



Glewwe, Eleanor. 2022. Substantive bias and the positional extension of major place contrasts. *Glossa: a journal of general linguistics* 7(1). pp. 1–37. DOI: <https://doi.org/10.16995/glossa.6537>



Open Library of Humanities

## Substantive bias and the positional extension of major place contrasts

Eleanor Glewwe, Grinnell College, US, [glewweel@grinnell.edu](mailto:glewweel@grinnell.edu)

---

The role of phonetic naturalness in biasing the learning of sound patterns remains an unsettled question in phonology. The present study investigates naturalness bias in phonotactic learning using a novel experimental paradigm that tests whether learners reproduce a phonetically-motivated phonotactic implicational about the distribution of major place contrasts in stops. Stops differing in place of articulation are easier to distinguish word-initially than word-finally, so place contrasts in word-final position should also exist in word-initial position. The reverse is not necessarily true. This implicational is typologically supported as well as motivated by perceptual naturalness. In two artificial grammar learning experiments, I exposed participants to place contrasts in stops either word-initially or word-finally and tested whether they extended the contrasts to the other word-edge position. Participants successfully learned to recognize novel words that fit the phonotactic pattern they had been trained on, but they were equally willing to extend the place contrasts in both directions, yielding no evidence for naturalness bias. These results contrast with those of a similar study that found asymmetric extension of the stop voicing contrast, supporting an effect of naturalness bias. Confusion data suggests that the reduction in perceptibility from word-initial to word-final position may be greater for stop voicing than for stop place of articulation. This difference may underlie the divergent results of the two studies, leading to the hypothesis that the strength of a substantive bias depends on the magnitude of its phonetic precursor.

---

*Glossa: a journal of general linguistics* is a peer-reviewed open access journal published by the Open Library of Humanities. © 2022 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

**OPEN ACCESS**



## 1. Introduction

Crosslinguistically, phonological patterns tend to be phonetically natural, meaning they are motivated by articulatory and perceptual phonetic tendencies (Stampe 1979; Archangeli & Pulleyblank 1994; Beguš 2019). There are two main proposals for why this is the case. The channel bias, or historical, explanation maintains that articulatory difficulty and misperception in the course of language transmission shape language change, leading over time to phonetically-motivated phonological patterns and processes (Ohala 1992; Blevins 2004). The substantive bias explanation claims that learners are influenced by a synchronic cognitive bias against phonetically unnatural phonological patterns which shapes language acquisition in such a way that phonetically natural patterns end up dominating in the phonological typology (Wilson 2006; Moreton & Pater 2012b; White 2017). While it has been argued that channel bias suffices to explain why phonology is largely phonetically natural, rendering it unnecessary and undesirable to posit a further substantive bias (Ohala 1974; Hale & Reiss 2000; Blevins 2004), it may be that both forces exist and influence the types of phonological patterns seen crosslinguistically (Hayes & Steriade 2004; Moreton 2008). Channel bias and substantive bias are difficult to distinguish, however, because both generally make the same predictions about which phonological patterns should be typologically frequent (Moreton 2008).

In arguing in favor of its existence, proponents of substantive bias have primarily relied on experimentation, particularly artificial grammar learning (AGL). If participants (infants, children, or adults) learn phonetically natural patterns better than phonetically unnatural patterns in AGL experiments, this provides evidence for a cognitive bias against acquiring phonetically unnatural patterns that may play a role in addition to the diachronic channel bias. The picture is currently inconclusive, however, because while some AGL experiments have uncovered learning advantages for phonetically motivated patterns (e.g. Wilson 2006; Finley 2012; White 2013), many others have yielded null results in which learners acquired natural and unnatural patterns equally well (e.g. Skoruppa & Peperkamp 2011; Do & Zsiga & Havenhill 2016; Lysvik 2018). Moreton & Pater's (2012b) review concluded that there was little experimental evidence for substantive bias, suggesting it is weak or absent. Since then, AGL experimentation has continued to turn up both positive and null results.

There are a number of potential explanations for why substantive bias effects are so elusive in AGL studies. If substantive bias is relatively subtle (and/or difficult for current AGL methodologies to detect), then it may only be uncovered in very sensitive experiments. In this case, we expect fairly frequent null results, and it is the steady accumulation of solid positive results that would prove that substantive bias does influence learners. On the other hand, if there is no synchronic bias against unnatural phonological patterns, then the positive results in the AGL literature must be spurious. Publication bias may have inflated their relative frequency, with studies

finding effects in the opposite direction of substantive bias potentially going unreported. Null results may also be underreported relative to their actual incidence. Between these two poles (substantive bias exists vs. there is no substantive bias), there could be more nuanced proposals. It might be that not all phonetic precursors to natural phonological patterns (i.e. perceptual or articulatory facts that make the patterns phonetically natural, perhaps in comparison to other logically possible patterns) give rise to substantive biases. Glewwe et al. (2018) contend that perceptual naturalness, but not articulatory naturalness, biases phonological learning. I consider a different possibility, namely, that the magnitude of a phonetic precursor affects the strength of the substantive bias it supports, or even whether that bias exists at all. If substantive biases vary in strength and only the strongest are detectible in AGL experiments, this could explain the patchwork of positive and null results.

This paper presents an AGL study testing for substantive bias in phonotactic learning. In two experiments, I exposed learners to major place contrasts in stops in either word-initial or word-final position and tested whether they extended the contrasts to the other word-edge position. Naturalness bias predicts that learners should extend such contrasts more from word-final position to word-initial position than vice versa, but in fact participants in my experiments extended equally in both directions. While this study cannot definitively distinguish between the abovementioned theories about why it is difficult to obtain substantive bias effects in AGL experiments, it can make a contribution. First, it represents another case of a null substantive bias result, helping to mitigate a potential publication bias. It is also useful to know *which* naturalness effects supported by phonetic precursors fail to emerge because the specific pattern of positive and null results may provide clues to the nature of substantive bias. This leads to the study's second contribution: it points to a potentially fruitful direction for the substantive bias research program. I compare the results of the present study to those of a parallel study that examined the stop voicing contrast rather than stop place contrasts and which did find an effect of naturalness. Because the experiments are so similar, many factors that might explain why a positive result was found in one case and a null result in another are controlled for. I present preliminary evidence that a difference in phonetic precursor size may underlie the divergent results of the two studies. More broadly, a new approach to untangling the contradictory picture in the AGL literature might involve devising ways of quantifying and comparing the magnitude of phonetic precursors in order to generate testable predictions about which substantive biases should be more or less likely to emerge in AGL experiments.

In the remainder of the paper, I introduce the phonotactic implicational that underpins the experiments (Section 2) and present the AGL study itself (Section 3). I then contrast its results with those of an analogous study (Glewwe 2021) exploring the stop voicing contrast and discuss how the studies in combination suggest a path forward in substantive bias research (Section 4). Section 5 concludes.

## 2. A phonotactic implicational: major place contrasts in stops

Past work on naturalness bias effects in phonological learning has looked more often at alternations than at phonotactics. When AGL experiments test for naturalness bias in phonotactic learning, they typically expose participants to stimuli that provide negative evidence for a particular phonotactic constraint (e.g. \* $[\alpha \text{ round}]$ [- $\alpha \text{ round}]$  in Skoruppa & Peperkamp 2011, \*[-son, + voice]# or \*[-son, -voice]# in Greenwood 2016) and then test how well participants appear to have learned this constraint. My approach differs in specifically testing the positional extension of phonemic contrasts. Rather than simply gauging how well learners master a particular phonotactic, I use a new paradigm in which I expose participants to certain phonemic contrasts in one position and then check to what extent they believe those same contrasts exist in another position. While the phonotactic patterns in the training data can also be expressed by a single phonotactic constraint, the additional element of comparison across contexts is also present, setting this design apart.<sup>1</sup>

The positional extension I tested is rooted in a phonotactic implicational concerning the positional distribution of major place contrasts. If a language has major place contrasts in stops post-vocally/in coda position (e.g. /ap/ vs. /at/ vs. /ak/), it should also have major place contrasts in stops pre-vocally/in onset position (e.g. /pa/ vs. /ta/ vs. /ka/), but not necessarily vice versa (Blevins 2004; O'Hara n.d.). This implicational is phonetically motivated. Based on Fujimura & Macchi & Streeter's (1978) finding that listeners attend more to CV transitions than to VC transitions in identifying the place of articulation of voiced stops, Steriade (1994; 2001a) proposed that differences in place of articulation are more perceptible pre-vocally ( $\_V$ ) than post-vocally ( $V\_$ ).<sup>2</sup> Blevins (2004) also synthesizes research demonstrating that the strongest cues to place of articulation in stops are in the CV transition and that place cues are weaker when the stop is not followed by a vowel. In other words, consonants differing in place of articulation are more perceptually similar (i.e. harder to distinguish) after a vowel than before a vowel. If a place contrast exists post-vocally, where it is harder to perceive, then it should also exist pre-vocally, where it is easier to perceive. It would be phonetically unnatural for a language to exhibit a contrast in a less salient position, perceptually speaking, while not exhibiting it in a more salient one.

---

<sup>1</sup> An earlier AGL study that also invited participants to extend a phonotactic pattern to a new context is Myers & Padgett 2014. In that case, participants were trained on a phrase-final restriction on obstruent voicing and tested on whether the restriction applied word-finally as well. The comparison was between phonological domains rather than positions in the word.

<sup>2</sup> More precisely, this applies to major place contrasts (labial vs. coronal vs. dorsal) in oral stops. The following discussion refers only to these contrasts. Other contrasts, such as that between retroflex and non-retroflex consonants, have different cues and therefore different patterns of perceptual similarity across positions (see Blevins 2004; Steriade 2001a).

The phonotactic implicational about the distribution of major place contrasts appears to hold crosslinguistically. Steriade (1994) observes that place neutralization typically occurs in coda position; similarly, Lombardi (2001) notes that many languages exhibit neutralization to an unmarked place in coda position. Further typological evidence for the preference for maintaining major place contrasts pre-vocally comes from the fact that coda consonants are more likely to be targets of place assimilation than onset consonants (Jun 1995). Referring specifically to word edges, Blevins (2004: 117) concludes, based on observations from dozens of genetically diverse languages, that “word-final neutralization of major place features is common, while word-initial neutralization of the same features is unattested.” This statement is perfectly consistent with the phonotactic implicational. O’Hara’s (n.d.) Word-Edge Consonant Database contains information about the word-initial and word-final consonant inventories of 174 languages from a wide variety of language families. Of the 57 languages in his database that exhibit exactly [p t k] word-finally, all also allow [p t k] (and sometimes other stops besides) word-initially. On the other hand, in a larger survey of 96 genetically-balanced languages that exhibit exactly [p t k] word-initially, O’Hara (2021) found that over half allow fewer stops in word-final position (with by far the most common pattern being a ban on all three stops word-finally).<sup>3</sup> These typological surveys also support the implicational whereby major place contrasts in word-final position entail those same contrasts word-initially but not vice versa.

If phonetic knowledge influences learners’ acquisition of phonological patterns and learners are biased toward phonetically natural phonological systems, they should make inferences about language in accordance with phonotactic implicational like this one. Consequently, I used the phonotactic implicational about the distribution of stops with different places of articulation to structure an AGL study testing for substantive bias.<sup>4</sup>

### 3. The positional extension of major place contrasts in stops

#### 3.1 Experiment 1

In Experiment 1, I exposed participants to artificial languages featuring a three-way place contrast in either word-initial (pre-vocalic) or word-final (post-vocalic) position and tested whether they assumed the contrast existed in the other position as well. If participants behave in a way that

---

<sup>3</sup> In conducting his typological study, O’Hara (2021) looked at the stop series with the greatest number of licit places of articulation in a given language. This was most often the plain voiceless stop series. Thus when he classifies a language as having exactly [p t k] word-initially, he means it has maximally a three-way major place distinction word-initially; that the language may also have e.g. the voiced stops [b d g] is not excluded.

<sup>4</sup> To be clear, I assume that substantive bias arises from the phonetic facts that make one phonological pattern more natural than another and not from the typology. Learners may possess implicit phonetic knowledge that could bias them toward some patterns over others, but they do not have implicit knowledge of the typology of the positional distribution of place contrasts in stops.

is consistent with substantive bias, they will extend the place contrasts more from word-final to word-initial position than from word-initial to word-final position.

### 3.1.1 Method

#### 3.1.1.1 Conditions

The first part of the experiment involved listening to words associated with pictures. There were two training conditions, Initial Contrast and Final Contrast. In the Initial Contrast condition, participants heard words beginning with labial, coronal, and dorsal stops (e.g. *pilan*, *tulir*, *kilun*), but in word-final position stops were only ever coronal (e.g. *wirut*). In the Final Contrast condition, participants heard words ending with labial, coronal, and dorsal stops (e.g. *rujap*, *wirut*, *wanuk*), but word-initial stops were always coronal (e.g. *tulir*).

I chose coronal place of articulation for the position in which stops did not contrast in place because coronal place is commonly considered unmarked (Kean 1975; Paradis & Prunet 1991; de Lacy 2006; Rice 2007). That said, Blevins (2004) disputes the special status of coronal place as the unmarked place and notes that, diachronically, when labial, coronal, and dorsal oral stops all neutralize it is generally to a glottal stop. Similarly, Lombardi (2001) observed that synchronic neutralization of place distinctions to glottal place (/h/ or /ʔ/) is common and proposed that /h/ and /ʔ/ are even less marked than relatively unmarked coronal sounds. De Lacy (2006) likewise considers glottal place to be the least marked place of articulation. Since my experiment participants were native English speakers, however, I used only English phonemes in the stimuli, which meant that the three-way major place contrast had to neutralize to labial, coronal, or dorsal place. Of these options, I chose coronal place as the least marked of the three.

O'Hara's (2021) typological study of the distribution of major place contrasts found that, of the 96 surveyed languages that have just a three-way major place contrast in stops word-initially, the vast majority either have all three places of articulation in stops word-finally as well or have no stops word-finally. O'Hara also found this all-or-nothing skew among languages that contrast more than three places of articulation for stops word-initially. In light of these surveys, the pattern exhibited in the artificial language of the Initial Contrast condition appears to be typologically rare, even though word-final place neutralization is common. In fact, the general case whereby word-initial place contrasts neutralize to a single place word-finally is not as rare as O'Hara's surveys would suggest because O'Hara specifically ignores [ʔ] in counting word-initial and word-final stop place contrasts. As Lombardi (2001) noted, word-final neutralization to glottal place occurs in many languages; Selayarese, for example, permits the stops /p b t d ʃ k g ʔ/ word-initially but only /ʔ/ word-finally (Mithun & Basri 1986). O'Hara's methodology therefore exaggerates the rarity of languages generally resembling the Initial Contrast language (even if specifically neutralizing a three-way place contrast to coronal place truly is rare).

On the other hand, expanding the typological picture to languages like Selayarese introduces further complications. Although [ʔ] is the stop with the least marked place of articulation (Lombardi 2001; de Lacy 2006), it has additional properties that distinguish it from other stops, such as its high sonority (de Lacy 2006). As a result, [ʔ] does not conform to the phonotactic implicational discussed in Section 2; there are languages that allow [ʔ] word-finally but not word-initially (de Lacy 2006). Its special behavior is precisely why O’Hara (2021) excluded [ʔ] from consideration in his typological surveys.

While Steriade (1994) argues that it is the relative perceptibility of contrasts that drives positional neutralization, the precise patterns in the crosslinguistic distribution of stop place contrasts may reflect additional forces (see Section 4 for further discussion). In Experiment 1, I wished to investigate whether the greater perceptibility of stop place of articulation in word-initial vs. word-final position affected learners’ acquisition of phonotactic systems, not whether a pattern’s typological frequency made it easier to learn. Typological complications notwithstanding, the artificial languages in the present study are the best means of testing for substantive bias rooted in the perceptibility of stop place contrasts with native English-speaking participants.

Returning to the conditions, then, the Initial Contrast language is natural in that it features major place contrasts in stops word-initially but not word-finally, a pattern consistent with the phonotactic implicational. The Final Contrast language, on the other hand, is unnatural: in having major place contrasts word-finally but not word-initially, it violates the implicational. **Table 1** shows which types of stops occurred in which positions in the two training conditions.

	#P	#T	#K	P#	T#	K#
<b>Initial Contrast</b>	✓	✓	✓	×	✓	×
<b>Final Contrast</b>	×	✓	×	✓	✓	✓

**Table 1:** Experiment 1 training conditions.

The test items were identical in the two conditions. In the test phase, participants in both the Initial Contrast and Final Contrast conditions heard words beginning and ending with labial, coronal, and dorsal stops.

### 3.1.1.2 Materials

The training and test items were all nonce words of the shape  $C_1VC_2VC_3$ . Either  $C_1$  or  $C_3$  was a stop drawn from [p t k b d g]. The other two Cs were sonorants drawn from [n l ɹ j w]. The glides [j] and [w] did not occur word-finally, and no item contained a [ji] or [wu] sequence.

In each word, all three Cs were different. The vowels were drawn from [i a u]. Labial, coronal, and dorsal stops were equally represented in the position with the place contrasts in training and across positions in test. Additionally, voiced and voiceless stops were equally represented across positions and places of articulation in both the training and test phases. In the Initial Contrast condition, half the training items with initial stops were members of minimal triplets (e.g. *bawan-dawan-gawan*). The same was true for training items with final stops in the Final Contrast condition (e.g. *lurip-lurit-lurik*). This was to encourage participants to notice that place of articulation was contrastive in stops in the given position. In the test phase, half of all items were members of minimal triplets for place of articulation. In both training and test, half of the items were iambs and half were trochees. Stress did not correlate with the position that featured the place contrasts or with the syllable in which the stop (always either word-initial or word-final) occurred. This was to prevent participants from associating the place contrasts with stress instead of position and to prevent stress from drawing undue attention to the stops. Stress was also not correlated with the voicing of the stop. **Table 2** shows some sample training items from the Initial Contrast condition. The full set of training and test stimuli for Experiment 1 are given in the Appendix.

#P	#T	#K	P#	T#	K#
pínir bilún ...	tínir dirúl ...	kínir gurúl ...		jáwit rujúd ...	

**Table 2:** Sample training items in the Initial Contrast condition.

The stimuli were recorded by a phonetically-trained male native speaker of American English who was naïve to the purpose of the experiment. Voiced stops were fully voiced, voiceless stops were aspirated, and word-final stops were released. Unstressed vowels were not reduced. Otherwise, pronunciation was as in American English. The stimuli were produced in isolation as if each item was a sentence unto itself. The stimuli were recorded in a sound-attenuated room using a head-mounted microphone. Recording was done using Audacity with a sampling rate of 22,050 Hz.

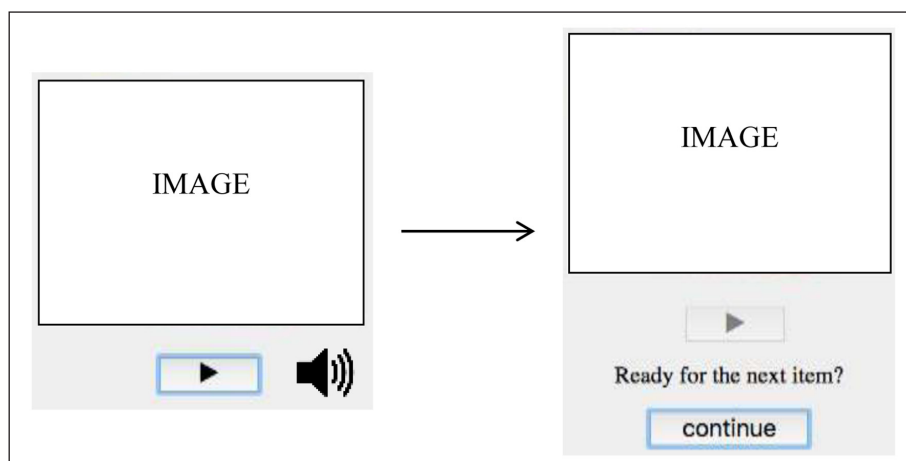
### 3.1.1.3 Procedure

The experiment was conducted online using Experigen (Becker & Levine 2013). All participants gave informed consent before proceeding with the experiment. Participants were instructed to wear headphones and to do the experiment in a quiet room. Written instructions informed them that they would be listening to some words of a new language. They would then be presented

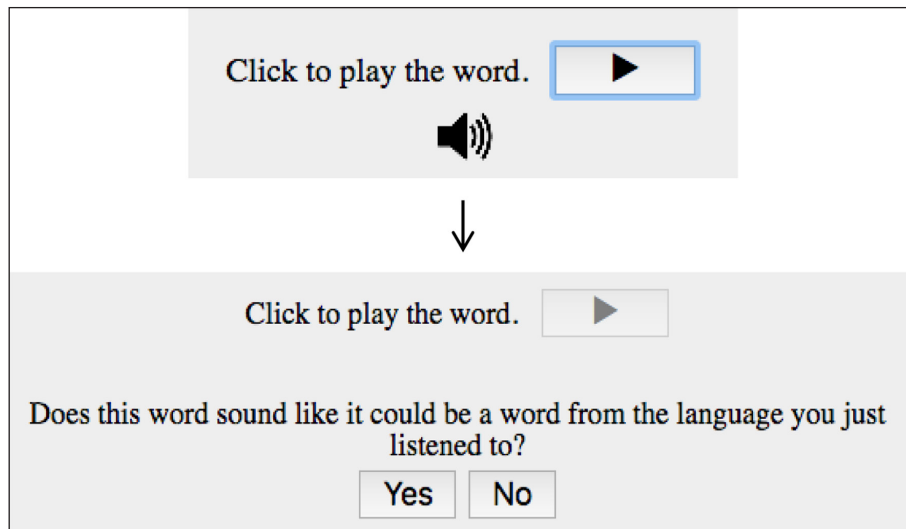


with additional words and asked whether those words sounded like they could also be from the language they had heard in the first part of the experiment. Before the experiment proper began, participants were prompted to play two sound files and type in the English word they had heard. The first word, *pad*, could be played multiple times so that participants could adjust the volume on their computer to a satisfactory level. The second word, *bat*, could only be played once. These test words were included to try to ensure participants were doing the experiment under acceptable listening conditions.

The experiment began with the training phase (see **Figure 1**). Participants proceeded at their own pace through two blocks of the same 40 training items. The order of the training items was randomized within each block for each participant, and there was no interruption between the two blocks. In each training trial, participants saw an image and clicked a button to hear the word for that image. Each sound file could be played only once. Participants then clicked to continue to the next training trial. They were encouraged to say the words out loud to help them learn the language. After the training phase, written instructions appeared telling participants that they would hear some additional words and should make their best guess as to whether each word sounded like it could also be a word from the language they had just listened to. There were no images in the test phase (see **Figure 2**). In each test trial, participants clicked a button to hear the sound file. They then had to click Yes or No to indicate whether they thought the word sounded like it could be from the language they had just listened to. There was a single test block consisting of 48 test items, which were the same in both conditions. Their order was randomized for each participant. At the end of the experiment, participants answered questions about their language background, whether they had noticed a pattern in their training language, and whether they had used a particular strategy in the test phase.



**Figure 1:** Event sequence for one training trial.



**Figure 2:** Event sequence for one test trial.

There were three types of item in the test phase. While the same set of test items was used in both conditions, which types test items fell into depended on the condition. Familiar Conforming items contained stops whose place of articulation and position conformed to the trained pattern (e.g. #P, #T, #K, and T# in the Initial Contrast condition) and were words that were heard in training. Novel Conforming items also contained stops whose place of articulation and position conformed to the trained pattern, but these words had not been heard in training. Finally, Novel Nonconforming items featured the place of articulation and position combinations not heard in training (e.g. P# and K# in the Initial Contrast condition). In each condition, there were 16 Familiar Conforming items, 16 Novel Conforming items, and 16 Novel Nonconforming items. **Table 3** gives sample test items of each type for the two training conditions.

	<b>Familiar Conforming</b>	<b>Novel Conforming</b>	<b>Novel Nonconforming</b>
<b>Initial Contrast</b>	pínir	panúl	nálup
<b>Final Contrast</b>	rujáp	nálup	panúl

**Table 3:** Sample test items for each training condition in Experiment 1.

### 3.1.1.4 Participants

The participants were native English speakers recruited through the UCLA Psychology Subject Pool. I excluded participants who reported that they were not in fact native English speakers (12), had taken more than one linguistics class (7), reported a history of speech or hearing impairments (4), gave an incorrect response to either of the two test words that preceded the experiment (0 in Experiment 1), or accepted all test items (2). After exclusions (25 out of 75 participants), there were 25 participants in each condition.

### 3.1.2 Predictions

While not central to the research question, a first prediction is that the acceptance rate of Familiar Conforming items should be greater than the acceptance rate of Novel Conforming items. In other words, participants should accept words they heard in the training phase more often than they accept words that fit their training pattern but are new to them.

Participants' acceptance rates of Novel Nonconforming items, relative to Novel Conforming items, indicate whether they have extended the place contrasts in stops to a new position in a given condition. For instance, if participants in the Initial Contrast condition accept test items with word-final labial and dorsal stops ( $P\#$  and  $K\#$ ), they have extended the word-initial place contrasts they encountered in training to word-final position. Recall the phonetic advantage favoring word-initial place contrasts: stop place of articulation is more perceptible at the beginning of a word than at the end of a word. On the basis of phonetic naturalness, then, a language with major place contrasts word-finally should also have major place contrasts word-initially, but not necessarily the other way around. Performance consistent with a substantive bias would be asymmetric extension: participants exposed to the place contrasts word-finally should extend them to word-initial position more than participants exposed to the place contrasts word-initially extend them to word-final position. Extension of the place contrasts to the other position manifests as erroneously accepting Novel Nonconforming items, so participants trained on the contrasts word-finally should more readily accept their Novel Nonconforming items ( $\#P$  and  $\#K$ ) than participants trained on the contrasts word-initially accept their own Novel Nonconforming items ( $P\#$  and  $K\#$ ).

Assuming similar performance in both conditions on Novel Conforming items, meaning that participants in both conditions learned to an equal degree what types of words did belong to their language, a higher acceptance rate of Novel Nonconforming items in the Final condition relative to the Initial Contrast condition would constitute evidence that learners are biased toward phonetically natural phonotactic systems. Put another way, if learning is substantively biased, participants should learn the unnatural Final Contrast language worse than the natural Initial Contrast language, as demonstrated by their mistakenly "filling in" the missing initial

labial and dorsal stops in their language.<sup>5</sup> If learners are not biased toward phonetically natural phonotactic systems, participants should accept Novel Nonconforming items at similar rates in both conditions since the two training patterns are of equal formal complexity.

### 3.1.3 Results

Figure 3 shows the acceptance rates of the three types of test item in each condition.<sup>6</sup> At a glance, it appears that participants behaved identically in both conditions.




Figure 3: Acceptance of test items by condition.

<sup>5</sup> A reviewer asked whether less successful learning of the Final Contrast language could also manifest as lower acceptance rates of Familiar Conforming or Novel Conforming items (i.e. less willingness to accept test items that fit the trained pattern). While not the focus of this paper, the general implementation of substantive bias that underpins my predictions here assumes there exist markedness constraints penalizing stops with (more) marked places of articulation (e.g. \*LAB, \*DORS). There also exist faithfulness constraints IDENT(place)/#\_ and IDENT(place)/\_# that respectively preserve place of articulation word-initially and word-finally. Learners' knowledge of the better perceptibility of stop place of articulation in initial position vs. final position causes IDENT(place)/#\_ to be obligatorily ranked above IDENT(place)/\_# (Steriade 2001b). The Final Contrast training language exhibits word-final labial, coronal, and dorsal stops, providing participants with evidence for the ranking IDENT(place)/\_# >> \*LAB, \*DORS. The bias-encoding ranking IDENT(place)/#\_ >> IDENT(place)/\_# would then, by transitivity, lead participants to erroneously accept words beginning with labial and dorsal stops too. This implementation does not predict any reluctance to accept words ending with labial, coronal, or dorsal stops. I therefore only expect a bias against the Final Contrast language to manifest as relatively higher acceptance of Novel Nonconforming items.

<sup>6</sup> Error bars represent  $\pm 1$  standard error.

I analyzed the results with mixed-effects logistic regressions using the `glmer()` function from the *lme4* package (Bates et al. 2015) in R. Because the factors of Familiarity and Conformity are not fully crossed, I carried out separate analyses of the conforming test items (both familiar and novel) and the novel test items (both conforming and nonconforming) (cf. Linzen & Gallagher 2017). I first fit a mixed-effects logistic regression to the conforming items (Familiar Conforming and Novel Conforming) with response (accept or reject) as the dependent variable, Familiarity (familiar vs. novel), Condition (Initial Contrast vs. Final Contrast), and their interaction as fixed effects, random intercepts for subject and item, and by-subject random slopes for Familiarity.<sup>7</sup> The model was dummy coded with novel as the reference level for Familiarity and Initial Contrast as the reference level for Condition. **Table 4** gives the fixed effects for this model. There was a significant main effect of Familiarity: Familiar Conforming items were accepted more often than Novel Conforming items ( $\beta = 1.526$ ;  $p < 0.001$ ). This is as predicted: participants should be better at accepting conforming items they heard in the training phase than conforming items that are new. The emergence of this expected effect confirms the experiment’s sensitivity to effects typically found in AGL studies. There was no significant effect of Condition, nor was there a significant interaction of Familiarity and Condition.

	$\beta$	SE	$z$	$p$
Intercept	0.547	0.214	2.552	0.011*
Familiarity = familiar (vs. novel)	1.526	0.291	5.242	<0.001***
Condition = Final Contrast (vs. Initial Contrast)	-0.167	0.266	-0.630	0.529
Familiarity $\times$ Condition	-0.157	0.339	-0.464	0.643

**Table 4:** Experiment 1—fixed effects of the conforming items model.

Turning now to novel items, the acceptance rates of Novel Conforming items were significantly above chance in both conditions, meaning that both Initial Contrast and Final Contrast participants correctly generalized to new words in their language. The acceptance rates of Novel Nonconforming items were significantly below chance in both conditions, meaning that both Initial Contrast and Final Contrast participants correctly rejected words

---

<sup>7</sup> I did not include by-item random slopes for Condition because although the test items were the same in both conditions, they did not fall into the same types in each condition. Thus the Familiar Conforming items in the Initial Contrast and Final Contrast conditions are not exactly the same sets of items, and the same is true for the Novel Conforming items.

not in their language.<sup>8</sup> This demonstrates successful phonotactic learning by participants in the experiment.

I fit a mixed-effects logistic regression to the novel items (Novel Conforming and Novel Nonconforming) with response (accept or reject) as the dependent variable, Conformity (conforming vs. nonconforming), Condition (Initial Contrast vs. Final Contrast), and their interaction as fixed effects, random intercepts for subject and item, and by-subject random slopes for Conformity. The model was dummy coded with nonconforming as the reference level for Conformity and Initial Contrast as the reference level for Condition. **Table 5** gives the fixed effects for this model. There was a significant main effect of Conformity: Novel Conforming items were accepted more often than Novel Nonconforming items ( $\beta = 1.094$ ;  $p < 0.001$ ). Of words they had never heard before, participants in both conditions preferred conforming items to nonconforming items. Crucially, this effect shows that participants did acquire the phonotactic pattern they were exposed to and that the experiment could detect this learning. There was no significant effect of Condition, and, of most interest, there was no significant interaction of Conformity and Condition. It is in this interaction that an effect of substantive bias would emerge: if learners favor phonetically natural phonotactic systems, there should be an interaction of Conformity and Condition such that participants in the Final Contrast condition accepted Novel Nonconforming items at a higher rate than participants in the Initial Contrast condition. Since this interaction is not significant, the results do not support substantive bias. Instead, participants learned the natural Initial Contrast language and unnatural Final Contrast language equally well, exhibiting no difference in the degree to which they “filled in” labial and dorsal stops in the position in which they had been withheld in training.

	$\beta$	SE	$z$	$p$
Intercept	-0.520	0.221	-2.354	0.019*
Conformity = conforming (vs. nonconforming)	1.094	0.241	4.534	<0.001***
Condition = Final Contrast (vs. Initial Contrast)	-0.057	0.309	-0.185	0.853
Conformity $\times$ Condition	-0.116	0.356	-0.326	0.744

**Table 5:** Experiment 1—fixed effects of the novel items model.

<sup>8</sup> Whether acceptance rates were significantly above chance was determined by fitting mixed-effects logistic regressions to the novel items (Novel Conforming and Novel Nonconforming) with the same random effects structure as in the model in Table 5 (see below) and changing the reference levels of the factors Condition and Conformity so that in each of the four models the intercept represented the acceptance rate for a different combination of Condition and Conformity. When the intercept represented Initial Contrast Novel Conforming or Final Contrast Novel Conforming items, it was significantly above chance (Initial Contrast:  $\beta = 0.575$ ,  $p = 0.004$ ; Final Contrast:  $\beta = 0.402$ ,  $p = 0.042$ ). When the intercept represented Initial Contrast Novel Nonconforming or Final Contrast Novel Nonconforming items, it was significantly below chance (Initial Contrast:  $\beta = -0.520$ ,  $p = 0.019$ ; Final Contrast:  $\beta = -0.577$ ,  $p = 0.009$ ).

In sum, participants did not fill in the place of articulation gaps in their training languages; rather, in both conditions, they successfully learned the gaps. The lack of interactions in the regression models means that there was no difference in participants' ability to generalize to new words or reject nonconforming words between the conditions.

### 3.1.4 Discussion

The participants in Experiment 1 did not reproduce the phonotactic implicational about the positional distribution of major place contrasts. That is, they did not extend the place contrasts more from word-final to word-initial position than from word-initial to word-final position. The experiment therefore yielded no evidence for substantive bias. Note that this was not because the experiment itself was entirely insensitive: it did uncover strong effects of familiarity and learning of the trained pattern. This suggests that if the relevant substantive bias does influence phonotactic learning, its effect must be much smaller than the effects the experiment was able to detect.

Another reason no substantive bias effect emerged could be that participants' learning was too explicit because the pattern was too easy to identify. Moreton & Pertsova (2016) demonstrated that learning in artificial phonotactic learning experiments can be more or less explicit depending on an experiment's design and procedure. Training with feedback, instructing participants to try to find a rule, and using "easily verbalizable" features all foster explicit learning while training without feedback, not mentioning rules in the instructions, and using features that are not easy to verbalize foster implicit learning (Moreton & Pertsova 2016: 277). In Experiment 1, there was no feedback of any kind in either the training phase or the test phase, and the instructions made no mention of rules. The feature of position (word-initial or word-final) is conceivably fairly easy to verbalize even for non-linguists, but segmental phonological features, including place of articulation, are likely much harder to verbalize. From the perspective of design, then, Experiment 1 looks more like experiments that attempt to foster implicit learning.

Moreton & Pertsova (2016) conducted an experiment in which participants had to learn to distinguish masculine and feminine words in an artificial language in which some phonological or semantic property was a perfect cue to gender (e.g. disyllabic words are masculine and trisyllabic words are feminine). There were two conditions: Feedback, whose design and procedure were meant to foster explicit learning, and No-Feedback, whose design and procedure were meant to foster implicit learning. Signs that participants are engaging in explicit learning include reporting seeking a rule and identifying the correct rule (Moreton & Pertsova 2016: 278). In the Feedback condition, 82% of participants reported seeking a rule during the training phase, and 61% reported using a rule to make their decisions during the test phase. In the No-Feedback condition, 56% of participants reported seeking a rule in training, and 43% reported using a rule in test. In my Experiment 1, participants were only asked about their strategy once, in the survey that followed the test phase. Most participants (17 out of 25 (68%) in the Initial Contrast condition and 19 out of 25 (76%) in

the Final Contrast condition) reported seeking a rule (defined as stating a connection between their responses and some property of the stimuli) to distinguish words that were in their language from words that were not in their language. The remaining participants reported a more intuitive strategy.<sup>9</sup> It is difficult to compare these survey results to those in Moreton & Pertsova's experiment due to differences in the training procedure and in the wording of the survey questions, but the extent of rule-seeking in Experiment 1 looks as though it lies in between the extent of rule-seeking in their Feedback (explicit learning) and No-Feedback (implicit learning) conditions. In their experiment, 14% of No-Feedback participants identified the correct rule while 29% of Feedback participants identified the correct rule. In Experiment 1, two participants (8%) found the correct rule in the Initial Contrast condition and four participants (16%) found the correct rule in the Final Contrast condition, where the correct rule was defined as one that mentioned labials and/or dorsals *and* position. In terms of the proportion of participants who identified the correct rule, Experiment 1 looks more like Moreton & Pertsova's No-Feedback condition, which was designed to promote implicit learning.

While Experiment 1 already incorporates design properties previously used to foster implicit learning and partly resembles Moreton & Pertsova's No-Feedback (implicit learning) condition in its results, Moreton & Pertsova demonstrated that participants may engage in explicit learning even in an experiment meant to encourage implicit learning. They also showed that changing properties of an experiment can affect what kind of learning participants engage in. Thus, even if Experiment 1 elicited a certain amount of implicit learning, it should still be possible to make learning even more implicit.

## 3.2 Experiment 2

I hypothesized that making the task in Experiment 1 harder might foster more implicit learning, which might in turn cause substantive bias to emerge. To test this, I carried out Experiment 2, a modified version of the positional extension of place contrasts experiment.

### 3.2.1 Method

Experiment 2 was identical to Experiment 1 except for three differences. First, I added 20 filler items in which all three consonants were sonorants (e.g. *lanir*) to the training phase. With the 40 critical items from Experiment 1, this yielded a total of 60 training items, one third of which were

---

<sup>9</sup> I tested for the substantive bias effect among the non-rule-seeking participants only (14 total; 8 in the Initial condition and 6 in the Final condition). The raw results appear promising: non-rule-seeking participants' mean acceptance rates of Novel Conforming items were similar (59% in the Initial condition and 61% in the Final condition), but the mean acceptance rate of Novel Nonconforming items was higher in the Final condition (47%) than in the Initial condition (37%). This looks like greater extension of the place contrasts from final position to initial position than vice versa. However, when I fit a mixed-effects logistic regression to the novel items (the same type of model as in Table 5), the interaction of Conformity and Condition that would demonstrate asymmetric extension and therefore support substantive bias was not significant ( $\beta = 0.566$ ;  $p = 0.323$ ). This is likely due to the small number of participants falling in the implicit learning category.



fillers. The all-sonorant fillers were intended to disguise the pattern, making it less obvious and thus leading to more implicit learning. The filler items were recorded by the same speaker who recorded the stimuli for Experiment 1. The full set of training and test stimuli for Experiment 2 are given in the Appendix.

Second, I increased the number of training blocks from two to three. This was done because the acceptance rate of Novel Conforming items in the Final Contrast condition was just barely significantly above chance in Experiment 1. Acceptance rates of Novel Conforming items must be above chance to show that participants have learned to accept the types of items they were trained on. Without this demonstration of learning, performance on Novel Nonconforming items is harder to interpret. I therefore increased the number of training blocks in Experiment 2 to try to push Novel Conforming acceptance rates up.

Finally, I removed the Familiar Conforming items from the test phase, leaving 32 test items in each condition.<sup>10</sup> This was done because after implementing the two changes above, pilot data showed no increase in the acceptance rates of Novel Conforming items. In particular, increasing the number of training blocks from two to three was not having the desired effect of boosting Novel Conforming acceptance rates. The presence of the Familiar Conforming items in the test phase may have been keeping down the acceptance rates of Novel Conforming items, since they constituted a class of test items that were even better (i.e. more acceptable) than the Novel Conforming items. In other words, the inclusion of test items familiar from training may have imposed a ceiling on the acceptance rate of items that still fit the training pattern but were unfamiliar. As the results will show, eliminating the Familiar Conforming items did lead to higher acceptance rates for Novel Conforming items.

The participants in Experiment 2 were native English speakers recruited through the UCLA Psychology Subject Pool. I applied the same exclusion criteria as in Experiment 1 and excluded 25 out of 67 participants. Of the 25 excluded participants, 10 were not in fact native English speakers, 7 had taken more than one linguistics class, 1 reported a history of speech or hearing impairments, and 7 gave an incorrect response to either of the two test words that preceded the experiment. After exclusions, there were 21 participants in each condition. None of the participants in Experiment 2 had participated in Experiment 1.

### 3.2.2 Predictions

The predictions for Experiment 2 are the same as the predictions for Experiment 1, except that Familiar Conforming items are not predicted to be accepted more than Novel Conforming items because there are no Familiar Conforming items. If learners are biased toward phonetically natural phonotactic systems, participants in the Final Contrast condition should more readily


---

<sup>10</sup> Which test items were Familiar Conforming differed between the two conditions, so removing the Familiar Conforming items meant that the test items were no longer identical in the two conditions in Experiment 2.

accept Novel Nonconforming items than participants in the Initial Contrast condition. This would show that learners are more likely to assume place contrasts exist word-initially if they have been taught they exist word-finally than they are to assume place contrasts exist word-finally if they have been taught they exist word-initially. If there is no substantive bias at work, acceptance rates of Novel Nonconforming items should not differ between the two conditions.

### 3.2.3 Results

**Figure 4** shows the acceptance rates of the two types of test items, Novel Conforming and Novel Nonconforming, in each condition. As mentioned above, removing the Familiar Conforming items from the test phase caused acceptance rates of both Novel Conforming and Novel Nonconforming items to increase relative to Experiment 1, which included Familiar Conforming items. In Experiment 1, the mean acceptance rate of Novel Conforming items was 61% in the Initial Contrast condition and 58% in the Final Contrast condition while in Experiment 2 the mean acceptance rate of Novel Conforming items was 66% in the Initial Contrast condition and 73% in the Final Contrast condition. Eliminating the Familiar Conforming test items therefore had the desired effect of pushing up acceptance rates of Novel Conforming items. In Experiment 2, these acceptance rates were significantly above chance in both conditions.<sup>11</sup> This demonstrates successful learning: participants were able to generalize to new words that matched the pattern they had been trained on.



**Figure 4:** Acceptance of test items by condition.

<sup>11</sup> This was determined by fitting a mixed-effects logistic regression to the test items (Novel Conforming and Novel Nonconforming) and changing the reference levels of the factors Condition and Conformity so that the model intercept represented the acceptance rate for the relevant combination of Condition and Conformity. When the intercept represented Initial Contrast Novel Conforming items or Final Contrast Novel Conforming items, it was significantly above chance (Initial Contrast:  $\beta = 0.739$ ,  $p < 0.001$ ; Final Contrast:  $\beta = 1.028$ ,  $p < 0.001$ ).

For Novel Nonconforming items, the mean acceptance rates in Experiment 1 were 39% in the Initial Contrast condition and 38% in the Final Contrast condition while the mean acceptance rates in Experiment 2 were 50% in the Initial Contrast condition and 48% in the Final Contrast condition. Eliminating the Familiar Conforming test items also pushed up acceptance rates of Novel Nonconforming items. In Experiment 1, Novel Nonconforming items were accepted at rates significantly below chance in both conditions, but in Experiment 2, the acceptance rates of Novel Nonconforming items did not differ from chance in either condition.<sup>12</sup> In other words, while participants in Experiment 1 were able to correctly reject words not in their language at above chance levels, participants in Experiment 2 were not. These participants did not go as far as filling in the place of articulation gaps in their training languages; rather, they did not know whether those gaps were there or not and behaved as though they were guessing in the test phase.

To test for asymmetric positional extension of the place contrasts, I fit a mixed-effects logistic regression to the test items (Novel Conforming and Novel Nonconforming) with response (accept or reject) as the dependent variable, Conformity (conforming vs. nonconforming), Condition (Initial Contrast vs. Final Contrast), and their interaction as fixed effects, random intercepts for subject and item, and by-subject random slopes for Conformity. **Table 6** gives the fixed effects for this model. There was a significant main effect of Conformity: Novel Conforming items were accepted more often than Novel Nonconforming items ( $\beta = 0.771$ ;  $p = 0.022$ ). As in Experiment 1, participants preferred novel words that fit the pattern they had been trained on to novel words that did not. There was no main effect of Condition and no interaction of Conformity and Condition.

	$\beta$	SE	$z$	$p$
Intercept	-0.032	0.268	-0.121	0.904
Conformity = conforming (vs. nonconforming)	0.771	0.337	2.286	0.022*
Condition = Final Contrast (vs. Initial Contrast)	-0.031	0.375	-0.081	0.935
Conformity $\times$ Condition	0.319	0.501	0.637	0.524

**Table 6:** Experiment 2—fixed effects of the model.

---

<sup>12</sup> When the intercept of the model described in fn. 11 represented Initial Contrast Novel Nonconforming items or Final Contrast Novel Nonconforming items, it did not differ significantly from chance (Initial Contrast:  $\beta = -0.032$ ,  $p = 0.904$ ; Final Contrast:  $\beta = -0.063$ ,  $p = 0.814$ ).

### 3.2.4 Discussion

Like Experiment 1, Experiment 2 found robust learning of the phonotactic patterns: participants clearly distinguished between conforming and nonconforming novel words. Experiment 2 also resembles Experiment 1 in its lack of a significant interaction of Conformity and Condition, meaning that there was no asymmetric extension of the place contrasts. Participants in the Final Contrast condition did not extend the place contrasts more from word-final to word-initial position than participants in the Initial Contrast condition extended the place contrasts from word-initial to word-final position. Experiment 2 thus did not yield evidence for substantive bias either. Though the actual acceptance rates of Novel Conforming and Novel Nonconforming items changed from Experiment 1 to Experiment 2 due to the removal of the Familiar Conforming items, the natural Initial Contrast language and the unnatural Final Contrast language were still learned equally well.

The main motivation for running Experiment 2 was to try to make participants' learning more implicit and thereby allow a substantive bias to emerge. To foster implicit learning, I added filler items to the training phase in hopes of making the distribution of labial, coronal, and dorsal stops less obvious. One might ask whether this modification actually caused more implicit learning in Experiment 2 as compared to Experiment 1. In Experiment 2, roughly half of participants (10 out of 21 (48%) in the Initial Contrast condition and 11 out of 21 (53%) in the Final Contrast condition) reported seeking a rule to distinguish words in their language from those not in their language. Only one participant in each condition found the correct rule. In Experiment 1, two thirds (Initial Contrast condition) or three quarters (Final Contrast condition) of participants reported seeking a rule, so there does seem to have been somewhat more implicit learning in Experiment 2. This did not lead to the emergence of substantive bias, however. The natural and unnatural patterns were still learned to an equal degree.

The basis of this study was the phonotactic implicational whereby major place contrasts in post-vocalic position entail major place contrasts in pre-vocalic position, but not vice versa. Together, Experiments 1 and 2 show that this implicational is not reproduced in an AGL paradigm. Instead, participants learned equally well a natural language in which place was contrasted in stops only word-initially and an unnatural language in which place was contrasted in stops only word-finally. Neither experiment's results provide support for substantive bias. Though Experiment 2 may have successfully encouraged more implicit learning, it still did not turn up evidence for substantive bias.

## 4. The positional extension of place contrasts vs. a voicing contrast

Experiments 1 and 2 tested for substantive bias in phonotactic learning by investigating whether learners were more likely to extend phonemic contrasts from a position in which they were less perceptible to a position in which they were more perceptible than vice versa. Such asymmetric

extension would support the hypothesis that phonetic naturalness biases phonological learning, but it was absent in the experiments. Interestingly, Glewwe (2021) conducted an analogous set of experiments examining a different phonemic contrast, namely, the voicing contrast in stops, and obtained different results. A parallel phonotactic implicational holds for stop voicing: a word-final stop voicing contrast entails a word-initial contrast, but the opposite is not necessarily true. This implicational is perceptually motivated in the same way the implicational about major place contrasts is.


The experiments in Glewwe (2021) had the same design as Experiment 1: the training phase exposed participants to a stop voicing contrast in either word-initial or word-final position, and the test phase invited participants to extend the stop voicing contrast to the other word-edge position. While the experiments in the present study simply had Initial Contrast and Final Contrast conditions, the design in Glewwe (2021) additionally manipulated which voicing value ([+voice] or [-voice]) stops neutralized to in the word-edge position without the voicing contrast. This resulted in four training conditions: the neutralizing-to-voiced conditions # {T, D}...D# (initial contrast) and #D...{T, D}# (final contrast) and the neutralizing-to-voiceless conditions # {T, D}...T# (initial contrast) and #T...{T, D}# (final contrast). Glewwe's (2021) three experiments also varied in the types of filler consonants the stimuli contained, with effects on the results.<sup>13</sup> However, the core research question was the same as in the current study: are learners more inclined to extend phonemic contrasts in stops from word-final position, where they are less perceptible, to word-initial position, where they are more perceptible, than from word-initial position to word-final position?

**Figure 5** shows the acceptance rates of Novel Nonconforming items for each condition in Glewwe's (2021) three experiments. Acceptance rates of Novel Conforming items are not shown, but in each experiment, they were significantly above chance and did not differ significantly between conditions. This makes it possible to test for asymmetric extension of the stop voicing contrast by comparing the acceptance rates of Novel Nonconforming items. Substantive bias predicts that participants in final contrast conditions should extend the voicing contrast to word-initial position more than participants in corresponding initial contrast conditions extend the contrast to word-final position; this would manifest as higher acceptance rates of Novel Nonconforming items in final contrast conditions than in corresponding initial contrast conditions. In **Figure 5**, each pair of solid and striped bars represents one final contrast condition (solid) and its corresponding initial contrast condition (striped).<sup>14</sup> The arrows point from the

<sup>13</sup> In Experiment 1, like in the place contrasts experiments, the non-stop consonants were voiced sonorants (e.g. *pímir*). In Experiment 2, the non-stop consonants were voiceless fricatives (e.g. *pífis*), and in Experiment 3, they included both voiced sonorants and voiceless fricatives (e.g. *dílur*, *kaśáf*).

<sup>14</sup> In these experiments, the number of participants in each condition ranged from 29 to 41, which is somewhat higher than in the place contrasts experiments, but the standard errors of the acceptance rates were similar between Glewwe (2021) and the present study.

acceptance rate that is predicted to be higher to the rate that is predicted to be lower. Across the three experiments, the acceptance rate of Novel Nonconforming items is higher for the final contrast condition than the initial contrast condition in five out of six cases. For comparison, in Experiment 1 of the present study, the acceptance rates of Novel Nonconforming items were 38% in the Final Contrast condition vs. 39% in the Initial Contrast condition, and in Experiment 2, they were 48% in the Final Contrast condition vs. 50% in the Initial Contrast condition. That is, the acceptance rates of Novel Nonconforming items were virtually identical (and numerically in the direction opposite that predicted by substantive bias).



**Figure 5:** Acceptance of Novel Nonconforming items by experiment and condition.

In Glewwe's (2021) study, extension of the stop voicing contrast was significantly greater from word-final position to word-initial position than from word-initial to word-final position in one of the three experiments (Experiment 2; see **Figure 5**), and a meta-analysis found a significant effect when the results of all three experiments were pooled. These results contrast with those of the present study, which found no asymmetric extension of major place contrasts in stops. It is unlikely that a third place contrasts experiment and a meta-analysis would uncover a substantive bias effect given that there was virtually no difference in the acceptance rates of Novel Nonconforming items between the Initial Contrast and Final Contrast conditions in Experiments 1 and 2. Additionally, the interaction testing for asymmetric extension in the present study's analyses never approached significance and was numerically in the wrong direction in Experiment 2.

Given the great similarity between the two sets of experiments, these divergent results are surprising. Glewwe (2021) did uncover evidence of perceptually-grounded substantive bias: though the effect appears subtle, learners were more likely to infer that the stop voicing contrast existed in the other word-edge position when not doing so meant accepting a phonetically unnatural phonotactic pattern. In the current study, on the other hand, learners were equally happy to accept natural and unnatural distributions of major place contrasts in stops. That is, their behavior showed no effect of substantive bias. In both studies, the hypothesized substantive bias was rooted in perceptual facts, so Glewwe et al.'s (2018) proposal that only perceptual naturalness, and not articulatory naturalness, synchronically biases learners cannot account for these different results.

One possible explanation for why a substantive bias effect emerged in the voicing contrast study but not in the place contrasts study would be a difference between voicing and place of articulation in terms of the size of the perceptibility differential across positions. If the difference between the perceptibility of stop voicing word-initially and word-finally is greater than the difference between the perceptibility of stop place word-initially and word-finally, then having a contrast only in word-final position, where it is harder to hear, should be more unnatural in the case of voicing than in the case of place of articulation, leading to more positional extension of the voicing contrast. If the perceptibility differential across word-edge positions is smaller for place of articulation, meaning that place contrasts in stops are only marginally harder to hear word-finally than word-initially, then there may be less pressure to avoid the unnatural sound pattern via positional extension, resulting in no effect in the present study.

White (2013) offered a similar explanation for a certain result his Experiment 1 obtained. He argued that learners are biased against saltation, an unnatural phonological pattern in which more distant sounds (e.g. [p]-[v]) alternate with each other while less distant sounds (e.g. [b]-[v], [f]-[v]) do not, because participants trained on a rule in which /p/ became [v] intervocally tended to assume that /b/ and /f/ also became [v].<sup>15</sup> This result could be due to complexity bias (Moreton & Pater 2012a) since /p/ → [v] is a two-feature change while /b/ → [v] and /f/ → [v] are only one-feature changes, except for the fact that participants were more likely to assume that /b/ became [v] than they were to assume that /f/ did. Since [b]~[v] and [f]~[v] alternations both involve only one feature change, complexity cannot explain this difference. Instead, White extracted confusion matrix data from Wang & Bilger 1973 that showed that [b] and [v] were more perceptually similar (confused more often) than [f] and [v]. Thus not only were learners biased against saltation, which could be attributed to complexity bias, but how much they extended to alternations that related perceptually similar sounds depended on *how* perceptually similar those sounds were, which could only reflect an effect of phonetic

---

<sup>15</sup> White's experiment also included coronal stops and fricatives [t d θ ð], but I use the labial case to exemplify.

substance. To be clear, participants were significantly inclined to “fill in” all the rules that would make the phonological system natural (i.e. /b/ → [v] and /f/ → [v]), but they filled in a rule more when the two sounds it would relate were more similar (that is, when *not* having the rule would arguably be more unnatural). White’s overall point was that there must still be a role for substantive bias, but his fine-grained results suggest that the strength of a substantive bias might depend on the magnitude of its phonetic precursor.

For this type of explanation to account for the difference between the current study and Glewwe’s (2021) study, there would have to be evidence that the decrease in perceptibility from word-initial position to word-final position actually is greater for the stop voicing contrast than it is for stop place contrasts. Confusion studies (e.g. Miller & Nicely 1955; Wang & Bilger 1973) can shed light on the relative perceptibility of different phonological features by measuring the confusability (and therefore the perceptual similarity) of sounds differing in those features. In order to induce errors, however, these studies typically present stimuli in noise. Noise appears to harm the perception of place of articulation more than the perception of voicing (Cole & Jakimik & Cooper 1978; Martin & Peperkamp 2017), so confusion data collected using masking noise may not provide an accurate picture of the actual relative perceptibility of place and voicing.

Wang & Bilger (1973) did conduct one confusion study in which stimuli were presented at different signal levels but without masking noise. They tested both CV and VC syllables, making it possible to compare the perceptibility of stop voicing and place of articulation across position (word-initial vs. word-final). **Tables 7** and **8** reproduce the confusion data from Wang & Bilger’s (1973) Tables 6 and 7 for the sounds relevant to the place contrasts and voicing contrast AGL experiments (/p t k b d g/) in the CV and VC contexts, respectively. Cells representing confusion of place of articulation are shaded while cells representing confusion of voicing are outlined in black.

As **Tables 7** and **8** show, confusing the voicing of a stop is generally more common in the word-final context than in the word-initial context, though this difference is driven by word-final voiced stops being confused for their voiceless counterparts. For place of articulation, there is no clear difference in consonant confusions between the CV and VC contexts. As a rough means of quantifying the difference in confusability of stop voicing and stop place across positions, I took the ratio of consonant confusions in final position to consonant confusions in initial position for both contrasts. For voicing, there were 151 confusions in the VC context and 85 in the CV context, yielding a ratio of 1.8. For place of articulation, there were 419 confusions in the VC context and 462 in the CV context, for a ratio of 0.9. These ratios suggest that the confusability of stop voicing increases much more from word-initial position to word-final position than does the confusability of stop place. There appears to be a substantial difference in the perceptibility of voicing across positions while there is little difference in the perceptibility of place of articulation. Indeed, in Wang & Bilger’s study, confusions for stop place were slightly more frequent in the CV context than in the VC context.



	p	t	k	b	d	g
p	773	38	35	9	14	6
t	33	783	27	7	19	6
k	71	89	585	11	14	18
b	21	11	8	587	28	17
d	9	10	4	32	771	24
g	4	4	8	27	41	764

**Table 7:** Consonant confusions in the CV context (added across all signal levels).

	p	t	k	b	d	g
p	763	23	80	9	3	10
t	25	795	31	1	9	6
k	66	36	753	2	7	18
b	55	12	28	608	19	87
d	9	16	7	10	763	13
g	16	10	44	22	7	762

**Table 8:** Consonant confusions in the VC context (added across all signal levels).

These confusion data do provide evidence that the decrease in perceptibility from word-initial position to word-final position is greater for stop voicing than for stop place, which in turn could explain why substantive bias emerged in Glewwe’s (2021) voicing contrast experiments but not in the place contrasts experiments reported on here. In the case of the present study, the difference in perceptibility of stop place of articulation word-initially vs. word-finally may not be large enough to drive more extension of stop place contrasts from final to initial position than from initial to final position. The substantive bias effect found in the voicing contrast study was already subtle, so if position affects the perceptibility of place of articulation less, the effect of substantive bias may be undetectable or even absent.

The evidence from Wang & Bilger should be treated with caution, though, because as a reviewer points out, they considered the confusability of these stops not just amongst themselves but with all the obstruents of English. On the one hand, based on the sounds they were exposed to in training and presented with in Novel Nonconforming test items, participants in the AGL experiments only needed to (implicitly) compare the perceptibility of voicing within stops with the same place of articulation or the perceptibility of place of articulation within stops with the

same voicing. On the other hand, when relying on Wang & Bilger's confusion study, it may be necessary to take into account, for example, confusions between [p] and [v] to capture the full picture of the perceptibility of voicing, or between [p] and [d] for the perceptibility of place of articulation. If voicing and place confusions that also involve changes in other features are included, the degree to which stop voicing is more perceptible word-initially vs. word-finally is much diminished, though stop place remains surprisingly more confusable word-initially than word-finally.<sup>16</sup>

Further perceptual studies are needed to determine whether the decrease in perceptibility from word-initial to word-final position really is more severe for stop voicing than stop place of articulation. These investigations might include targeted confusion studies or difficult discrimination tasks like Martin & Peperkamp's (2017) that induce enough errors (without masking noise) to allow for comparison between different features. More generally, conducting more phonetic studies conceived expressly to quantify phonetic precursors may yield new explanations for some of the uneven results in the experimental literature. Moreover, such studies could generate new predictions for future AGL studies. All of this would enrich substantive bias research.

I now briefly consider some other possible reasons for the null substantive bias result in Experiments 1 and 2 vs. the positive result in the voicing contrast study, as well as the broader issues those reasons bring to the fore. The first reason is related to the difference between stop place and voicing discussed above but additionally takes language experience into account. That is, what if position has relatively little effect on the perceptibility of stop place of articulation specifically for native English speakers, and this is why Experiments 1 and 2 did not find an effect? Traditional conceptions of substantive bias within Optimality Theory encode the preference for natural phonology in an innate universal constraint set, which contains some constraints but not others, in accordance with universal phonetic facts (Hale & Reiss 2000; Greenwood 2016).<sup>17</sup> Under this view, substantive bias predictions and effects should not vary across languages. Alternatively, the phonetic knowledge that biases learners toward natural patterns could itself be learned (Hayes & Wilson 2008: 425; White 2013: 167). For White (2013), the phonetic knowledge that drives substantive bias comes from the P(erceptual)-map (Steriade 2001b), a repository, within the grammar, of information about perceptual similarity. For instance, the P-map would encode the fact that stops differing in place of articulation are more perceptually similar word-finally than word-initially. Interestingly, Steriade (2001b) leaves open the possibility that the

---

<sup>16</sup> This is echoed by Wang & Bilger's (1973) calculations for percent information transmitted by obstruent voicing and place when there was no masking noise (see their Tables 12 and 13). Percent information transmitted for voicing was 57.1% in the CV context and 54.5% in the VC context, giving voicing just a slight advantage word-initially. Percent information transmitted for place (when treated as a 3-way contrast between labial, coronal, and dorsal) was 50.6% in the CV context but rose to 56.5% in the VC context.

<sup>17</sup> Substantive bias could also be encoded via fixed rankings (or relative weightings) of constraints (see fn. 5).

details of the perceptual knowledge stored in the P-map may be language-dependent rather than universal. To exemplify with a different phonological contrast, what the P-map says about the perceptual similarity of voiced and voiceless stops in different positions may depend on which cues (release cues, closure cues, etc.) listeners rely on most to identify stop voicing in a particular language. If the P-map can vary crosslinguistically in this way, then the substantive bias predictions it generates may be stronger or weaker for different languages.

In addition to raw acoustic cues, other factors that might shape the perceptibility of phonological contrasts in a given language could potentially affect the strength of language-specific substantive biases too. For instance, a contrast's importance for differentiating words in a particular language (i.e. its functional load in that language) may affect its distinctiveness for native speakers (Cole & Jakimik & Cooper 1978; Martin & Peperkamp 2017). A contrast's functional load could also differ by position (e.g. word-initial vs. word-final), with consequences for position-specific perceptibility, and the degree of such differences could vary by language.

I do not expect any of these language-dependent factors to ever generate predictions that a phonetically unnatural pattern should be favored over a natural one. It seems unlikely that the functional load of a particular contrast would be significantly greater in a position where, acoustically, it was harder to perceive than in a position where it was easier to perceive precisely because, whatever the reason, phonological systems tend to be phonetically natural. Moreover, even if a contrast's functional load were to boost its perceptibility in an unexpected position, this would likely not override the influence of raw acoustic cues. With respect to the P-map, universal properties of production and perception will ensure that the general phonetic knowledge it contains is largely the same across languages; what can vary are details like how much the perceptibility of a contrast improves from one position to another.

In sum, universal phonetic precursors have the potential to give rise to substantive biases, which may then play a role in shaping the phonological typology, but the phonetic implementation (of e.g. various contrasts) in specific languages may affect the strength of those biases in native speakers if the phonetic knowledge underpinning them is learned. This may be especially true in adult speakers, who are the most frequent participants in AGL studies. If the perceptibility of major place contrasts in stops is less influenced by position specifically in English, it might explain why Experiments 1 and 2 did not uncover a naturalness bias. This discussion points to the importance of conducting experiments examining the same hypothesized substantive biases in different languages with distinct phonetic implementations. Positive results would show that a substantive bias can become active in learners with access to the right phonetic knowledge, which in turn would mean that that bias could potentially have influenced the phonological typology.

Finally, I will touch on one last possible reason for the difference between the place contrasts and voicing contrast studies, namely, that the explanation is in some way connected to the typology. Half of the initial contrast conditions in Glewwe's (2021) voicing experiments (i.e.

the  $\#\{T, D\}\dots T\#$  conditions) exhibited a typologically very common pattern (Blevins 2004): they were final devoicing languages in phonotactic form. While the neutralizing-to-voiced initial contrast conditions ( $\#\{T, D\}\dots D\#$ ) represented a virtually unattested pattern (Blevins 2004), at least positional neutralization of stop voicing is commonplace. As previously mentioned, the Initial Contrast conditions in the present study exhibited a typologically rare pattern: languages with major place contrasts in stops word-initially rarely neutralize them to a non-glottal place word-finally. While I contend that substantive bias arises from phonetic precursors rather than typological asymmetries, it is conceivable that some overlooked factor related to the weaker typological support for the design of Experiments 1 and 2 is at the root of their null substantive bias results. What that factor might be is unclear. That said, there may be reasons unrelated to the perceptibility of stop place contrasts for the tendency for languages to allow either all their initial stops or none of them in word-final position (O’Hara’s (2021) all-or-nothing skew). Of the languages in O’Hara’s sample that allow none, some simply ban all word-final consonants (e.g. Fijian (Blevins 2004)). Others, in line with the crosslinguistic preference for more sonorous codas (de Lacy 2006), only allow word-final consonants with higher sonority than stops (e.g. Mandarin (Lee & Zee 2003)). It therefore seems plausible that many languages that contrast stops at multiple places of articulation word-initially but prohibit all these stops word-finally do so for independent reasons and not because stop place contrasts are harder to perceive word-finally. This of course does not explain the “all” side of the skew, that is, why languages that *do* allow word-final stops rarely neutralize them to a single non-glottal place. In this case, the answer might have to do with perceptibility. If, as I hypothesized above, the perceptibility of stop place is less diminished in word-final position than that of stop voicing (and not just in English), there could simply be less pressure (from channel bias or substantive bias) to neutralize stop place contrasts word-finally, so fewer such languages arise in the typology.

## 5. Conclusion

I conducted an AGL study that tested for substantive bias in phonotactic learning by examining whether participants exposed to place of articulation contrasts in stops in a given position were more likely to extend those contrasts to a new position where they are more perceptible than to a new position where they are less perceptible. Although the experiments found strong effects of phonotactic learning, participants extended the place contrasts equally in both directions, yielding no support for substantive bias. This lack of an effect might be taken as another entry in the long list of null results that have cast doubt on the substantive bias hypothesis. A closely related study, however, addressed the same question by examining extension of the stop voicing contrast and did uncover an effect of substantive bias, suggesting that it was something specific about stop place of articulation that led to the null finding in this study.

It is possible that the degree to which perceptibility is diminished from word-initial position to word-final position is smaller for stop place of articulation than for stop voicing; preliminary evidence from confusion data supplies some support for this idea. This would reduce the relative unnaturalness of only having place contrasts word-finally, thereby resulting in more symmetrical extension of the place contrasts across positions. The present study reinforces the elusiveness of substantive bias and may support the hypothesis that its role in shaping phonological grammars is subtle. In combination with the voicing contrast study, though, it points to a more nuanced hypothesis: the strength of a substantive bias (or whether it exists at all) varies with the magnitude of its phonetic precursor.<sup>18</sup> Substantive biases may need relatively extreme phonetic precursors in order to be detectible in an AGL experiment, which might explain the frequency of null results. (The effect of substantive bias could be larger in natural language acquisition, where all learning is presumably implicit.) Testing this hypothesis will require conducting phonetic studies to precisely quantify differences in the size of phonetic precursors in order to establish clear predictions about which substantive biases should emerge more or less strongly. How to estimate and compare the magnitude of different phonetic precursors and the extent to which this is feasible (e.g. across perception and articulation) are questions for future research. Confusion studies and discrimination tasks could be used to compare perceptual similarity in different contexts. Moreton (2008; 2010) drew on instrumental studies measuring F1 or F0 to argue for the equal magnitude of phonetic precursors to various phonological dependencies involving vowel height, consonant voicing, and/or tone; his approach could serve as a starting point for quantifying and comparing articulatory precursors.

In yielding null substantive bias results while very similar experiments did not, the two experiments reported on here suggest a promising frontier in the further characterization of the nature and scope of substantive bias. Taking the magnitude of phonetic precursors into account may shed new light on existing AGL results and can furnish novel predictions for future experiments. If this line of inquiry proves fruitful, it could eventually help determine which phonetically natural tendencies in the typology might stem in part from substantive bias and which can be attributed entirely to channel bias.

---

<sup>18</sup> This is different from Archangeli & Pulleyblank's (1994) hypothesis that the typological prevalence of a phonological pattern is correlated with the robustness of its phonetic precursor. In a sense, it is a more circumscribed hypothesis since it makes predictions only about substantive bias strength and not about the typology, which can also be shaped by channel bias and complexity bias (Moreton 2008; 2012).

## Appendix

The stimuli are given in a broad IPA transcription (<j> represents [j], <a> represents [a], and <r> represents [ɾ]). Shaded items are familiar (heard in training and test). Bolded items belong to minimal triplets.

### Experiment 1 Initial Contrast Condition Stimuli

Training	#P	#T	#K	P#	T#	K#
	pínir pilán pulír pírul bínar bilún bawán búníl	tínir tajál tulír túwan dínar dirúl dawán dáníl	kínir kilún kulír kíral gínar gurúl gawán gírul		jáwit wirút lurít lánat nílad julád rujúd nírad  wíjat nilút wajút rúwat wínud nalíd jarúd lúnid	
Test	#P	#T	#K	P#	T#	K#
	pínir pilán panúl páwir bawán búníl búlar burún	tínir tajál tanúl túrul dawán dáníl dúlar dajún	kínir kilún kanúl kájar gawán gírul gúlar galín	nilíp nálup jávip rujáp wánub warfb rujúb jávab	nilít níjat jáwit wirút wánud luwíd rujúd nírad	nilík líruk jávik wanúk wánug nurág rujúg línig

### Experiment 1 Final Contrast Condition Stimuli

Training	#P	#T	#K	P#	T#	K#
		<b>tínir</b> tajál tulír túwan <b>dínar</b> dirúl <b>dawán</b> <b>dánil</b>  tíjar tinúl talín tújal dánir dalúr duján díwil		<b>jáwip</b> <b>rujáp</b> <b>luríp</b> júnup <b>nílab</b> lijúb <b>rujúb</b> <b>jáwab</b>	<b>jáwit</b> <b>wirút</b> <b>lurít</b> lánat <b>nílad</b> julád <b>rujúd</b> <b>nírad</b>	<b>jáwik</b> <b>wanúk</b> <b>lurík</b> nálak <b>nílag</b> jalág <b>rujúg</b> <b>línig</b>
Test	#P	#T	#K	P#	T#	K#
	<b>pínir</b> pilán <b>panúl</b> páwir <b>bawán</b> búnil <b>búlar</b> burún	<b>tínir</b> <b>tajál</b> <b>tanúl</b> túrul <b>dawán</b> <b>dánil</b> <b>dúlar</b> dajún	<b>kínir</b> kilún <b>kanúl</b> kájar <b>gawán</b> gírul <b>gúlar</b> galín	<b>nilíp</b> nálup <b>jáwip</b> <b>rujáp</b> <b>wánub</b> waríb <b>rujúb</b> <b>jáwab</b>	<b>nilít</b> níjat <b>jáwit</b> <b>wirút</b> <b>wánud</b> luwíd <b>rujúd</b> <b>nírad</b>	<b>nilík</b> líruk <b>jáwik</b> <b>wanúk</b> <b>wánug</b> nurág <b>rujúg</b> <b>línig</b>

### Experiment 2 Initial Contrast Condition Stimuli

Training	#P	#T	#K	P#	T#	K#	Fillers
	pínir pilán pulír pírul bínar bilún bawán búnill	tínir tajál tulír túwan dínar dirúl dawán dánill	kínir kilún kulír kíral gínar gurúl gawán gírul		jáwit wirút lurít lánat nílad julád rujúd nírada  wíjat nilút wajúta rúwata wínud nalíd jarúda lúnida		níjal rílun jáwar rújun líwir náwal júnar wírín lánir járan lirín wijúl janír ralún luwín nirál rijún wanúl liwár warín
Test	#P	#T	#K	P#	T#	K#	
	panúl páwir búlar burún	tanúl túrul dúlar dajún	kanúl kájar gúlar galín	nilíp nálup wánub jávíp rujáp waríb rujúb jávab	nilít níjat wánud luwíd	nilík líruk wánug jávík wanúk nurág rujúg línig	



## Experiment 2 Final Contrast Condition Stimuli

Training	#P	#T	#K	P#	T#	K#	Fillers
		<b>tínir</b> tajál <b>tulír</b> túwan <b>dínar</b> dirúl <b>dawán</b> dánil  tíjar tinúl talín tújal dánir dalúr duján díwil		<b>jáwip</b> rujáp <b>luríp</b> júnap <b>nílab</b> lijúb <b>rujúb</b> jáwab	<b>jáwit</b> wirút <b>lurít</b> lánat <b>nílad</b> julád <b>rujúd</b> nírad	<b>jáwik</b> wanúk <b>lurík</b> nálak <b>nílag</b> jalág <b>rujúg</b> línig	níjal rílun jáwar rújun líwir náwal júnar wírin lánir járan lirín wijúl janír ralún luwín nirál rijún wanúl liwár warín
Test	#P	#T	#K	P#	T#	K#	
	<b>panúl</b> pilán <b>búlar</b> páwir <b>bawán</b> búnil burún <b>pínir</b>	<b>tanúl</b> túrul <b>dúlar</b> dajún	<b>kanúl</b> kilún <b>gúlar</b> kájar <b>gawán</b> gírul galín <b>kínir</b>	<b>nilíp</b> nálup <b>wánub</b> waríb	<b>nilít</b> níjat <b>wánud</b> luwíd	<b>nilík</b> líruk <b>wánug</b> nurág	

## Ethics and consent

This research was certified exempt from IRB review by the UCLA Institutional Review Board (IRB#16-001655).

## Acknowledgements

I would like to thank two anonymous reviewers for their thoughtful comments and for helpful suggestions that strengthened the paper. Thanks also to associate editor Juliet Stanton; to Robert Daland, Bruce Hayes, Megha Sundara, and Kie Zuraw; to Charlie O'Hara; and to audiences at the 2017 Annual Meeting on Phonology (AMP 5), the 2018 LSA Annual Meeting, and the UCLA Phonology Seminar. All remaining errors are my own.

## Competing interests

The author has no competing interests to declare.

---

## References

- Archangeli, Diana & Pulleyblank, Douglas. 1994. *Grounded phonology*. Cambridge, MA: MIT Press.
- Bates, Douglas & Mächler, Martin & Bolker, Benjamin M. & Walker, Steven C. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1). 1–48. DOI: <https://doi.org/10.18637/jss.v067.i01>
- Becker, Michael & Levine, Jonathan. 2013. Experigen – an online experiment platform. <http://becker.phonologist.org/experigen>.
- Beguš, Gašper. 2019. Post-nasal devoicing and the blurring process. *Journal of Linguistics* 55(4). 689–753. DOI: <https://doi.org/10.1017/S002222671800049X>
- Blevins, Juliette. 2004. *Evolutionary phonology: The emergence of sound patterns*. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511486357>
- Cole, Ronald A. & Jakimik, Jola & Cooper, William E. 1978. Perceptibility of phonetic features in fluent speech. *The Journal of the Acoustical Society of America* 64(1). 44–56. DOI: <https://doi.org/10.1121/1.381955>
- de Lacy, Paul. 2006. *Markedness: Reduction and preservation in phonology*. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511486388>
- Do, Youngah & Zsiga, Elizabeth & Havenhill, Jonathan. 2016. Naturalness and frequency in implicit phonological learning. Talk presented at the 90<sup>th</sup> Annual Meeting of the Linguistic Society of America, Washington, D.C.
- Finley, Sara. 2012. Typological asymmetries in round vowel harmony: Support from artificial grammar learning. *Language and Cognitive Processes* 27(10). 1550–1562. DOI: <https://doi.org/10.1080/01690965.2012.660168>

- Fujimura, O. & Macchi, M. J. & Streeter, L. A. 1978. Perception of stop consonants with conflicting transitional cues: A cross-linguistic study. *Language and Speech* 21(4). 337–346. DOI: <https://doi.org/10.1177/002383097802100408>
- Glewwe, Eleanor. 2021. Substantive bias and the positional extension of a stop voicing contrast. Ms. Grinnell College.
- Glewwe, Eleanor & Zymet, Jesse & Adams, Jacob & Jacobson, Rachel & Yates, Anthony & Zeng, Ann & Daland, Robert. 2018. Substantive bias and the acquisition of final (de)voicing patterns. Talk presented at the 92<sup>nd</sup> Annual Meeting of the Linguistic Society of America, Salt Lake City.
- Greenwood, Anna. 2016. *An experimental investigation of phonetic naturalness*. Santa Cruz, CA: University of California, Santa Cruz dissertation.
- Hale, Mark & Reiss, Charles. 2000. “Substance abuse” and “dysfunctionalism”: Current trends in phonology. *Linguistic Inquiry* 31(1). 157–169. DOI: <https://doi.org/10.1162/002438900554334>
- Hayes, Bruce & Steriade, Donca. 2004. Introduction: the phonetic bases of phonological Markedness. In Hayes, Bruce & Kirchner, Robert & Steriade, Donca (eds.), *Phonetically based phonology*, 1–33. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511486401.001>
- Hayes, Bruce & Wilson, Colin. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39(3). 379–440. DOI: <https://doi.org/10.1162/ling.2008.39.3.379>
- Jun, Jongho. 1995. *Perceptual and articulatory factors in place assimilation: An optimality theoretic approach*. Los Angeles, CA: University of California, Los Angeles dissertation.
- Kean, Mary-Louise. 1975. *The theory of markedness in generative grammar*. Cambridge, MA: Massachusetts Institute of Technology dissertation.
- Lee, Wai-Sum & Zee, Eric. 2003. Standard Chinese (Beijing). *Journal of the International Phonetic Association* 33(1). 109–112. DOI: <https://doi.org/10.1017/S0025100303001208>
- Linzen, Tal & Gallagher, Gillian. 2017. Rapid generalization in phonotactic learning. *Laboratory Phonology: Journal of the Association for Laboratory Phonology* 8(1). 24. DOI: <https://doi.org/10.5334/labphon.44>
- Lombardi, Linda. 2001. Why place and voice are different: Constraint-specific alternations in Optimality Theory. In Lombardi, Linda (ed.), *Segmental Phonology in Optimality Theory: Constraints and Representations*, 13–45. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511570582.002>
- Lysvik, Julian K. 2018. An Artificial Language Learning experiment finds no bias against word-final voicing. Poster presented at the *Twenty-Sixth Manchester Phonology Meeting*, Manchester.
- Martin, Alexander & Peperkamp, Sharon. 2017. Assessing the distinctiveness of phonological features in word recognition: Prelexical and lexical influences. *Journal of Phonetics* 62. 1–11. DOI: <https://doi.org/10.1016/j.wocn.2017.01.007>
- Miller, George A. & Nicely, Patricia E. 1955. An analysis of perceptual confusions among some English consonants. *Journal of the Acoustical Society of America* 27(2). 338–352. DOI: <https://doi.org/10.1121/1.1907526>

- Mithun, Marianne & Basri, Hasan. 1986. The phonology of Selayarese. *Oceanic Linguistics* 25(1/2). 210–254. DOI: <https://doi.org/10.2307/3623212>
- Moreton, Elliott. 2008. Analytic bias and phonological typology. *Phonology* 25(1). 83–127. DOI: <https://doi.org/10.1017/S0952675708001413>
- Moreton, Elliott. 2010. Underphonologization and modularity bias. In Parker, Steve (ed.), *Phonological Argumentation: Essays on Evidence and Motivation* (Advances in Optimality Theory), 79–101. London: Equinox.
- Moreton, Elliott. 2012. Inter- and intra-dimensional dependencies in implicit phonotactic learning. *Journal of Memory and Language* 67(1). 165–183. DOI: <https://doi.org/10.1016/j.jml.2011.12.003>
- Moreton, Elliott & Pater, Joe. 2012a. Structure and substance in artificial-phonology learning, part I: Structure. *Language and Linguistics Compass* 6(11). 686–701. DOI: <https://doi.org/10.1002/lnc3.363>
- Moreton, Elliott & Pater, Joe. 2012b. Structure and substance in artificial-phonology learning, part II: Substance. *Language and Linguistics Compass* 6(11). 702–718. DOI: <https://doi.org/10.1002/lnc3.366>
- Moreton, Elliott & Pertsova, Katya. 2016. Implicit and explicit processes in phonotactic learning. In Scott, Jennifer & Waughtal, Deb (eds.), *Proceedings of the 40th Boston University Conference on Language Acquisition (BUCLD 40)*, 277–290. Somerville, MA: Cascadilla Press.
- Myers, Scott & Padgett, Jaye. 2014. Domain generalisation in artificial language learning. *Phonology* 31(3). 399–433. DOI: <https://doi.org/10.1017/S0952675714000207>
- Ohala, John J. 1974. Phonetic explanation in phonology. In Bruck, Anthony & Fox, Robert A. & La Galy, Michael W. (eds.), *Papers from the Parasession on Natural Phonology*, 251–274. Chicago: Chicago Linguistic Society.
- Ohala, John J. 1992. What's cognitive, what's not, in sound change. In Kellermann, Günter & Morrissey, Michael D. (eds.), *Diachrony within synchrony: Language history and cognition. [Duisburger Arbeiten zur Sprach- und Kulturwissenschaft 14]*, 309–355. Frankfurt am Main: Peter Lang Verlag.
- O'Hara, Charles Patrick. 2021. *Soft biases in phonology: Learnability meets grammar*. Los Angeles, CA: University of Southern California dissertation.
- O'Hara, Charlie. n.d. Word-edge consonant database. Available at <https://github.com/charlieohara/wordedgeconsonantdatabase>. (Accessed 2021-09-09.)
- Paradis, Carole & Prunet, Jean-François (eds.). 1991. *Phonetics and Phonology 2: The special status of coronals: internal and external evidence*. San Diego: Academic Press.
- Rice, Keren. 2007. Markedness in phonology. In de Lacy, Paul (ed.), *The Cambridge Handbook of Phonology*, 79–97. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/CBO9780511486371.005>
- Skoruppa, Katrin & Peperkamp, Sharon. 2011. Adaptation to novel accents: Feature-based learning of context-sensitive phonological regularities. *Cognitive Science* 35. 348–366. DOI: <https://doi.org/10.1111/j.1551-6709.2010.01152.x>

- Stampe, David. 1979. *A Dissertation on Natural Phonology*. New York: Garland Publishing, Inc.
- Steriade, Donca. 1994. Positional neutralization and the expression of contrast. Ms. University of California, Los Angeles.
- Steriade, Donca. 2001a. Directional asymmetries in place assimilation: A perceptual account. In Hume, Elizabeth & Johnson, Keith (eds.), *The Role of Speech Perception in Phonology*, 219–250. San Diego: Academic Press.
- Steriade, Donca. 2001b. The phonology of perceptibility effects: The P-map and its consequences for constraint organization. Ms. University of California, Los Angeles.
- Wang, Marilyn D. & Bilger, Robert C. 1973. Consonant confusions in noise: A study of perceptual features. *Journal of the Acoustical Society of America* 54(5). 1248–1266. DOI: <https://doi.org/10.1121/1.1914417>
- White, James. 2017. Accounting for the learnability of saltation in phonological theory: A maximum entropy model with a P-map bias. *Language* 93(1). 1–36. DOI: <https://doi.org/10.1353/lan.2017.0001>
- White, James Clifford. 2013. *Bias in phonological learning: Evidence from saltation*. Los Angeles, CA: University of California, Los Angeles dissertation.
- Wilson, Colin. 2006. Learning phonology with a substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* 30. 945–982. DOI: [https://doi.org/10.1207/s15516709cog0000\\_89](https://doi.org/10.1207/s15516709cog0000_89)

