An HH-conforming instance was defined as two high vowels or two low vowels separated only by consonants and prosodic symbols; a non-conforming instance was a high and a low vowel, in either order, separated only by zero or more consonants (of any sort, not just obstruents) and prosodic symbols. An HV-conforming instance was defined as a high vowel followed by a voiced obstruent, or a low vowel by a voiceless one, separated only by zero or more prosodic symbols; a non-conforming instance was defined similarly, with 'high' and 'low' interchanged. A VV-conforming instance was defined as two obstruents, both voiceless or both voiced, occurring initially in two successive syllables; a non-conforming instance was the same, except that the two obstruents disagreed in voicing. A single segment could participate in more than one instance of the same pattern, e.g. as the second vowel in an HH-conforming instance and the first vowel in an HH-non-conforming one. A single word could contribute multiple instances of a pattern. Separate counts were made from the entire CELEX corpus (17·9 million word tokens) and from the spoken-English subcorpus (1·3 million), and were tabulated both with and without frequency weighting.

To test the gradient-constraint hypothesis, the tabulated frequencies were used to find the marginal probabilities (e.g. that the first of two successive vowels will be high), which were then multiplied to yield the expected frequency of conforming instances assuming independence. The results are shown in Table XIII. The observed/expected ratios are in every case close to 1, regardless of corpus or weighting. There is thus no clear difference between the HH and VV patterns, on the one hand, and the HV pattern, on the other, in the degree of support which they receive in the English lexicon.

A second possibility is that English learners might acquire a 'covert ranking' (Davidson *et al*. 2004) between constraints which are inactive in English, but are crucially ranked in the experimental grammar, so that the original English ranking is closer to the HH and VV rankings than to the

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

	corpus type			
	combined written and spoken		spoken only	
pattern	equally weighted	frequency-weighted	equally weighted	frequency-weighted
HH	51,416	3,133,331	16,204	193,506
HV	47,925	4,902,618	14,741	319,060
VV	17,896	92,479	5,188	54,623

Table XIV
Occurrence of words containing at least one non-conforming instance in the English lexicon (CELEX).

HV ranking. That could happen if HV-non-conformity were more frequent in English than HH- and VV-non-conformity, and the learner incrementally demoted initially high-ranked constraints against each of the three patterns as each non-conforming datum was encountered (Boersma & Hayes 2001). In simulations with the Gradual Learning Algorithm (Boersma 1998), the unit of learning data is typically the word; hence, I counted words containing at least one HH-, HV- or VV-non-conforming instance. Table XIV shows the results. Contrary to hypothesis, the learner would encounter VV-non-conforming words much less often than HH- or HV-non-conforming ones, whereas HH- and HV-non-conforming words are similar to each other in frequency. (Very similar results are obtained if we count individual instances of non-conformity within a word, rather than the non-conformity of whole words.)

A third way in which experience of English might explain the experimental results is if the structure of the English lexicon makes HH- or VV-conforming familiarisation items especially memorable, hence especially effective in influencing responses during the test phase. CVC nonsense words with dense English lexical neighbourhoods are recalled better than those with sparse ones (Roodenrys & Hinton 2002, Storkel *et al*. 2006). It is not known whether the same holds for CVCV non-words, but let us assume for the sake of argument that it does, and check whether the HH- and VV-conforming experimental items have more neighbours than the HV-conforming ones. Following the studies just cited, two words were treated as neighbours if their segmental representations differed by at most one insertion, deletion or substitution. Average lexical neighbourhood size in CELEX was computed over all HH-conforming experimental stimuli, all HV-conforming ones and all VV-conforming ones. Words whose CELEX corpus frequency was zero were excluded. Table XV shows that, regardless of corpus type or frequency weighting, the HH-conforming items have the *smallest* neighbourhoods, while the VV-conforming items have the largest, contrary to expectation.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

	corpus type			
	combined written and spoken		spoken only	
pattern	equally weighted	frequency-weighted	equally weighted	frequency-weighted
HH	0·9	4·1	0·3	5·3
HV	1·2	12·1	0·6	14·1
VV	1·3	12·4	0·6	14·3

Table XV

Average neighbourhood size for pattern-conforming experimental stimuli in the English lexicon (CELEX). Frequency-weighted counts are weighted by occurrences per million words in the specified corpus.

In none of these three ways – frequency of non-conforming words, ratio of observed to expected conforming instances and size of neighbourhood – do the lexical statistics of English favour the HH and VV patterns over the HV pattern. There are many other ways to count, and perhaps some of them would find such a bias (though I know of none that do). However, any alternative statistical proposal based on such a bias would have to be neither ad hoc nor post hoc, but motivated by a theory of the learner and the task – a theory which would also have to explain why that same bias did *not* induce a preference for HH- and VV-conforming items in the HV Condition. In the interim, I conclude that the analytic biases observed in the experiment were not acquired from experience with English. The issue can only settled in the end by testing speakers of different languages.

8 General discussion

8.1 Summary of empirical results

Phonological height-height patterns are typologically more frequent than height-voice patterns. This asymmetry is not attributable to a difference in the magnitudes of their phonetic precursors. It is also not well explained by differences in the effect of compensation for coarticulation, nor by differences in within-language frequency of occurrence of the two precursors. Experiment 1 found a learning bias which, if it operates in nature the same way it did in the lab, could produce the observed typological excess of height-height over height-voice patterns.[6] This finding agrees

[6] It is not known at this point how lab-learned phonotactics relates to natural language phonotactics. Artificial phonotactic restrictions can be learned very quickly in a lab situation, in a matter of tens of trials, and easily changed in response to a change in training data (Taylor & Houghton 2005, Chambers *et al.* 2006). Natural language phonotactic restrictions are so resistant to change that they often

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

with other studies which have found cognitive analogues of typological asymmetries such as onset preference and coda avoidance (Schane *et al.* 1974), sonority sequencing (Pertz & Bever 1975, Berent *et al.* 2007, Moreton *et al.*, in press), patterns of assimilation (Wilson 2003a, b, Davidson *et al.* 2004), vowel harmony (Mintz & Walker 2006) and implicational relations in palatalisation (Wilson 2006). Where the new results go beyond the old is in equating, rather than manipulating, the phonetic 'naturalness' of the phonological patterns.

What is the specific *content* of the analytic bias responsible for the HH/HV asymmetry? We are in a position to evaluate several hypotheses. It cannot be a bias for phonetically natural patterns over phonetically unnatural ones, since both the HH and HV patterns are phonetically natural in the sense of having robust phonetic precursors. It also cannot be the case that analytic bias favours exactly those patterns which are typologically frequent (or, equivalently, that analytic bias can be reliably inferred from typology), since Experiment 2 found evidence that long-range voice-voice dependencies can be learned more readily than height-voice dependencies, even though the typological frequencies of the two patterns are both very low. A third possibility is that repetitions of the exact same segment are favoured over other patterns, but this is not supported by the results of either experiment.

The most interesting remaining possibilities have to do with the featural symmetry of the HH, VV and HV patterns. On the one hand, analytic bias might favour within-tier (vowel-to-vowel or consonant-to-consonant) dependencies over between-tier dependencies. On the other, it might favour single-feature dependencies over those involving two different features. The former possibility is contradicted in the present case by evidence that a height-height dependency is learned better than a dependency between the height of one vowel and the backness of another (Moreton, in preparation), leaving the latter as the most promising direction for future research. It leads to a number of interesting questions, among them: does the phonetic content of the features matter (e.g. is height-place treated differently from height-voice)? Is there a general relationship between the difficulty of a pattern and the number of features involved? What kinds of learning algorithm *make* a single-feature dependency easier to learn than a two-feature one (Moreton, in press)?

8.2 Theoretical implications

I know of no author who explicitly denies the *existence* of channel or analytic bias. Where opinions differ is in the emphasis placed on each as an

cause illegal stimuli to be misperceived as legal ones (e.g. Hallé *et al.* 1998). There is as yet little evidence that artificial phonotactics can affect segmental perception (but see Ohala & Feder 1994). Further research is clearly needed. This study provides some of that further research by investigating whether short-term phonotactic learning resembles natural phonotactics in what kinds of patterns it favours.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

effective factor in creating the kind of typological differences that linguists typically confront – differences in frequency between common processes, and others that are minimally different from them. The position that analytic bias is typologically ineffectual has been stated most clearly by Haspelmath (1999: 204–205):

> This does not, of course, mean that there is no UG, no innate mental organ that is specialized for linguistic skills. Clearly, there are universal properties of language that probably cannot be derived from constraints on language use, e.g. the fact that grammars generally do not contain numerical specifications (e.g. 'a word may be at most 15 segments long'); or indeed the fact that humans use fairly rigid grammatical rules to begin with, rather than arranging morphemes in a random way and leaving interpretation to pragmatics … But these features of language are so general that they have little to do with the grammarian's everyday work.

The present results tell against that hypothesis. There is nothing formally outlandish about the HV phonological pattern. It is just as ordinary, from a featural perspective, as, for example, the widespread ban on postnasal voiceless obstruents (Pater 2004). The conclusion that follows from the present results is that analytic bias, all by itself, is capable of creating non-trivial asymmetries without assistance from channel bias.

A somewhat weaker hypothesis is that analytic bias is not involved in that most striking of all typological facts, the predominance of phonetically 'natural' phonological patterns over phonetically 'unnatural' ones. This proposal is often stated as a parsimony argument (e.g. Hale & Reiss 2000: 162, Blevins 2004: 52). However, we have just seen evidence that analytic bias can affect typology, and there is elsewhere evidence that humans have analytic biases which involve phonetic substance or 'naturalness' (Schane *et al*. 1974, Pertz & Bever 1975, Saffran & Thiessen 2003, Wilson 2003a, b, 2006, Davidson *et al*. 2004, Chambers *et al*. 2006, Mintz & Walker 2006, Berent *et al*. 2007, Moreton *et al*., in press). If analytic bias can affect typology when 'naturalness' is *not* an issue, as in the present study, it is reasonable to think that it can affect typology when naturalness *is* an issue. Indeed, it would be unparsimonious to expect otherwise.

On the other hand, that does not mean that analytic bias can be read directly off of the typological facts, as is tacitly or explicitly assumed in most UG-based approaches (McCarthy 1988, 2002: 108–120, Prince & Smolensky 1993: 5), since Experiment 2 found an analytic bias which does not correspond to a typological asymmetry. Nor can analytic bias be inferred from 'phonetic naturalness' in the sense of precursor robustness, since the HH and HV patterns had equally robust precursors but differed in analytic bias. That result is particularly interesting in connection with the hypothesis that Universal Grammar is 'phonetically grounded' (see e.g. the papers in Hayes *et al*. 2004). If the hypothesis is correct, the present results imply one of two things. Either some simple precursors do

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

not give rise to corresponding constraints, or else learning mechanisms have more difficulty finding some rankings than others. In either case, an explanation is needed.

Researchers working in frameworks which rely principally on analytic or channel bias have expressed doubt about the single-factor focus, but have hesitated to abandon it (Pater 2004: 284–285, Przezdziecki 2005: 26–27). I believe that the main reason for this reluctance is a concern that admitting both factors will only make a hard problem even harder. Because so much of typology can be fit by a well-tailored theory using either factor alone, it seems hopeless to use typological data to decide among the enormous number of possible two-factor theories. What the present findings mean for linguistic theory is that, first, neither channel nor analytic bias can be safely neglected in explaining typology (Hyman 2001, Myers 2002, Kiparsky 2008; see also Cole & Iskarous 2001, Blevins 2006: 246, Wilson 2006), but that, second, it is possible to acknowledge both factors and still arrive at a firm conclusion in particular cases – as well as generating new questions and testable hypotheses.

We have to ask what the contributions of analytic and channel bias are and how they can be distinguished empirically. In particular, we seek restrictive hypotheses about how the two factors interact that offer some hope of controlling the explosion of possible explanations for any given typological fact. These are research problems for the long term, but here are some concrete initial suggestions.

One very restrictive hypothesis is that analytic bias is decisive only when precursor robustness is not. In all of the cases discussed in this paper, the phonological patterns differ in frequency, while the precursors match (or nearly match) in robustness: height-height and height-voice, tone-tone and voice-tone, and voice-tone and height-tone. It is logically possible that the patterns might match in frequency while the precursors differ in robustness, or that the typological difference is opposite to the precursor difference. It is an open question whether selective learning can offset or reverse the effects of differential precursor robustness, and, if so, in what circumstances. Diachronic conspiracies, as well as mismatches between perception and sound change, suggest likely places to look.

A second restrictive hypothesis is that Universal Grammar determines which patterns are attestable and which are unattestable ('hard' typology), whereas precursor biases determine which of the attestable patterns are actually attested ('soft' typology). The proposal has been made in a number of places (Hale & Reiss 2000, Hyman 2001, Myers 2002, Buckley 2003, Kiparsky 2006). The evidence that any (reasonably simple) pattern is genuinely unlearnable is very slim, the only case of which I know being the non-adjacent syllable dependency studied extensively by Newport & Aslin (2004). However, the question has been little studied, and future research may turn up more cases. Underphonologisation and diachronic conspiracies will be informative in deciding where to look for them.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Appendix: Participant instructions

The following instructions were presented at the beginning of the experiment:

> Welcome!
>
> This experiment is about learning to recognize words in two artificial languages. You will first study words of the language; then you will be tested on how well you can recognize them.
>
> The study phase goes like this. The computer pronounces a word for you. You pronounce it back, trying to match the pronunciation as closely as possible. Then you click on a button that says 'Next' to get the next word. This will go on for a long time. The computer will record your speech, but it will not tell you how accurate your pronunciation was.
>
> After the study phase, there will be a test. The computer will say two words. One is a word of the language you studied; the other is not. You should choose the one that you think was in the language you studied – click '1' if it was the first word, '2' if it was the second word. If you can't tell, make your best guess. The computer will then play you the next pair of words, until you have finished the test.
>
> Message boxes like this one will appear when needed to remind you what's coming next. When you're ready, click 'Continue' to begin the experiment.

At the beginning of the break, the following text appeared, and remained on the screen until participant proceeded to the second Language condition.

> You have reached the end of the test phase for Language A. It's time for a break! Some music will start shortly. When the music ends, please click 'Continue' to go on to Language B.

At the start of each familiarisation phase, participants received the following reminder:

> You're about to begin the study phase. The computer will pronounce a word for you. You should repeat it aloud, trying to match the pronunciation as closely as you can. Then click the 'Next' button to go on.
>
> If you make a mistake, don't worry about it; just go right on to the next word.
>
> When you're ready to begin the study phase, click 'Continue'.

At the start of each test phase, they received the following reminder:

> You have reached the end of the study phase.
>
> Now comes the test phase. The computer will say two words. One is a word of the language you studied; the other is not. You should choose the one that you think was in the language you studied – click '1' if it was the first word, '2' if it was the second word. If you can't tell for sure, you should make your best guess. After you have answered, the computer will play you another pair of words, and so on until you have finished the test.
>
> If you make a mistake, don't worry about it; just keep right on going.
>
> Please click 'Continue' when you are ready to start the test phase.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

REFERENCES

Altmann, Gabriel & Werner Lehfeldt (1980). *Einführung in die quantitative Phonologie*. Bochum: Studienverlag Brockmeyer.
Anderson, Stephen R. (1981). Why phonology isn't 'natural'. *LI* **12**. 493–539.
Archangeli, Diana & Douglas Pulleyblank (1994). *Grounded phonology*. Cambridge, Mass.: MIT Press.
Arensen, Jon (1982). *Murle grammar*. Juba, Sudan: University of Juba & Summer Institute of Linguistics.
Asher, R. E. (1985). *Tamil*. London: Routledge.
Baayen, R. H., R. Piepenbrock & L. Gulikers (1995a). *The CELEX lexical database*. Release 2. [CD-ROM]. Philadelphia: Linguistic Data Consortium, University of Pennsylvania.
Baayen, R. H., R. Piepenbrock & L. Gulikers (1995b). *English linguistic guide*. In Baayen *et al*. (1995a). File EUG_LET.PS.
Bach, Emmon & Robert T. Harms (1972). How do languages get crazy rules? In Robert P. Stockwell & Ronald K. S. Macaulay (eds.) *Linguistic change and generative theory*. Bloomington: Indiana University Press. 1–21.
Baković, Eric (2000). *Harmony, dominance, and control*. PhD dissertation, Rutgers University.
Barnes, Jonathan (2002). *Positional neutralization: a phonologization approach to typological patterns*. PhD dissertation, University of California, Berkeley.
Beardsley, Amy Neel & Walter L. Cullinan (1987). Speech sample type and children's segmental durations. *JPh* **15**. 29–38.
Beddor, Patrice Speeter & Rena Arens Krakow (1999). Perception of coarticulatory nasalization by speakers of English and Thai: evidence for partial compensation. *JASA* **106**. 2868–2887.
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.
Beddor, Patrice Speeter, Rena Arens Krakow & Stephanie Lindemann (2001). Patterns of perceptual compensation and their phonological consequences. In Hume & Johnson (2001b). 55–78.
Bender, M. L. (1974). Phoneme frequencies in Amharic. *Journal of Ethiopian Studies* **12**. 19–24.
Berent, Iris, Donca Steriade, Tracy Lennertz & Vered Vaknin (2007). What we know about what we have never heard: evidence from perceptual illusions. *Cognition* **104**. 591–630.
Bermúdez-Otero, Ricardo (2006). Phonological change in Optimality Theory. In Keith Brown (ed.) *Encyclopedia of language and linguistics*. 2nd edn. Vol. 9. Oxford: Elsevier. 497–505.
Blevins, Juliette (2004). *Evolutionary Phonology: the emergence of sound patterns*. Cambridge: Cambridge University Press.
Blevins, Juliette (2006). Reply to commentaries. *Theoretical Linguistics* **32**. 245–256.
Blumenfeld, Lev (2006). *Constraints on phonological interactions*. PhD dissertation, Stanford University.
Boersma, Paul (1998). *Functional phonology: formalizing the interaction between articulatory and perceptual drives*. PhD dissertation, University of Amsterdam.
Boersma, Paul (2004). Review of Bruce Tesar & Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, Mass.: MIT Press. Downloaded April 2004 from http://www.fon.hum.uva.nl/paul/papers/TesarSmolenskyReview.pdf.
Boersma, Paul & Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *LI* **32**. 45–86.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Boersma, Paul & David Weenink (2005). *Praat : a system for doing phonetics by computer* (version 4.3.02). Retrieved March 2005 from http://www.praat.org/.

Bogoras, Waldemar (1922). Chukchee. In Franz Boas (ed.) *Handbook of American Indian languages*. Part 2. Washington, D.C.: Government Printing Office. 631–903.

Bole-Richard, Rémy (1983). *Systématique phonologique et grammaticale d'un parler ewe : le gen-mina du sud-togo et sud-bénin*. Paris: L'Harmattan.

Bonatti, Luca L., Marcela Peña, Marina Nespor & Jacques Mehler (2005). Linguistic constraints on statistical computations: the role of consonants and vowels in continuous speech processing. *Psychological Science* **16**. 451–459.

Buckley, Eugene (2003). Children's unnatural phonology. *BLS* **29**. 523–534.

Chambers, J. K. (1973). Canadian raising. *Canadian Journal of Linguistics* **18**. 113–135.

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

Chambers, Kyle E., Kristine H. Onishi & Cynthia Fisher (2006). A vowel is a vowel: generalizing newly-learned phonotactic constraints to novel contexts. Ms, University of Rochester.

Chang, Steve S., Madelaine C. Plauché & John J. Ohala (2001). Markedness and consonant confusion asymmetries. In Hume & Johnson (2001b). 79–101.

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

Clements, G. N. & Elizabeth V. Hume (1995). The internal organization of speech sounds. In Goldsmith (1995). 245–306.

Cole, Jennifer & Khalil Iskarous (2001). Effects of vowel context on consonant place identification: implications for a theory of phonologization. In Hume & Johnson (2001b). 103–122.

Creel, Sarah C., Elissa L. Newport & Richard N. Aslin (2004). Distant melodies: statistical learning of nonadjacent dependencies in tone sequences. *Journal of Experimental Psychology : Learning, Memory, and Cognition* **30**. 1119–1130.

Crowther, Court S. & Virginia A. Mann (1992). Native language factors affecting use of vocalic cues to final consonant voicing in English. *JASA* **92**. 711–722.

Cutler, Anne, Andrea Weber, Roel Smits & Nicole Cooper (2004). Patterns of English phoneme confusions by native and non-native listeners. *JASA* **116**. 3668–3678.

Darcy, Isabelle, Franck Ramus, Anne Cristophe, Katherine Kinzler & Emmanuel Dupoux (in press). Phonological knowledge in compensation for native and non-native assimilation. In Frank Kügler, Caroline Féry & Ruben van de Vijver (eds.) *Variation and gradience in phonetics and phonology*. Berlin: Mouton de Gruyter.

Davidson, Lisa, Peter Jusczyk & Paul Smolensky (2004). The initial and final states: theoretical implications and experimental explorations of Richness of the Base. In René Kager, Joe Pater & Wim Zonneveld (eds.) *Constraints in phonological acquisition*. Cambridge: Cambridge University Press. 321–368.

Dawson, Willa (1980). *Tibetan phonology*. PhD dissertation, University of Washington.

de Jong, Kenneth & Bushra Zawaydeh (2002). Comparing stress, lexical focus, and segmental focus: patterns of variation in Arabic vowel duration. *JPh* **30**. 53–75.

Dell, Gary S., Kristopher D. Reed, David R. Adams & Antje S. Meyer (2000). Speech errors, phonotactic constraints, and implicit learning: a study of the role of experience in language production. *Journal of Experimental Psychology : Learning, Memory, and Cognition* **26**. 1355–1367.

Dettweiler, Stephen H. (2000). Vowel harmony and neutral vowels in c'Lela. *Journal of West African Languages* **28:1**. 3–18.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Dutoit, T., V. Pagel, N. Pierret, F. Bataille & O. van der Vreken (1996). The MBROLA project: towards a set of high quality speech synthesizers free of use for non commercial purposes. *Proceedings of the International Conference on Spoken Language Processing* (*ICSLP*). Vol. 3. 1393–1396. Available (March 2008) at http://www.asel.udel.edu/icslp/cdrom/authindx.htm.

Ebert, Karen H. (1976). *Sprache und Tradition der Kera* (*Tschad*). Vol. 2: *Lexikon/Lexique*. Berlin: Reimer.

Ebert, Karen H. (1979). *Sprache und Tradition der Kera* (*Tschad*). Vol. 3: *Grammatik*. Berlin: Reimer.

Egerod, Søren (1956). *The Lungtu dialect : a descriptive and historical study of a South Chinese idiom*. Copenhagen: Munksgaard.

Fischer-Jørgensen, Eli (1972). *PTK* et *BDG* français en position intervocalique accentuée. In Albert Valdman (ed.) *Papers in linguistics and phonetics to the memory of Pierre Delattre*. The Hague: Mouton. 143–200.

Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe (2004). Similarity avoidance and the OCP. *NLLT* **22**. 179–228.

Gakuru, Mucemi, Frederick K. Iraki, Roger Tucker, Ksenia Shalanova & Kamenda Ngugi (no date). Design of speech data base for unit selection in Kiswahili TTS. Available (March 2008) from www.llsti.org/pubs/Kiswahili_Speech_Database.pdf.

Goldrick, Matthew (2004). Phonological features and phonotactic constraints in speech production. *Journal of Memory and Language* **51**. 586–603.

Goldsmith, John (ed.) (1995). *The handbook of phonological theory*. Cambridge, Mass. & Oxford: Blackwell.

Gordon, Matthew (2004). Syllable weight. In Hayes *et al*. (2004). 277–312.

Gordon, Raymond G., Jr. (ed.) (2005). *Ethnologue : languages of the world*. 15th edn. Dallas: SIL International. http://www.ethnologue.com.

Haas, Mary (1946). A grammatical sketch of Tunica. In Cornelius Osgood (ed.) *Linguistic structures of Native America*. New York: Viking Fund. 337–366.

Hale, Mark & Charles Reiss (2000). 'Substance abuse' and 'dysfunctionalism': current trends in phonology. *LI* **31**. 157–169.

Hallé, Pierre A., Juan Segui, Uli Frauenfelder & Christine Meunier (1998). Processing of illegal consonant clusters: a case of perceptual assimilation? *Journal of Experimental Psychology : Human Perception and Performance* **24**. 592–608.

Hansson, Gunnar Ólafur (2001). *Theoretical and typological issues in consonant harmony*. PhD dissertation, University of California, Berkeley.

Hansson, Gunnar Ólafur (2004). Long-distance voicing agreement: an evolutionary perspective. *BLS* **30**. 130–141.

Haspelmath, Martin (1999). Optimality and diachronic adaptation. *Zeitschrift für Sprachwissenschaft* **18**. 180–205.

Hayes, Bruce, Robert Kirchner & Donca Steriade (eds.) (2004). *Phonologically based phonology*. Cambridge: Cambridge University Press.

Hayes, Bruce & Donca Steriade (2004). Introduction: the phonetic bases of phonological markedness. In Hayes *et al*. (2004). 1–33.

Hetzron, Robert (1969). *The verbal system of Southern Agaw*. Berkeley & Los Angeles: University of California Press.

Holt, Lori L. (2005). Temporally nonadjacent nonlinguistic sounds affect speech categorization. *Psychological Science* **16**. 305–312.

Hombert, Jean-Marie, John J. Ohala & William G. Ewan (1979). Phonetic explanations for the development of tones. *Lg* **55**. 37–58.

Hualde, José Ignacio (1989). Autosegmental and metrical spreading in the vowel-harmony systems of northwestern Spain. *Linguistics* **27**. 773–805.

Hualde, José Ignacio (1991). *Basque phonology*. London & New York: Routledge.

Hume, Elizabeth & Keith Johnson (2001a). A model of the interplay of speech perception and phonology. In Hume & Johnson (2001b). 3–26.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

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: Anma Libri. 407–418.
Hyman, Larry M. (1998). Positional prominence and the 'prosodic trough' in Yaka. *Phonology* **15**. 41–75.
Hyman, Larry M. (2001). The limits of phonetic determinism in phonology: *NC revisited. In Hume & Johnson (2001b). 141–185.
Jacobsen, William H., Jr. (1964). *A grammar of the Washo language*. PhD dissertation, University of California, Berkeley.
Kämpfe, Hans-Rainer & Alexander P. Volodin (1995). *Abriß der tschuktschischen Grammatik auf der Basis der Schriftsprache*. Wiesbaden: Harrassowitz.
Kavitskaya, Darya (2002). *Compensatory lengthening : phonetics, phonology, diachrony*. London & New York: Routledge.
Kawahara, Shigeto (2005). Voicing and geminacy in Japanese: an acoustic and perceptual study. *University of Massachusetts Occasional Papers in Linguistics* **31**. 87–120.
Keane, Elinor (2001). *Echo words in Tamil*. PhD dissertation, Oxford University.
Kiparsky, Paul (1995). The phonological basis of sound change. In Goldsmith (1995). 640–670.
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.
Koenig, Laura L. & Areti Okalidou (2003). Stress effects on coarticulation in English and Greek. Poster presented at the 145th Meeting of the Acoustical Society of America, Nashville.
Koo, Hahn & Jennifer Cole (2006). On learnability and naturalness as constraints on phonological grammar. In Antonis Botinis (ed.) *Proceedings of ISCA Tutorial and Research Workshop on Experimental Linguistics*. Athens: University of Athens. 174–177.
Lampp, Claire & Heidi Reklis (2004). Effects of coda voicing and aspiration on Hindi vowels. Poster presented at the 146th Meeting of the Acoustical Society of America, New York.
Lotto, Andrew J. & Lori L. Holt (2006). Putting phonetic context effects into context: a commentary on Fowler (2006). *Perception and Psychophysics* **68**. 178–183.
Lubker, James & Thomas Gay (1982). Anticipatory labial coarticulation: experimental, biological, and linguistic variables. *JASA* **71**. 437–448.
McCarthy, John J. (1983). Phonological features and morphological structure. In John F. Richardson, Mitchell Marks & Amy Chukerman (eds.) *Papers from the parasession on the interplay of phonology, morphology, and syntax*. Chicago: Chicago Linguistic Society. 135–161.
McCarthy, John J. (1988). Feature geometry and dependency: a review. *Phonetica* **45**. 84–108.
McCarthy, John J. (2002). *A thematic guide to Optimality Theory*. Cambridge: Cambridge University Press.
Macmillan, Neil A. & C. Douglas Creelman (1990). *Detection theory : a user's guide*. Cambridge: Cambridge University Press.
Manuel, Sharon Y. (1990). The role of contrast in limiting vowel-to-vowel coarticulation in different languages. *JASA* **88**. 1286–1298.
Mintz, Toby & Rachel Walker (2006). Infants' sensitivity to vowel harmony and its role in word segmentation. Paper presented at the 80th Annual Meeting of the Linguistic Society of America, Albuquerque.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Moreton, Elliott (2004). Realization of the English postvocalic [voice] contrast in F_1 and F_2. *JPh* **32**. 1–33.
Moreton, Elliott (in press). Underphonologization and modularity bias. In Steven Parker (ed.) *Phonological argumentation*. London: Equinox.
Moreton, Elliott & Shigeaki Amano (1999). Phonotactics in the perception of Japanese vowel length: evidence for long-distance dependencies. In *Proceedings of the 6th European Conference on Speech Communication and Technology*, Budapest, 2679–2682.
Moreton, Elliott, Gary Feng & Jennifer L. Smith (in press). Syllabification, sonority, and perception: new evidence from a language game. *CLS* **41**.
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. 37–64.
Myers, Scott (2002). Gaps in factorial typology: the case of voicing in consonant clusters. Ms, University of Texas, Austin.
Newport, Elissa L. & Richard N. Aslin (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology* **48**. 127–162.
Nikolaeva, Irina & Marina Tolskaya (2001). *A grammar of Udihe*. Berlin & New York: Mouton de Gruyter.
Ohala, John J. (1990). The phonetics and phonology of aspects of assimilation. In John Kingston & Mary E. Beckman (eds.) *Papers in laboratory phonology I : between the grammar and physics of speech*. Cambridge: Cambridge University Press. 258–275.
Ohala, John J. (1993). The phonetics of sound change. In Charles Jones (ed.) *Historical linguistics : problems and perspectives*. London: Longman. 237–278.
Ohala, John J. (1994a). Hierarchies of environments for sound variation; plus implications for 'neutral' vowels in vowel harmony. *Acta Linguistica Hafniensia* **27**. 371–382.
Ohala, John J. (1994b). Towards a universal, phonetically-based theory of vowel harmony. *Proceedings of the 3rd International Conference on Spoken Language Processing* (*ICSLP 94*). Vol. 2. Yokohama: Acoustical Society of Japan. 491–494.
Ohala, John J. (2005). Phonetic explanations for sound patterns: implications for grammars of competence. In William J. Hardcastle & Janet Mackenzie Beck (eds.) *A figure of speech : a Festschrift for John Laver*. Mahwah, NJ: Erlbaum. 23–38.
Ohala, John J. & Deborah Feder (1994). Listeners' normalization of vowel quality is influenced by 'restored' consonantal context. *Phonetica* **51**. 111–118.
Onishi, Kristine H., Kyle E. Chambers & Cynthia Fisher (2002). Learning phonotactic constraints from brief auditory experience. *Cognition* **83**. B13–B23.
Palmer, F. R. (1959). The verb classes of Agau (Awiya). *Mitteilungen des Instituts für Orientforschung* **7**. 270–297.
Parkinson, Frederick B. (1996). *The representation of vowel height in phonology*. PhD dissertation, Ohio State University.
Paster, Mary (2004). Vowel height harmony and blocking in Buchan Scots. *Phonology* **21**. 359–407.
Pater, Joe (2004). Austronesian nasal substitution and other NÇ effects. In John J. McCarthy (ed.) *Optimality Theory in phonology : a reader*. Malden, Mass.: Blackwell. 271–289.
Pearce, Mary (2003). *Vowel harmony in Kera* (*Chadic*). MA dissertation, University College London.
Pearce, Mary (2005). Kera tone and voicing. *UCL Working Papers in Linguistics* **17**. 61–82.
Pertz, D. L. & T. G. Bever (1975). Sensitivity to phonological universals in children and adults. *Lg* **51**. 149–162.
Pitkin, Harvey (1984). *Wintu grammar*. Berkeley: University of California Press.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Port, Robert F. (1981). Linguistic timing factors in combination. *JASA* **69**. 262–274.
Port, Robert F. & Rosemarie Rotunno (1979). Relation between voice-onset time and vowel duration. *JASA* **66**. 654–662.
Prince, Alan & Paul Smolensky (1993). *Optimality Theory : constraint interaction in generative grammar*. Ms, Rutgers University & University of Colorado, Boulder, Published 2004, Malden, Mass. & Oxford: Blackwell.
Przezdziecki, Marek (2005). *Vowel harmony and coarticulation in three dialects of Yoruba : phonetics determining phonology*. PhD dissertation, Cornell University.
Pycha, Anne, Pawel Nowak, Eurie Shin & Ryan Shosted (2003). Phonological rule-learning and its implications for a theory of vowel harmony. *WCCFL* **22**. 423–435.
R Development Core Team (2005). R: a language and environment for statistical computing. Vienna: R Foundation for Statistical Computing. http://www.r-project.org.
Roodenrys, Steven & Melinda Hinton (2002). Sublexical or lexical effects on serial recall of nonwords? *Journal of Experimental Psychology : Learning, Memory, and Cognition* **28**. 29–33.
Rose, Sharon & Rachel Walker (2004). A typology of consonant agreement as correspondence. *Lg* **80**. 475–531.
Rousset, Isabelle (2004). *Structures syllabiques et lexicales des langues du monde*. PhD dissertation, Université Grenoble III.
Saffran, Jenny R. (2002). Constraints on statistical language learning. *Journal of Memory and Language* **47**. 172–196.
Saffran, Jenny R. (2003). Statistical language learning: mechanisms and constraints. *Current Directions in Psychological Science* **12**. 110–114.
Saffran, Jenny R. & Erik D. Thiessen (2003). Pattern induction by infant language learners. *Developmental Psychology* **39**. 484–494.
Sanders, Nathan (2003). *Opacity and sound change in the Polish lexicon*. PhD dissertation, University of California, Santa Cruz.
Schane, Sanford A., Bernard Tranel & Harlan Lane (1974). On the psychological reality of a natural rule of syllable structure. *Cognition* **3**. 351–358.
Schiffman, Harold F. (1999). *A reference grammar of spoken Tamil*. Cambridge: Cambridge University Press.
Seiden, William (1960). Chamorro phonemes. *Anthropological Linguistics* **2:4**. 6–35.
Seidl, Amanda & Eugene Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development* **1**. 289–316.
Sigurd, B. (1968). Rank-frequency distribution for phonemes. *Phonetica* **18**. 1–15.
Sohn, Ho-Min (1971). A-raising in Woleaian. *University of Hawaii Working Papers in Linguistics* **3:8**. 15–35.
Sohn, Ho-Min (1975). *Woleaian reference grammar*. Honolulu: University Press of Hawaii.
Sohn, Hyang-Sook (1986). Toward an integrated theory of morphophonology: vowel harmony in Korean. *Studies in the Linguistic Sciences* **16:2**. 157–184.
Steriade, Donca (2001a). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. Ms, University of California, Los Angeles.
Steriade, Donca (2001b). Directional asymmetries in place assimilation: a perceptual account. In Hume & Johnson (2001b). 219–250.
Storkel, Holly L., Jonna Armbrüster & Tiffany P. Hogan (2006). Differentiating phonotactic probability and neighborhood density in adult word learning. *Journal of Speech, Language, and Hearing Research* **49**. 1175–1192.
Suárez, Jorge A. (1983). *La lengua tlapaneca de Malinaltepec*. Mexico City: Universidad Nacional Autónoma de México.
Summers, W. van (1987). Effects of stress and final-consonant voicing on vowel production: articulatory and acoustic analysis. *JASA* **82**. 847–863.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press

Tamaoka, Katuo & Shogo Makioka (2004). Frequency of occurrence for units of phonemes, morae, and syllables appearing in a lexical corpus of a Japanese newspaper. *Behavioral Research Methods, Instruments, and Computers* **36**. 531–547.

Taylor, Conrad F. & George Houghton (2005). Learning artificial phonotactic constraints: time course, durability, and relationship to natural constraints. *Journal of Experimental Psychology : Learning, Memory, and Cognition* **31**. 1398–1416.

Thomas, Erik R. (2000). Spectral differences in /ai/ offsets conditioned by voicing of the following consonant. *JPh* **28**. 1–25.

Vagges, K., F. E. Ferrero, E. Magno-Caldognetto & C. Lavagnoli (1978). Some acoustic characteristics of Italian consonants. *Journal of Italian Linguistics* **3:1**. 69–85.

Vainio, Martti (1996). Phoneme frequencies in Finnish text and speech. In A. Iivonen & A. Klippi (eds.) *Studies in logopedics and phonetics 5*. Publications of the Department of Phonetics, University of Helsinki, Series B: Phonetics, Logopedics and Speech Communication. 181–192.

Walker, Rachel (2005). Weak triggers in vowel harmony. *NLLT* **23**. 917–989.

Weismer, Gary (1979). Sensitivity of voice-onset-time (VOT) measures to certain segmental features in speech production. *JPh* **7**. 197–204.

Wilson, Colin (2003a). Experimental investigation of phonological naturalness. *WCCFL* **22**. 533–546.

Wilson, Colin (2003b). Analytic bias in artificial phonology learning: consonant harmony *vs*. random alternation. Handout from presentation at the Workshop on Markedness and the Lexicon, MIT.

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

Wiswall, Wendy J. (1991). *Partial vowel harmonies as evidence for a Height node*. PhD dissertation, University of Arizona.

Wolf, Catherine G. (1978). Voicing cues in English final stops. *JPh* **6**. 299–309.

https://doi.org/10.1017/S0952675708001413 Published online by Cambridge University Press