Substantive Bias in Artificial Phonology Learning

Shuang Zheng[1], Youngah Do[1]

[1]Department of Linguistics, The University of Hong Kong

Abstract

This paper reviews the research on substantive bias in phonological learning since the publication of Moreton and Pater's paper on structure and substance in artificial phonology learning. The studies are categorized into vowel patterns, consonant patterns, and suprasegmental patterns. The review highlights advancements in experimental paradigms, a more precise definition of phonetic naturalness, and an expanded exploration of various phonological phenomena in the investigation of substantive bias over the past decade. Based on the review findings, we suggest that future research on the impact of substantive bias on phonological learning should include an examination of the articulatory and perceptual foundations of each phonological pattern, along with an analysis of the similarities in features, articulation, and perception.

Keywords: phonology; phonetics; learning bias; substantive bias; artificial language learning paradigms

Substantive Bias in Artificial Phonology Learning

1. Introduction

In Moreton & Pater (2012a, 2012b), the first comprehensive review of structure and substance in artificial phonology learning, it was argued that the influence of phonetic substance is at best weaker than structural bias. A major direction of subsequent works has been to unravel the nature of substantive bias and reveal its exact effect on phonological learning. The present paper reviews the research on substantive bias subsequent to the publication of Moreton & Pater (2012b) and compares it with the findings included in Moreton & Pater (2012b).

The structure of the current review mirrors that of Moreton & Pater (2012b) by dividing available studies into three categories: vowels, consonants, and suprasegments. There is a separate section on the asymmetries between consonants and vowels in Moreton & Pater (2012b) but not in the current review. Only one study since then demonstrated a learning advantage of consonant repetition over vowel repetition in word recall tests (Basnak & Ota, 2024), which was inconsistent with the earlier findings in the series of experiments by Juan Toro (Nevins & Toro, 2007; Toro et al., 2008a, 2008b; Pons & Toro, 2010). In addition, some studies involving more complicated natural language phenomena, such as opaque interactions in Prickett (2019) and Japanese Rendaku in Tanaka (2024), are not included in the current review either. The phenomena they examine often closely interact with other structural properties of language, such as morphology and syntax, and the substantive bias effect there may be confounded by structural bias.

Substantive bias studies on vowel patterns have exclusively focused on vowel harmony. In the section on vowel patterns (Section 2), we show that experimental paradigms designed to

investigate substantive bias have consistently developed. Based on this, we discuss the sensitivity of substantive bias to methodological choices of experiments. In contrast to the studies on vowel patterns, studies on consonant patterns experimented with a wider variety of phonological phenomena, such as nasalization, voicing, and saltatory alternations. In the section on consonant patterns (Section 3), we argue that the explorations of various consonant patterns have allowed us to consider the relevance of phonetic precursor strength to the substantive bias effect. The section on suprasegmental patterns (Section 4) is partitioned into tone patterns and stress patterns, both of which demonstrated consistent supporting evidence of substantive bias, contrary to the results on segmental patterns. We conclude by proposing future directions for substantive bias research in phonological learning.

2. Vowel patterns

One prominent typological asymmetry used to investigate substantive bias concerns the comparison of vowel harmony and vowel disharmony. This is presumably because these two patterns share similar levels of complexity, while vowel harmony is better supported by phonetic substance: The assimilation of vowel features reduces the number of gestures needed in production and, at the same time, provides extra cues to the quality of vowels in perception (Kimper, 2017; Ohala, 1994). Even though the two studies included in Moreton & Pater (2012b), i.e., Pycha et al. (2003) and Skoruppa & Peperkamp (2011), failed to show the effect of phonetic naturalness in learning vowel harmony, both studies envisioned a positive result with potential experimental design upgrades. Pycha and colleagues proposed that a statistically significant difference between groups could be found with a larger subject pool. Skoruppa and Peperkamp suggested a production task as a follow-up study on the grounds that phonological learning tends to be modality-specific (Warker et al., 2009). To draw a more comprehensive

conclusion about the status of substantive bias in the acquisition of vowel harmony versus vowel disharmony, subsequent studies introduced more nuanced tests, more naturalistic training, and iterative and communicative learning, to their investigations.

A series of experiments by Alexander Martin employed more refined tests. In Martin & Peperkamp (2020)[1], each participant received two tests, once immediately after exposure to an artificial language and once approximately twelve hours afterward either with or without sleep. Participants' performance in the test phase immediately after exposure showed a learning advantage of the harmony group on both familiar and novel items. In addition, the asymmetry in learning remained stable in the retest session either with or without sleep, highlighting the impact of substantive bias on both short-term and long-term memory. Martin and White (2021) required learners to make more complicated phonological generalizations. Participants were presented with single-suffix words in the training phase but tested on both single-suffix and double-suffix words. No evidence of different learning between harmony and disharmony was found in the single-suffix trials. Nonetheless, participants in the disharmony group gave significantly fewer correct responses (in the form of [+F][-F][+F][-F]) in the double-suffix trials compared to those in the harmony group, as they tended to prefer patterns where the two suffixes agreed in backness (in the form of [+F][-F][+F][+F]). This implies that learners inferred a harmonic generalization in the absence of clear evidence. Both studies replicated Pycha et al. (2003) with a larger subject pool (173 in Martin & Peperkamp (2020) and 120 in Martin & White (2021)), but only Martin and Peperkamp (2020) obtained positive results in the testing session directly comparable to Pycha et al. (2003). One possible explanation for the diverging results is that, while Pycha and colleagues (2003) and Martin and White (2021) used native English stimuli, Martin and Peperkamp (2020) exposed English-speaking participants to French stimuli, which inadvertently encouraged pure phonetic listening and discouraged

[1] First reported as Martin (2017), Chapter 3.4.

explicit sound-to-orthography mapping (Steele et al., 2015). This signifies how specific experimental design choices can overshadow the substantive bias effect.

Other studies on vowel harmony exhibit a growing endeavor to train subjects with more naturalistic input, despite the inherent limitations of the artificial language learning paradigm. Baer-Henney and colleagues (2015) introduced exceptions to the input, i.e., a minority proportion of training items (e.g., [mɪt ~ mɪtu] with vowel backness disharmony) did not conform to the target alternation pattern (e.g., [mɪt ~ mɪty] with vowel backness harmony), mirroring the variability found in natural languages. Their results demonstrate that participants relied more on phonetic substance when the exposure was shorter and the input was more variable, introducing high uncertainty to learning. Building upon this, later studies utilized variable input to heighten uncertainty in the learning process, demonstrating the impact of substantive bias with French-speaking adults (Martin, 2017), English-speaking adults (Mooney & Do, 2018), and Cantonese-speaking children (Do & Mooney, 2022).

Martin (2017) incorporated production tests as proposed by Skoruppa and Peperkamp (2011) and variable input following Baer-Henney and colleagues (2015). The author conducted a production experiment, a perception experiment, and a replication for each experiment while altering the input variability across participant groups (all harmonic, mostly harmonic, all disharmonic, mostly disharmonic, and control). A significant difference was found between the two variable input groups in the original production experiment but not in its replication. In addition, the performances of the two categorical input groups appeared significantly different in the first perception experiment, but this asymmetry was not observed in the follow-up experiment. The incompatible results again support the idea that the substantive bias effect is relatively weak and easily affected by random noise which is why substantive bias tests require large sample sizes to be well-powered. Importantly, the results allude to how the testing modality affects learning: The overall accuracy was lower in the perception task than in the

production task. Martin (2017) hypothesized that mixing correct and incorrect options in a forced-choice perception task might have confused participants. This was particularly true for the two variable groups lacking a confident generalization from training, which explained why substantive bias only manifested in the two categorical groups. On the other hand, such confusion was weaker in a production task, so the learning outcomes of the two variable input groups exhibited a substantive bias effect. In Mooney and Do (2018), adults maintained the probabilistic variation in their production relative to their input, regardless of whether the exposure was harmony-dominant or disharmony-dominant. This null substantive bias result was consistent with the production replication in Martin (2017), which had a similar variable input and the same number of participants (32 in two conditions). In Do & Mooney (2022), child participants in the harmony-dominant condition also reflected the input variability in the forced-choice test. Nevertheless, children in the disharmony-dominant condition completely reversed the distribution to a harmony-dominant system in their selections. This is in line with the morphosyntax studies reporting that children and adults have disparate strategies in dealing with variable input (Hudson Kam & Newport 2005, 2009), although their regularization behavior may not be the same across phonology and syntax.

Finley (2021) and Huang and Do (2023) explored learners' preference for vowel harmony by creating languages with variable alternations consisting of both alternating and non-alternating suffixes. In Finley (2021, Experiment 1), the stem + alternating suffix combinations were always harmonic (e.g., [keti-me], [kuto-mo]) while the stem + non-alternating suffix combinations complied with vowel harmony half of the time (e.g., [keti-go], [kuto-go]). Participants were more likely to select the harmonic words not only for the alternating suffix but also for the non-alternating suffix, indicating their inclination towards phonetically grounded patterns. Huang and Do (2023) assigned participants to one of four experimental conditions: categorically harmonic, categorically disharmonic, variably harmonic

where the non-alternating disharmonic suffix appeared 25% of the time, and variably disharmonic where the non-alternating harmonic suffix appeared 25% of the time. While there was no effect of condition in categorical learning, the variably harmonic condition achieved higher accuracy rates than the variably disharmonic condition. Moreover, the non-alternating suffix in the variably harmonic condition was modulated toward harmony, as in Finley (2021). In both studies, uncertainty in input might have caused more learning challenges, forcing learners to rely more on phonetic substance, and ultimately, becoming an overarching trigger for substantive bias in a lab setting.

Finally, Yu & Do (2022)[2] and Yu (2024) investigated vowel (dis)harmony in iterated learning and communicative learning, respectively, as an attempt to simulate the emergence of typological asymmetries in diachronic changes. Yu and Do (2022) constructed 8 generations of 8 transmission chains for vowel harmony and disharmony respectively. The learning outcomes of the previous generation were used as training materials for the next generation. There were only qualitative misalignments between conditions, such as one vowel harmony chain that remained stable at 100% accuracy rates since the fourth generation with no such tendency observed from any of the vowel disharmony chains. No statistically significant difference was found. In Yu (2024, Experiment 2), participants interacted in pairs through an online matching game after completing identical individual training and testing sessions. There was a marginal difference between conditions in the individual testing session. However, the disharmony condition showed a significant decrease in disharmony responses before and after interaction. No such change was found in the harmony condition, leading to a significant difference in selecting the target pattern between conditions in the interaction session. The author speculated that participants might have felt more burdened by communicative efficiency compared to

[2] Also reported as Experiment 1 in Yu (2024).

passively passing down the information in Yu and Do (2022). The burden of communicative efficiency resulted in high cognitive demands that amplified the effect of substantive bias.

Our review shows various reports that support the substantive bias hypothesis regarding vowel harmony versus disharmony since the review by Moreton and Pater (2012b). The analysis of the link between the experimental design and the impact of substantive bias increases our confidence that the substantive effect can be observed under the right experimental conditions, integrating more sophisticated assessments and creating more naturalistic training protocols. We also acknowledged that research has been conducted on iterative and communicative learning in laboratory settings. However, the precise impact of transmission and communication on learning with substantive bias is still unclear and requires further investigation.

Fruitful results have also been reported on more specific feature-based typological asymmetries related to vowel harmony. Finley (2012) examined the preference for a non-high trigger vowel in rounding harmony (Kaun, 2004). The round feature is perceptually more extreme for high vowels than mid vowels, so mid vowels receive more perceptual benefit from rounding and are more likely to trigger rounding harmony. While participants in the mid vowel condition (e.g., [bede-mi], [gobo-go]) successfully acquired the exposure pattern, the performance of the high vowel condition (e.g., [bidi-ge], [gubu-mu]) was at a chance level, signifying a failure in learning. Kimper (2015) followed up on Finley (2012) and tested participants' generalization of rounding harmony from the exposure condition to the untrained condition. High-performing learners exhibited successful generalization from the high vowel condition to the mid vowel condition. The generalization in the opposite direction was not successful. Both the learning and generalization results confirmed the presence of substantive bias in the asymmetric distribution of rounding harmony based on vowel trigger height. Aside from that, Kaun (2004) described a preference for a high target vowel in rounding harmony,

which was alluded to in Finley & Badecker (2009, Experiment 3). However, the authors found that the training on high vowel suffixes (e.g., [bide-mi], [podo-mu]) or low vowel suffixes (e.g., [bide-ge], [podo-go]) did not have an effect on participants' ability to identify harmonic words with the exposure suffix (e.g., [-mi], [-mu]) or generalize the harmonic pattern to the reserved suffix (e.g., [-ge], [-go]).

Mooney and Do (2018) and Do and Mooney (2022) explored the preference for (1) a non-high trigger vowel and (2) a high target vowel in rounding harmony with a different experimental setting from Finley's works. While Finley and colleagues focused on recreating the typological asymmetry in a laboratory setting, which required categorical input and semantically null stimuli, Do and Mooney were primarily interested in how free variation was acquired in rounding (dis)harmony patterns. Accordingly, they trained participants with variable input and semantically meaningful stimuli. Their findings contrasted with those of Finley & Badecker (2009) and Finley (2012), presumably due to the different experimental designs reflecting different research questions. Among both English-speaking adults and Cantonese-speaking children, the rate of rounding harmony application in high target vowels was significantly higher than that in mid target vowels. However, no effect of trigger height was found. It is possible that the biases for non-high trigger vowels and high target vowels in rounding harmony tend to surface in different elicitation environments.

Finley & Badecker (2012, Experiment 1) questioned whether the preference for front vowel targets over back vowel targets in height harmony (Linebaugh, 2007) is rooted in substantive bias. The height harmony that lowers [u] to [o] decreases the perceived roundedness of the target vowel and is thus dispreferred. Alternatively, the height harmony involving [i] and [e], which share the same constriction location, is articulatorily easier than that involving [u] and [o] with different constriction locations. Consistent with the phonetic grounding, participants in the front vowel condition (e.g., [dunig-i], [bemeg-e]) were unwilling to extend

height harmony to back vowels. In contrast, participants in the back vowel condition (e.g., [dunig-u̲], [bemeg-o̲]) selected significantly more harmonic responses with front vowels than the control condition but no more harmonic responses with back vowels than the control condition.

The above-mentioned studies tried to categorize vowel harmony into more specific types, i.e., rounding harmony or height harmony, with regard to relevant phonological features. All except one study yielded positive results. Moreover, Mooney and Do (2018), who found equivalent acquisition of vowel harmony and disharmony, were able to discover a difference in the learning of rounding harmony by vowel target height. These results suggest that exploring particular patterns and features related to vowel harmony can enhance our comprehension of the phonetic properties responsible for each typological asymmetry. Since vowel harmony involves all major phonological features, including round, back, high, and tense (van der Hulst & van de Weijer, 1995), we believe that vowel harmony continues to offer an ideal pattern for uncovering nuanced variations of the substantive bias effect related to specific phonological features.

3. Consonant patterns

Research in the domain of vowels almost exclusively dealt with vowel harmony alternations. Research in the domain of consonants, on the other hand, included both alternation and phonotactic studies of diverse phenomena. Within these, studies concerning nasal quality and voicing quality appeared to dominate in number.

On nasal quality, Lin (2016)[3] probed into a universal implicational relation which states that if obstruents are targets of nasalization in a language, so are the sonorants, but not vice versa (Walker, 2011). Participants were introduced to positive evidence of singular and plural pairs (e.g., [asaka ~ ãsãka]) where nasalization targeted either [s] or [w] but was always blocked by [k]. The [s] target group showed a high nasalization rate on [w], whereas the [w] target group performed poorly on the generalization to [s], in line with the typology. Strutjen and colleagues (2018) asked whether the substantive bias underlying vowel nasalization is based on production or perception, considering that the low nasal vowel [ã] is articulatorily easier but perceptually less salient than the non-low nasal vowels [ɛ̃] and [ĩ]. Participants heard the nasalization of one of three vowels ([a], [ɛ], or [i]) during exposure and selected between the oral form and the nasalized form of all three vowels in the test. The low vowel group had significantly fewer correct responses regarding all three vowels than the other two groups, signifying that vowel nasalization learning was influenced by a perceptual-based substantive bias[4]. Huang and Do (2021b) compared the [n] to [l] alternation triggered by non-high-front vowels as the pattern motivated by both perception and articulation and the alternation triggered by high-front vowels as the unmotivated pattern. The alternation rates between the two categorical input conditions were not significantly different. However, the phonetically motivated condition with variable input was significantly more likely to accept the alternation than the unmotivated condition with variable input. The authors proposed that substantive bias is triggered by uncertainty in learning, such as input variability.

As for voicing quality, van de Vijver and Baer-Henney (2014) investigated the production of articulatorily motivated intervocalic voicing by three different age groups (5-year-olds, 7-year-olds, and adults). There was an overgeneralization of intervocalic voicing in

[3] Experiment 1 largely failed due to an overcomplicated design, which was simplified in Experiment 2. Only results from Experiment 2 are reported here.
[4] Considering that Strutjen et al. (2018) employed a perception test, it was also possible that the articulatory factors in vowel nasalization were not activated in the test. We advise a cautious interpretation of their results.

children's production of nonce nouns while the proportion of overgeneralization decreased with age. The authors contended that children rely more on phonetic grounding than adults when making generalizations to nonce words because their language input is relatively small and unreliable. Do and colleagues (2016) and Do and Havenhill (2020) experimented with articulatorily natural postnasal voicing as opposed to postnasal devoicing (Hayes & Stivers, 1996). The acquisition of the two patterns was not significantly different when participants were exposed to categorical input and tested with forced-choice tasks in Do et al. (2016). Nevertheless, production in training significantly facilitated the choice of postnasal voicing as the majority input pattern in Do & Havenhill (2020). The implication of these findings is two-fold: First, substantive bias can be triggered by input uncertainty; second, substantive bias based on articulatory factors can be triggered by production during acquisition. Both of these implications substantiate the idea that substantive bias is sensitive to the methodological choices made in experimentation.

The voicing phenomenon that received the most attention among the studies on substantive bias was the crosslinguistic avoidance of word-final voiced obstruents, as voiced obstruents require more articulatory effort than voiceless ones and the final position further reduces their perceptual distinctness (Lysvik, 2020). The markedness of word-final voicing has been examined in both phonotactic and alternation patterns. Myers and Padgett (2014, Experiment 1) first recreated the phonotactics asymmetry by familiarizing participants with a language where only [s] or [z] was a possible utterance-final obstruent while neither [s] nor [z] was a possible utterance-medial word-final obstruent. The overall performance of the devoicing group was significantly better than that of the voicing group. Furthermore, participants in the devoicing group extended the generalization from utterance-final (e.g., [santa pi<u>s</u>]) to word-final positions (e.g., [santa pi<u>s</u> mizupu]), whereas the voicing group was unwilling to do so. The authors conjectured that final devoicing was more learnable than final

voicing due to substantive bias. Thereafter, Greenwood (2016, Experiment 2) exposed participants to isolated words with final (de)voicing. The two groups did not have significantly different performances when the input speech was casual and fast. Nevertheless, the voicing group gave significantly more correct judgments than the devoicing group when the input speech was hyperarticulated, contradicting the substantive bias hypothesis. On top of that, participants' performance in the test correlated with their perception in an identification phase after training and testing. The author thus posited that any performance difference between groups was not a result of substantive bias but a simple success or failure in perception[5]. This observation can point to the potential interplay between inductive bias and misperception-induced confusion, which should receive more attention in the interpretation of substantive bias results.

In addition to phonotactics, the learning of alternations involving voicing has been investigated and the results are mixed. Myers and Padgett (2014, Experiment 2) confirmed that learners could generalize utterance-final devoicing alternations to word-final positions. After exposure to singular ~ plural utterances where the singular nouns with voiceless stop endings always occurred at utterance-final positions (e.g., [bitomi gi<u>t</u> ~ bitomi gi<u>d</u>i] and [bitomi ki<u>p</u> ~ bitomi ki<u>p</u>i]), participants accepted significantly more alternations with utterance-medial word-final devoicing (e.g., [bitomi te<u>p</u> nama ~ bitomi te<u>b</u>i name]) compared to word-final voicing (e.g., [bitomi te<u>b</u> name ~ bitomi te<u>p</u>i name]). Glewwe and colleagues (2018)[6] followed this paradigm and changed the stimuli to singular ~ plural words where the final stop was either always voiceless (e.g., [tuli<u>p</u> ~ tuli<u>p</u>-i], [mule<u>p</u> ~ mule<u>b</u>-i]) or always voiced (e.g., [tuli<u>b</u> ~ tuli<u>p</u>-

[5] Greenwood (2016) hypothesized the English participants' misperception of the voiceless final patterns might have come from influences of the English /z/ exemplars or the English lexicon. However, the author acknowledged that neither hypothesis could explain the positive substantive bias results in Myers & Padgett (2014, Experiment 1). Glewwe (2019) and Lysvik (2020) argued that final devoicing is articulatorily motivated, and thus production tests are required for Greenwood (2016) to find genuine substantive bias or channel bias effects.

[6] Also reported as Experiment 5 in Glewwe (2019).

i], [muleb̲ ~ muleb̲-i]). Significantly better performance was found in the voicing condition than in the devoicing condition, again in the opposite direction of the substantive bias hypothesis. The negative results were attributed to a complexity bias at play: Owing to the presence of sonorant-final fillers, participants in the voicing condition might have learned a simple constraint *[-voice]#, whereas participants in the devoicing condition would have needed to learn a more complex composition of constraints *[-son, +voice]#[7]. Lysvik (2020) re-examined final (de)voicing in alternations without the potential confounds in previous works and discovered that participants in the devoicing condition were significantly more willing to produce and accept alternating plural forms (e.g., [rusub̲u] in [rusup̲ ~ rusub̲u]) than those in the voicing condition (e.g., [rusup̲u] in [rusub̲ ~ rusup̲u]). The author speculated that participants had a harder time acquiring word-final voicing than devoicing, but, instead of devoicing word-final obstruents in singular forms, they abandoned the alternation in plural forms altogether. Evjen (2021) implemented Lysvik's (2020) design in a transmission study. There was a trend for the devoicing condition to prefer alternation more than the voicing condition, resembling the findings of Lysvik (2020).

Some other studies compared nasalization with voicing in phonotactics and alternations. On the one hand, obstruents are disfavored compared to nasals in word-final phonotactics, because the stop burst was harder to perceive than nasality in coda positions. Greenwood (2016, Experiment 3) demonstrated that participants learning a nasal coda language (e.g., [pɑn̲.kɑ]) reached significantly higher accuracy in the judgment task than those learning an obstruent coda language (e.g., [pɑt̲.kɑ]) when the input speech was fast and casual. On the other hand, devoicing as opposed to nasalization is a favored repair strategy for final voiced obstruents in alternations, as devoicing is a perceptually minimal alternation compared to nasalization

[7] Myers & Padgett (2014) and Greenwood (2016) also used sonorants as fillers exclusively. Glewwe (2019) suspected that structural bias also contributed to the negative results in Greenwood (2016). There was no generalizable explanation as to why Myers & Padgett (2014) still acquired positive results in the presence of structural bias.

(Steriade, 2001). In Albright & Do (2017), participants trained on singular ~ plural pairs with an equal amount of final devoicing (e.g., [deɪp̲ ~ deɪb̲-i:]) and nasalization (e.g., [tɹi:m̲] ~ [tɹi:b̲-i:]) showed a significant preference for devoicing over nasalization in forced-choice tests, again supporting the substantive bias hypothesis.

On top of typological tendencies found at word-final positions, substantive bias research has delved into positional asymmetries of phonological patterns' distributions. For example, Glewwe (2018)[8] examined the substantive bias underlying an implicational relation that if a language has word-final obstruent voicing contrast, it will also contrast voicing word-initially, but not vice versa (Steriade, 1997). Results showed that participants exposed to word-final voicing contrasts (*#T or *#D) were more likely to infer a contrast in the reserved position (i.e. word-initial) than those exposed to word-initial voicing contrasts (*T# or *D#), confirming the substantive bias hypothesis. Glewwe (2022)[9] investigated a parallel implicational relation concerning the place of articulation: If a language has major place contrasts (labial, coronal, and dorsal) in stops in the coda position, it will most likely have major place contrasts in stops in the onset position, but not vice versa (Blevins, 2004). However, there was no emergence of substantive bias, i.e., participants from initial and final contrast conditions were equally capable of accepting conforming items and rejecting nonconforming ones, contrary to the findings of Glewwe (2018). The author conjectured that the perceptual difference of place contrasts across positions might be weaker than that of voicing contrasts, which led to word-final place contrasts being only marginally harder to perceive than word-initial place contrasts, leaving learners under less pressure to avoid word-final place contrasts. Finally, Finley (2017) asked whether learners' acquisition of phonological metathesis is governed by the syllable structure constraint

[8] Also reported as Experiment 3 and 4 in Glewwe (2019). Glewwe (2018, Experiment 1) used sonorants as fillers exclusively and introduced structural bias to the design, similar to Glewwe et al. (2018). The author changed all fillers to voiceless fricatives in Experiment 2. Only results from Glewwe (2018, Experiment 2) are reported here.

[9] Also reported as Experiment 1 and 2 in Glewwe (2019).

that onsets are preferred over codas. When learners were exposed to arbitrary metatheses (/r/ coda + C onset or /l/ coda + C onset), they generalized freely to phonologically grounded metatheses (C+/r/ or C+/l/ complex onsets). The generalization in the opposite direction was significantly weaker, in accordance with the substantive bias hypothesis.

The contrastive results of Glewwe (2018) on voicing and Glewwe (2022) on place of articulation suggest that different phonetic precursors have different strengths in the acquisition of phonological patterns. In other words, changes in a single feature, depending on which feature, may or may not be strong enough to cause learnability differences in the lab. Albright and Do (2017) advised that the varying strengths of different phonetic precursors could be better understood by considering perceptual similarity. In the following section, we present substantive bias studies that assessed phonological and perceptual similarity in addition to featural similarity and discuss how they serve the elicitation of substantive bias in the lab.

White (2014) provided experimental evidence for a substantive bias against saltatory alternations primarily driven by perceptual motivations. In Experiment 1, participants familiarized with potentially saltatory patterns (e.g., [p ~ v]) extended the alternation to include intermediate sounds (e.g., [b ~ v] and [f ~ v]) and effectively transformed the learning outcomes to non-saltatory patterns, while participants familiarized with non-saltatory patterns (e.g., [b ~ v]) rarely extended the alternation. Moreover, the acceptance rate for stop alternations (e.g., [b ~ v]) was significantly more than that for fricative alternations (e.g., [f ~ v]), which could not be explained by structural bias but could be interpreted as a bias favoring perceptually minimal alternations (Steriade, 2001). Furthermore, the explicitly saltatory condition (e.g., [p ~ v] and [b ~ b]) in Experiment 2 yielded a similar proportion of alternation as the potentially saltatory condition in Experiment 1, indicating that participants were biased against saltatory alternation even with unambiguous evidence in the input. White & Sundara (2014) replicated the results with 12-month-old infants. The authors found that infants' looking time for dissimilar and

similar sounds differed significantly after exposure to perceptually dissimilar sounds (e.g., [p ~ v]) but not after exposure to perceptually similar sounds (e.g., [b ~ v]). The infants' capability to generalize alternations from dissimilar to similar sounds but not the other way around is consistent with a substantive bias favoring alternations between perceptually similar sounds.

Stave and colleagues (2013) and Smolek and Kapatsinski (2018) [10] approached articulatorily motivated saltatory alternation with palatalization patterns. The authors reckoned that alveolar palatalization ([t ~ tʃ]) and velar palatalization ([k ~ tʃ]) involve articulatorily more minor changes than labial palatalization ([p ~ tʃ]) because [t], [k] and [tʃ] are all lingual gestures while [p] is a labial gesture (Yun, 2006). Production and judgment results both established that the labial condition was not more likely to palatalize [p] than [t] or [k], while the alveolar condition and the velar condition preferred the respective place of articulation they were trained on[11]. Furthermore, the bias against labial palatalization was stronger in production than in judgment, substantiating the articulatory basis of the bias against labial palatalization.

Finley (2022) questioned whether similarity-based generalizations of spirantization patterns are determined by articulatory or perceptual factors. Pairs of sounds that differ in voicing quality are articulatorily more similar than those differing in places of articulation under the calculation of shared natural classes in English (Frisch, 1996; Frisch et al., 2004) but perceptually less similar than those differing in places of articulation since they are less mutually confusable (Wang & Bilger, 1973). Finley's (2022) results indicated that the generalization of spirantization patterns based on voicing (e.g., from [p ~ f], [t ~ s] to [b ~ v], [d ~ z]) was significantly more successful than the generalization based on place of articulation (e.g., from [t ~ s], [d ~ z] to [p ~ f], [b ~ v]). The preference for articulatorily minimal alternations aligned with the articulatorily motivated phonological similarity bias.

[10] Stave et al. (2013) tested palatalization in an unnatural context (before [-a]) while Smolek & Kapatsinski (2018) in a natural context (before [-i]). But their results converged and thus are reported together here.
[11] Contrary to White (2014), Smolek and Kapatsinski (2018) did not find an overgeneralization of saltatory alternations to non-saltatory ones and accredited the discrepancy to the threshold training in White's study.

The specific notion of "similarity" in artificial phonology learning dates back to Skoruppa et al. (2011), which illustrated how featural similarity (i.e. single feature change, double feature changes, or multiple feature changes) relates to the learnability of alternations. In Cristia et al. (2013) and White (2014), learners' phonological generalizations align with both featural similarity and perceptual or articulatory similarity. White (2014), Stave and colleagues (2013), Smolek and Kapatsinski (2018), and Finley (2022) further distinguished how perceptual and articulatory similarity play a role in the generalization of alternations. The development from featural similarity to perceptual and articulatory similarity is illustrative of a refined understanding of phonetic naturalness and, accordingly, substantive bias. In research prior to Moreton & Pater (2012a, 2012b), we observe that phonological naturalness tends to conflate phonetic substance and featural complexity, e.g., an alternation pattern involving fewer featural changes is also more grounded on phonetic substance than the one involving more featural changes. Subsequent to the division of substantive and structural bias in Moreton & Pater (2012a, 2012b), scholars have been more mindful of controlling the featural complexity of patterns when investigating the role of phonetic substance in learning, which we believe contributed to the steadily accumulating positive substantive bias results. As far as we see, the current problem is that not all phonetic precursors are salient enough to have an impact on phonological learning. The studies on perceptual and articulatory similarity (Albright & Do, 2017; Finley, 2022; White, 2014) may guide future research on substantive bias, emphasizing a detailed examination of alternation patterns characterized by shared natural classes and mutual confusability of features. Along with this, there was a call for a separation between perceptually-grounded and articulatorily-grounded natural patterns due to the inconsistent outcomes within the two categories and the unique research modality required for the learning of each pattern (Do & Havenhill, 2020; Glewwe, 2019).

4. Suprasegmental patterns

The artificial language learning experiments on suprasegmental patterns, though smaller in number compared to the segmental studies, have consistently yielded positive results. Two categories of suprasegmental patterns have been the focus of these experiments: tone and stress.

Chen (2020) looked into the tonal phonotactic tendency to avoid rising tones at non-domain-final positions, which has its basis in articulation: The same magnitude of pitch raising is more difficult to achieve in a non-domain final syllable. It was found that learners of the language that bans non-final rising tones (*NonFinalR) exhibited significantly higher consistency in their grammatical judgments across two testing sessions than learners of the language that bans non-final high tones (*NonFinalH)[12]. Nevertheless, the results were ambiguous between structural bias and substantive bias, which was further addressed by Chen (2024). The follow-up study constructed languages banning non-final falling tones (*NonFinalF) and low tones (*NonFinalL), respectively, in addition to *NonFinalR and *NonFinalH. Only the *NonFinalR language learners demonstrated implicit knowledge of the target pattern by performing significantly better than the control group learners while being unconfident about their judgment. Since *NonFinalR and *NonFinalF must be structurally equivalent, the current review regards the preference for *NonFinalR as driven by substantive bias.

[12] It should be noted that the main effect of condition was significant in predicting response consistency but not significant in predicting response accuracy. In addition to the tonal constraint, all participants were simultaneously exposed to a natural segmental constraint that bans retroflex consonants in the languages. According to the author, participants' judgments on items conforming or nonconforming to *NonFinalR were consistent with that of the natural segmental constraint, while the same could not be said for *NonFinalH, which further supported the argument that *NonFinalR is a natural tonal constraint.

Kao (2017, Experiment 1)[13] presented evidence for substantive bias in the acquisition of a natural pattern where the high tone is preserved in tone elision against an unnatural pattern where the low tone is preserved. Among English participants, the natural condition showed significantly better performance on both familiar and novel items than the unnatural condition. Among Mandarin participants, those who learned the natural pattern were more likely to generalize to novel items than those who learned the unnatural pattern, although their performance on familiar items was equivalent. The Mandarin participants' performance further demonstrates how substantive bias can be better elicited in generalization to novel items compared to recognition of familiar items.

Huang & Do (2021a) uncovered a preference for right-dominant tone deletion as opposed to left-dominant tone deletion. The two rules are both uni-directional and structurally equivalent, but only right-dominant tone deletion is phonetically motivated: The duration of the left syllable is relatively insufficient, so the contour tone on the left syllable is more apt to change. Both Mandarin and Cantonese speakers learned the right-dominant pattern significantly better than the left-dominant pattern, agreeing with the substantive bias hypothesis.

The works on stress patterns also invariably demonstrated the role of phonetic substance in learning. Carpenter (2010) compared a natural rule that stresses the first low vowel, or else the first vowel, with an unnatural rule that stresses the first high vowel, or else the first vowel. The author manipulated all applicable correlates of stress – length, intensity, and pitch – on critical vowels to maximally neutralize the phonetic differences of items between conditions. With the perceptibility of the two rules controlled for, any difference in the acquisition of the two rules could be ascribed to substantive bias. Among both native English speakers and native French speakers, the stress low vowel group received significantly higher

[13] Kao (2017, Experiment 2) attempted to assess the substantive bias effect behind natural and unnatural contour tone formation strategies but failed to obtain interpretable results owing to the higher complexity of contour tone formation. The results relevant to Experiment 2 are not reported here.

scores on novel items than the stress high vowel group. The French speaker even exhibited a preference for the stress low vowel pattern in the familiar item tests. Carpenter (2016a) replicated the study with 9- and 10-year-old English-speaking children and found that the stress high vowel group was more inclined to choose the stress-low vowel pattern when presented with both in forced-choice tests. With a similar stimuli manipulation, Carpenter (2016b) further inquired whether the weight-to-stress principle that heavy syllables are more prone to attract stress than light syllables is motivated by substantive bias. Among both English and French participants, the stress heavy syllable group scored significantly higher than the stress light syllable group on novel items. The findings of all three experiments were well-aligned with the substantive bias hypothesis.

Notably, Greenwood (2016, Experiment 1) hypothesized that the weight-to-stress principle is a result of misperception, and biased learning can be elicited if and only if the specific perceptual advantage of the natural pattern over the unnatural pattern is recreated in the lab. To this end, the author preserved the phonetic differences of critical syllables – rhyme duration and pitch – in the stimuli manipulation. Participants attended a learning experiment based off of Carpenter (2010) and a perception experiment on the critical items. The proportion of correct trials in the stress light syllable condition was at chance level and significantly lower than that in the stress heavy syllable condition. Moreover, only the participants in the stress light syllable condition displayed a worse perception of words. On that account, the worse acquisition of a pattern could have been a result of the misperception of the pattern, instead of substantive bias. Seeing that potential coarticulation effects in the stimuli production could have led to an across-the-board misperception among the learners, Greenwood (2016) advocated for detailed descriptions of the stimuli recording in future works.

That being said, the substantive bias results on suprasegmental patterns were still consistently positive, contrary to the results on segmental patterns. However, the comparable

results of Experiment 1 on stress and Experiment 3 on coda sonority in Greenwood (2016) imply that the impact of phonetic substance on segmental and suprasegmental learning may not be fundamentally different. Huang & Do (2021a) also claimed that the nature of substantive bias in the learning of suprasegmental and segmental phonology seems to be similar. Our preliminary hypothesis for the discrepancy in substantive bias results across the two domains stems from the observation that the acquisition of suprasegmental patterns could be comprehensively harder. Artificial language patterns, including the phonetically unnatural ones, were designed to be easily learnable in the span of an experiment. It was first proposed by Finley (2012) and espoused again by Glewwe (2019) that participants learning the unnatural pattern might have already caught up with that of the natural pattern by the time they entered the testing session, leading to the null results in segmental studies. Contrastively, suprasegmental studies appear to have lengthier training sessions (Carpenter 2010, 2016a, 2016b; Greenwood, 2016, Experiment 1) compared to segmental studies and involved completely categorical inputs (all studies in Section 4), and yet certain patterns could not be successfully acquired by the end of training (Kao, 2017, Experiment 2). Therefore, the positive results in suprasegmental studies might have been elicited at the point when the acquisition of the natural and unnatural patterns had not converged. We have not identified any studies directly comparing the learnability of segmental and suprasegmental patterns through varying durations of training, which may provide clues to the distinct substantive bias effects across the two domains. Further research in this area is warranted.

5. Discussion

This review evaluates the research conducted since the publication of Moreton & Pater (2012b) on substantive bias in phonological learning. Based on the review, we argue that the

substantive bias effect emerges under appropriate experimental parameters and with a well-considered selection of phonological patterns that encompass diverse phonetic precursors. In this section, we propose two directions for future research on substantive bias.

Firstly, as our review has shown, researchers have implemented several improvements in laboratory settings, including nuanced tests, naturalistic training, and the incorporation of iterative and communicative learning. These advancements address broader concerns that artificial language learning experiments may not fully capture the complexities of the human learning process. Nevertheless, we have not fully understood the extent to which each methodological choice affected the results. For example, it has been shown that phonological learning is modality-specific (Warker et al., 2009), and production aids in the successful elicitation of substantive bias on some occasions (Do & Havenhill, 2020; Martin, 2017), while it is yet to test whether production and perception affect articulatorily-grounded and perceptually-grounded natural patterns in distinctive ways and how misproduction and misperception should be taken into account in the analysis.

Secondly, the revision of the definition of phonetic naturalness and its associated concept of substantive bias successfully separated it from structural bias. However, we still observe that different phonetic precursors influence the effects of substantive bias to varying degrees. We believe the future of substantive bias research lies in a more intricate division of phonological patterns. For instance, we need to evaluate not only featural similarity but also articulatory and perceptual similarity. We then need to probe into the phonetic bases of a typological asymmetry (articulatory or perceptual or both), though we acknowledge that these two aspects are frequently interconnected, making it challenging to differentiate between them. Finally, further examination is warranted to determine whether and how the strength of the phonetic precursor differs among segmental patterns and between segmental and

suprasegmental patterns. We believe further research in these directions will provide a better understanding of the nature and characteristics of substantive bias in phonological learning.

References

Albright, A., & Do, Y. (2017). A substantive bias for perceptually minimal alternations in Artificial Grammar learning. *Old World Conference on Phonology (OCP), 2017*. https://hub.hku.hk/bitstream/10722/274779/1/Abstract.pdf?accept=1

Baer-Henney, D., Kügler, F., & van de Vijver, R. (2015). The Interaction of Language-Specific and Universal Factors During the Acquisition of Morphophonemic Alternations With Exceptions. *Cognitive Science*, *39*(7), 1537–1569. https://doi.org/10.1111/cogs.12209

Basnak, J., & Ota, M. (2024). Learnability Advantage of Segmental Repetitions in Word Learning. *Language and Speech*, 00238309231223909. https://doi.org/10.1177/00238309231223909

Blevins, J. (2004). *Evolutionary phonology: The emergence of sound patterns*. Cambridge University Press. https://books.google.com/books?hl=en&lr=&id=-AsIRQT2VFQC&oi=fnd&pg=PP1&dq=blevins+2004&ots=OiRQpZ9W5a&sig=dLT6yBl1vdPlAl6VPz3t08EdNh4

Carpenter, A. C. (2010). A naturalness bias in learning stress. *Phonology*, *27*(3), 345–392. https://doi.org/10.1017/S0952675710000199

Carpenter, A. C. (2016a). Learning natural and unnatural phonological stress by 9-and 10-year-olds: A preliminary report. *Journal of Child Language Acquisition and Development*, *4*(2), 62–77.

Carpenter, A. C. (2016b). The Role of a Domain-specific Language Mechanism in Learning Natural and Unnatural Stress. *Open Linguistics*, *2*(1). https://doi.org/10.1515/opli-2016-0006

Chen, T.-Y. (2020). An inductive learning bias toward phonetically driven tonal phonotactics. *Language Acquisition*, *27*(3), 331–361. https://doi.org/10.1080/10489223.2020.1769630

Chen, T.-Y. (2024). Revisiting the phonological learnability of* NonFinalR: A large-scale experimental study. *Proceedings of the 19th Conference on Laboratory Phonology*. 19th Conference on Laboratory Phonology. https://labphon.org/sites/default/files/labphon19/Papers/LabPhon19_paper_5.pdf

Cristia, A., Mielke, J., Daland, R., & Peperkamp, S. (2013). Similarity in the generalization of implicitly learned sound patterns. *Laboratory Phonology*, *4*(2). https://doi.org/10.1515/lp-2013-0010

Do, Y., & Havenhill, J. (2020). Production and Substantive Bias in Phonological Learning. *Proceedings of the Annual Meetings on Phonology*. https://doi.org/10.3765/amp.v9i0.4925

Do, Y., & Mooney, S. (2022). Variation awaiting bias: Substantively biased learning of vowel harmony variation. *Journal of Child Language*, *49*(2), 397–407. https://doi.org/10.1017/S0305000920000719

Do, Y., Zsiga, E., & Havenhill, J. (2016). Naturalness and frequency in implicit phonological learning. *Talk Presented at the 90th Annual Meeting of the Linguistic Society of America*. 90th Annual Meeting of the Linguistic Society of America.

Evjen, M. (2021). *The transmission of unnatural phonology-Testing final voicing and final devoicing through iterated artificial language learning* [Master's Thesis, University of Oslo]. https://www.duo.uio.no/handle/10852/87285

Finley, S. (2012). Typological asymmetries in round vowel harmony: Support from artificial grammar learning. *Language and Cognitive Processes*, *27*(10), 1550–1562. https://doi.org/10.1080/01690965.2012.660168

Finley, S. (2017). Learning metathesis: Evidence for syllable structure constraints. *Journal of Memory and Language*, *92*, 142–157.

Finley, S. (2021). Learning Exceptions in Phonological Alternations. *Language and Speech*, *64*(4), 991–1017. https://doi.org/10.1177/0023830920978679

Finley, S. (2022). Generalization to Novel Consonants: Place Versus Voice. *Journal of Psycholinguistic Research*, *51*(6), 1283–1309. https://doi.org/10.1007/s10936-022-09897-1

Finley, S., & Badecker, W. (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language*, *61*(3), 423–437. https://doi.org/10.1016/j.jml.2009.05.002

Finley, S., & Badecker, W. (2012). Learning biases for vowel height harmony. *Journal of Cognitive Science*, *13*(3), 287–327.

Frisch, S. (1996). *Similarity and frequency in phonology* [Northwestern University]. https://search.proquest.com/openview/8d5c449d59f3476d8f22f1b9ad9934ee/1?pq-origsite=gscholar&cbl=18750&diss=y

Frisch, S. A., Pierrehumbert, J. B., & Broe, M. B. (2004). Similarity Avoidance and the OCP. *Natural Language & Linguistic Theory*, *22*(1), 179–228. https://doi.org/10.1023/B:NALA.0000005557.78535.3c

Glewwe, E. (2018). Complexity Bias and Substantive Bias in Phonotactic Learning. *Poster at Annual Meeting on Phonology*, 5–7. https://eglewwe.com/wp-content/uploads/2019/01/lsa-2019-handout.pdf

Glewwe, E. (2022). Substantive bias and the positional extension of major place contrasts. *Glossa: A Journal of General Linguistics*, *7*(1), Article 1. https://doi.org/10.16995/glossa.6537

Glewwe, E. R. (2019). *Bias in phonotactic learning: Experimental studies of phonotactic implicationals* [PhD Thesis, University of California, Los Angeles].

https://search.proquest.com/openview/f367523c801c1cb8affba876965d96d8/1?pq-origsite=gscholar&cbl=18750&diss=y

Glewwe, E., Zymet, J., Adams, J., Jacobson, R., Yates, A., Zeng, A., & Daland, R. (2018). Substantive bias and the acquisition of final (de)voicing patterns. *92nd Annual Meeting of the Linguistic Society of America*. 92nd Annual Meeting of the Linguistic Society of America.

Greenwood, A. (2016). *An experimental investigation of phonetic naturalness* [PhD Thesis, University of California, Santa Cruz]. https://search.proquest.com/openview/2e8f7163a82cedd8fe1598da98413a13/1?pq-origsite=gscholar&cbl=18750

Hayes, B. and Stivers, T., 1996. The phonetics of postnasal voicing. Working Papers, Department of Linguistics, University of California, Los Angeles, CA.

Huang, T., & Do, Y. (2021a). Phonetically Grounded Structural Bias in Learning Tonal Alternations. *Frontiers in Psychology*, *12*. https://www.frontiersin.org/articles/10.3389/fpsyg.2021.705766

Huang, T., & Do, Y. (2021b). Substantive Bias and Variation in the Acquisition of/n/∼/l/Alternation. *Proceedings of the Annual Meetings on Phonology*. http://journals.linguisticsociety.org/proceedings/index.php/amphonology/article/view/5161

Huang, T., & Do, Y. (2023). Substantive bias and variation in the acquisition of vowel harmony. *Glossa: A Journal of General Linguistics*, *8*(1), Article 1. https://doi.org/10.16995/glossa.9313

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing Unpredictable Variation: The Roles of Adult and Child Learners in Language Formation and Change. *Language*

Learning and Development, *1*(2), 151–195. https://doi.org/10.1080/15475441.2005.9684215

Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, *59*(1), 30–66. https://doi.org/10.1016/j.cogpsych.2009.01.001

Kao, S. (2017). *Phonological Learning Bias in Tone Patterns* [PhD Thesis, State University of New York at Stony Brook]. https://search.proquest.com/openview/decdc6bf4a3884d69df7e010d5718514/1?pq-origsite=gscholar&cbl=18750

Kaun, A. (2004). The typology of rounding harmony. In B. Hayes, R. Kirchner, & D. Steriade (Eds.), *Phonetically based phonology*. Cambridge University Press. https://books.google.com/books?hl=en&lr=&id=LtYrdhxAO_QC&oi=fnd&pg=PA87&dq=the+typology+of+rounding+harmony&ots=X4aJH33mwR&sig=RmNgPYh3IdjfmtcOjNjhzbfPmDY

Kimper, W. (2015). Asymmetric generalisation of harmony triggers. *Proceedings of the Annual Meetings on Phonology*. http://journals.linguisticsociety.org/proceedings/index.php/amphonology/article/view/3662

Kimper, W. (2017). Not crazy after all these years? Perceptual grounding for long-distance vowel harmony. *Laboratory Phonology*, *8*(1). https://www.journal-labphon.org/articles/10.5334/labphon.47/

Lin, Y.-L. (2016). *Sonority effects and learning bias in nasal harmony* [PhD Thesis, University of Toronto (Canada)]. https://search.proquest.com/openview/5cc777ac668d8298538527d4d6fafc59/1?pq-origsite=gscholar&cbl=18750

Linebaugh, G. D. (2007). *Phonetic grounding and phonology: Vowel backness harmony and vowel height harmony* [PhD Thesis, University of Illinois at Urbana-Champaign]. https://www.ideals.illinois.edu/items/83923

Lysvik, J. K. (n.d.). *Where does naturalness in phonology come from? - Insights from artificial language learning* [PhD Thesis]. University of Oslo.

Martin, A. (2017). *Biases in phonological processing and learning* [PhD Thesis, Université Paris sciences et lettres]. https://theses.hal.science/tel-01939096

Martin, A., & Peperkamp, S. (2020). Phonetically natural rules benefit from a learning bias: A re-examination of vowel harmony and disharmony. *Phonology*, *37*(1), 65–90. https://doi.org/10.1017/S0952675720000044

Martin, A., & White, J. (2021). Vowel Harmony and Disharmony Are Not Equivalent in Learning. *Linguistic Inquiry*, *52*(1), 227–239. https://doi.org/10.1162/ling_a_00375

Mooney, S., & Do, Y. (2018). Learners change artificial languages to constraint free variation in line with typological principles. *Old World Conference on Phonology (OCP)*. Old World Conference on Phonology (OCP).

Moreton, E., & Pater, J. (2012a). Structure and Substance in Artificial-phonology Learning, Part I: Structure. *Language and Linguistics Compass*, *6*(11), 686–701. https://doi.org/10.1002/lnc3.363

Moreton, E., & Pater, J. (2012b). Structure and Substance in Artificial-Phonology Learning, Part II: Substance. *Language and Linguistics Compass*, *6*(11), 702–718. https://doi.org/10.1002/lnc3.366

Myers, S., & Padgett, J. (2014). Domain generalisation in artificial language learning. *Phonology*, *31*(3), 399–433.

Nevins, A., & Toro, J. M. (2007, May 21). Consonants as skeleta, vowels as glue. *Paper Presented at the Ljubljana Linguistics Circle*. Ljubljana Linguistics Circle, University of Ljubljana.

Ohala, J. J. (1994). Towards a universal, phonetically-based, theory of vowel harmony. *ICSLP*, *3*, 491–494.

Pons, F., & Toro, J. M. (2010). Structural generalizations over consonants and vowels in 11-month-old infants. *Cognition*, *116*(3), 361–367.

Prickett, B. (2019). Learning biases in opaque interactions. *Phonology*, *36*(4), 627–653.

Pycha, A., Nowak, P., & Shin, E. (2003). Phonological Rule-Learning and Its Implications for a Theory of Vowel Harmony. *Proceedings of the 22nd West Coast Conference on Formal Linguistics*, *22*.

Skoruppa, K., Lambrechts, A., & Peperkamp, S. (2011). The role of phonetic distance in the acquisition of phonological alternations. *Proceedings of the 39th Meeting of the North-East Linguistics Society (NELS)*. The 39th Meeting of the North-East Linguistics Society (NELS)), Amherst, Massachusetts. https://repository.essex.ac.uk/4251/1/SkoruppaetalNELS.pdf

Skoruppa, K., & Peperkamp, S. (2011). Adaptation to Novel Accents: Feature-Based Learning of Context-Sensitive Phonological Regularities. *Cognitive Science*, *35*(2), 348–366. https://doi.org/10.1111/j.1551-6709.2010.01152.x

Smolek, A., & Kapatsinski, V. (2018). What happens to large changes? Saltation produces well-liked outputs that are hard to generate. *Laboratory Phonology*, *9*(1). https://www.journal-labphon.org/article/id/6227/

Stave, M., Smolek, A., & Kapatsinski, V. (2013). Inductive bias against stem changes as perseveration: Experimental evidence for an articulatory approach to output-output

faithfulness. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *35*(35). https://escholarship.org/content/qt293733rg/qt293733rg.pdf

Steele, A., Denby, T., Chan, C., & Goldrick, M. (2015). Learning non-native phonotactic constraints over the web. *ICPhS*. https://www.internationalphoneticassociation.org/icphs-proceedings/ICPhS2015/Papers/ICPHS0258.pdf

Steriade, D. (1997). *Phonetics in phonology: The case of laryngeal neutralization*. UCLA.

Steriade, D. (2001). Directional asymmetries in place assimilation: A perceptual account. *The Role of Speech Perception in Phonology*, 219–250.

Strütjen, K., Baer-Henney, D., Indefrey, P., & van de Vijver, R. (2018). *Perceptual bias in learning a vowel nasalization pattern*.

Tanaka, Y. (2024). Learning biases in proper nouns. *Phonology*, 1–32.

Toro, J. M., Nespor, M., Mehler, J., & Bonatti, L. L. (2008). Finding Words and Rules in a Speech Stream: Functional Differences Between Vowels and Consonants. *Psychological Science*, *19*(2), 137–144. https://doi.org/10.1111/j.1467-9280.2008.02059.x

Toro, J. M., Shukla, M., Nespor, M., & Endress, A. D. (2008). The quest for generalizations over consonants: Asymmetries between consonants and vowels are not the by-product of acoustic differences. *Perception & Psychophysics*, *70*(8), 1515–1525. https://doi.org/10.3758/PP.70.8.1515

van de Vijver, R., & Baer-Henney, D. (2014). Developing biases. *Frontiers in Psychology*, *5*, 83070.

van der Hulst, H., & van de Vijver, J. (1995). Vowel Harmony. In J. Goldsmith, J. Riggle, & A. C. L. Yu (Eds.), *The Handbook of Phonological Theory* (1st ed.). Wiley. https://doi.org/10.1002/9781444343069

Walker, R. (2011). Nasal Harmony. In M. van Oostendorp (Ed.), *The Blackwell companion to phonology*. Wiley-Blackwell.

Wang, M. D., & Bilger, R. C. (1973). Consonant confusions in noise: A study of perceptual features. *The Journal of the Acoustical Society of America*, *54*(5), 1248–1266.

Warker, J. A., Xu, Y., Dell, G. S., & Fisher, C. (2009). Speech errors reflect the phonotactic constraints in recently spoken syllables, but not in recently heard syllables. *Cognition*, *112*(1), 81–96.

White, J. (2014). Evidence for a learning bias against saltatory phonological alternations. *Cognition*, *130*(1), 96–115. https://doi.org/10.1016/j.cognition.2013.09.008

White, J., & Sundara, M. (2014). Biased generalization of newly learned phonological alternations by 12-month-old infants. *Cognition*, *133*(1), 85–90.

Yu, B. (2024). *The impact of the naturalness bias on vowel harmony patterns in transmission and communication* [Master's Thesis]. The University of Hong Kong.

Yu, B., & Do, Y. (2022). The Transmission of Vowel Harmony and Vowel Disharmony: An Iterated Learning Study. *Proceedings of the Annual Meetings on Phonology*. https://doi.org/10.3765/amp.v10i0.5439

Yun, G. H. (2006). *The interaction between palatalization and coarticulation in Korean and English* [The University of Arizona]. https://search.proquest.com/openview/eb631c45990a90663af8ad855652d61a/1?pq-origsite=gscholar&cbl=18750&diss=y&casa_token=ylKHop9C9C4AAAAA:hjigg7KddVN7Ob8OOuv9o70EfJdGgH0yUSR2VnIGOhgQvYT14XoStP1ysd_aa8r9xdzbgkUxU8pP