2024. *Proc Ling Soc Amer* 9(1). 5743. https://doi.org/10.3765/plsa.v9i1.5743.

Phonological learning is asymmetrical between prefixes and suffixes

Darby Grachek & Elsi Kaiser*

Abstract. Phonological theories typically treat prefixes and suffixes as if phonological processes apply uniformly to both. However, previous studies have found clear asymmetries between the behavior of both affixes. We argue that this is due to phonological processes being easier to learn in suffix position rather than prefix position. To test this, an artificial language learning task was used to evaluate whether phonological learning is symmetrical between affixes. Results pattern in the direction of the hypothesis, namely, that phonological learning is asymmetrically facilitated in suffix position over prefix position, providing support to general cognitive mechanisms being responsible for strong cross-linguistic tendencies.

Keywords. Phonology; morphophonology; learning bias; artificial language learning; affixes; prefix; suffix; asymmetry

1. Introduction. Typically, phonological theories treat affixes the same way (Nespor & Vogel 1986). However, several previous studies have found asymmetries between prefixes and suffixes in both typological (Dryer & Haspelmath 2013; Hyman 2002) and phonological behavior (Elkins 2020; White et al. 2018). As evidence of typological asymmetry, a study from the World Atlas of Language Structures (WALS) (Dryer & Haspelmath 2013) finds that only 16% of a diverse set of languages (969 total) use primarily prefixes, as opposed to 55% which use primarily suffixes. Evidence of a phonological asymmetry has been shown in the way that prefixes and suffixes participate in phonological processes. Specifically, prefixes are often left out of phonological processes that the stem and suffix regularly participate in (Elkins 2020).

Our experiment contributes to the understanding of this asymmetry by investigating whether a phonological learning bias underlies the differences in phonological behavior between prefixes and suffixes. Our results show that participants learn and apply phonological processes more regularly to suffixes than to prefixes. This suggests that phonological learning is facilitated slightly more across suffix-stem boundaries than prefix-stem boundaries.

To illustrate the phonological asymmetry that has been observed between prefixes and suffixes, (1) shows an exceptional segmental pattern involving Russian prefixes. In this pattern, there is a word-internal process which palatalizes consonants preceding high and mid-front vowels. This is shown in (1a) to apply to suffixes, demonstrating that suffixes participate in word-internal phonological processes. There is also a word-external velarization process in Russian, which velarizes the consonant following high/front vowels instead of palatalizing them (shown in (1b)). If prefixes are identical to suffixes in terms of their phonological behavior, they should also be subject to the word-internal palatalization process. However, as shown in (1c), they are instead subject to the word-external velarization process. This showcases a typical pattern of exceptionality for prefixes—they tend to behave as if they are outside the word domain, while suffixes tend to undergo word-internal processes.

* Huge thank you Stephanie Shih and Louis Goldstein for their continuous encouragement and guidance throughout this project. We would also like to thank both the USC PhonLunch group and Psycholinguistics Lab for providing insights to both the content and presentation of these materials. Authors: Darby Grachek, University of Southern California (grachek@usc.edu) & Elsi Kaiser, University of Southern California (emkaiser@usc.edu).

(1) Russian palatalization and velarization (from Gribanova 2008[6-7] summarized by Elkins 2020)

a. *Word-internal palatalization*

/obid + e/	→	[obʲidʲe]	‘offense.DAT’	
/alʲt + ist/	→	[alʲtʲist]	‘viola player’	

b. *Word+word velarization*

/ugol ivana/	→	[ugolˠɨvana]	‘Ivan’s corner’	*[ugolʲivana]
/sad iriny/	→	[satˠɨriny]	‘Irina’s garden’	*[satʲiriny]

c. *Prefixes trigger velarization*

/ot + iskatʲ/	→	[otˠɨskatʲ]	‘find.INF’	*[otʲiskatʲ]
/ob + ide/	→	[obˠɨde]	‘about Ida’	*[obʲide]

Phonological exceptionality associated with prefixes is widespread according to Elkins (2020), who cites several more examples of prefixes behaving as if they are external to the phonological domain of the stem. The phonological asymmetry is especially interesting because it leaves open questions about what sort of affixal information is relevant to phonological theories. Several studies investigating the behavior of prefixes and suffixes cite possible differences in the morphological relation between the stem and each affix type (e.g. Hyman 2008; White et al. 2018; Wynne et al. 2021). This body of work suggests that prefixes are not as related to the stem as suffixes are, and a general structure reflecting this is proposed by Hyman (2008), which can be seen in (2).

(2) Proposed morphological structure for stems, prefixes, and suffixes (from Hyman 2008:323)

✓ [prefix [stem suffix]]   X [prefix stem suffix]   X [[prefix stem] suffix]

Instead of assuming that both affix types behave uniformly as part of the stem, theories of phonology may be able to more easily account for exceptional prefix behavior by considering that prefixes and suffixes may be slightly different from one another. The structure in (2) is one way to think about the relation between prefixes, stems, and suffixes. Additionally, the growing psycholinguistic and learning bias literature provide an opportunity for phonological theories to make more direct connections between human language processing/cognition and prefix/suffix asymmetries. Though structural explanations for the asymmetries have been offered (e.g. Hyman 2008), our study explores a further possibility, namely that the observed asymmetries could be due to a phonological learning bias. If this is the case, we should see learning differences between phonological processes in different affix positions. Simply put, maybe the reason that prefixes are exceptional more often is because phonological processes are harder to learn in prefix position than in suffix position. This view will be pursued by considering evidence from an artificial language learning experiment, where we asked participants to learn a harmony pattern that applies to both prefixes and suffixes. Our results suggest that there is some limited evidence for a phonological learning bias which favors learning across the suffix-stem boundary over the prefix-stem boundary.

2. Background.

2.1. LANGUAGE LEARNING BIASES. Strong cross-linguistic tendencies are sometimes hypothesized to arise because language learning can be biased toward a particular pattern (Culbertson 2012). In the case of the prefix/suffix asymmetry, some have suggested that humans may be biased against using prefixes, which causes prefixes to be less prevalent than suffixes across the world's languages. Hupp et al. (2009) show that English speakers judge sequences of words, musical notes, and shapes to be more similar when they share common features at the beginning of the sequence, rather than the end. Hupp et al. (2009) conclude that there must be a learning bias against anything which disturbs a speaker's ability to recognize the beginning of a sequence. When the sequence in question is a word, prefixes are predicted to be less favored, since they would interfere with the beginning of the word sequence. This has the the potential to explain the prefix-suffix asymmetry as well. It is possible that due to the strong preference to preserve the left edge of the stem, phonological processes what would affect that portion of the stem are dispreferred. This may be why phonological processes frequently fail to apply across prefix-stem boundaries, but not across suffix-stem boundaries. It also follows that learners may find it more difficult to learn phonological processes which would apply across prefix-stem boundaries.

Martin & Culbertson (2020) argue against a universal phonological learning bias by investigating Kîîtharaka, a Bantu language spoken in eastern Kenya. Kîîtharaka is a typologically rare language, because like many Bantu languages, it uses more prefixation than suffixation. Interestingly, Martin & Culbertson (2020) obtain results that go against Hupp et al. (2009), finding that sequences which were similar at the end were judged to be more related than sequences which were similar at the beginning. Based on these results, Martin & Culbertson (2020) suggest that the preference for suffixes found in Hupp et al. (2009) is not a universal phonological learning bias, but instead a tendency that follows from the patterns in ones native language.

It is also possible that differences in linguistic structure (unique to individual languages) might cause phonological learning biases to surface differently. In the case of English, the prominence of suffixation means that there is more variation at the ends of words than at the beginning. This means that for English speakers, the left edge of the stem is more reliable for word recognition than the right edge. The opposite is true for Kîîtharaka — more prominent use of prefixes means there is more variation the left edge of words, making the right edge of words more useful for word recognition. The collective results from Hupp et al. (2009) and Martin & Culbertson (2020) show that speakers tend to rely on the portion of a word that varies less often in order to identify complex words. If we interpret Martin and Culbertson's (2020) results in this way, then they still provide evidence that is compatible with a phonological learning bias in their study on Kîîtharaka speakers, they just show a preference for the right edge instead of the left edge.

2.2. LANGUAGE PROCESSING. A processing-based account for the asymmetry between prefixes and suffixes has been proposed in the literature by Hawkins & Cutler (1988). Similar to Hupp et al. (2009) and Martin & Culbertson (2020), Hawkins & Cutler (1988) highlight the importance of the stem in word recognition of morphologically complex forms. In order to modify a stem with an affix, Hawkins & Cutler (1988) claim that the stem must be recognized by the speaker first. The need to recognize a stem before its affixes is predicted to be more difficult in a prefixed word than a suffixed one. Since prefixes are ordered before the stem, the speaker must hold the prefix in their working memory while they wait to hear the following stem in order to access it in the lexicon. It is only after the speaker has accessed the lexical entry for the stem that they can apply the modification from the prefix. This is argued to take more time in prefixed words than in suffixed words, since modifications to stems from suffixes occur after the stem (allowing the stem to be recognized before the affix is ever encountered).[1]

2.3. PHONOLOGICAL LEARNING. While Hawkins & Cutler (1988), Hupp et al. (2009) and Martin & Culbertson (2020) investigate prefix/suffix asymmetries in the word processing domain, White et al. (2018) focus more on the phonological domain. Using an artificial language learning study, they investigate behavioral biases associated with prefixes and suffixes in local and non-local harmony patterns. They ask whether affix type can influence the likelihood that a participant would infer a local or non-local harmony pattern. Speakers of several different languages were tested (Dutch, English, French, German, and Greek), receiving input that was ambiguous between a local and non-local frontness/backness harmony pattern.

Their results showed that participants did in fact have a greater preference for local harmony with suffixes rather than prefixes. However, they also report differences in the robustness of this preference across languages. While English, German, and French speakers behaved as their hypothesis predicted, Dutch speakers, although having the strongest preference for local suffix harmony, also had a tendency to accept items with non-local prefix harmony and non-local suffix harmony. Differences in responses across different languages may indicate that there is actually a baseline difference in vowel harmony learning due to the affix location alone, with locality being an additional effect on top of that. Given the result of White et al. (2018), and evidence for increased difficulty in recognizing prefixed words (Hawkins & Cutler 1988; Hupp et al. 2009), it is important to ask whether phonological processes are also more difficult for speakers to apply to prefixes than suffixes.

2.4. RESEARCH QUESTIONS. We conducted a study investigating whether the learning of phonological alternations is easier in prefix position than in suffix position using an artificial language learning task. We ask whether there is a difference in speakers' ability to learn phonological patterns between affixes of different types. Specifically, are phonological patterns harder to learn in prefix position than in suffix position? If this is the case, it would provide an explanation for the frequent phonological exceptionality found in prefix position which is tethered to the language learning and processing biases argued for in previous work (Hawkins & Cutler 1988; Hupp et al. 2009; White et al. 2018).

3. Methods.

3.1. PARTICIPANTS. Responses were collected from 50 English-speaking participants online via Prolific (Palan & Schitter 2018). Participants were paid $9 for their participation. Since Martin & Culbertson (2020) have shown that language specific differences can affect how learning biases surface, only native speakers of English were surveyed in this study. Three participants were excluded from data analysis since they reported being non-native speakers of English. An additional five participants were excluded for not reaching the 80% accuracy threshold on filler items. One more participant met the filler accuracy threshold, but answered with only 'yes' responses

[1] This is presumably true for the majority of the world's languages, since many languages use suffixation more frequently than prefixation (e.g. Dryer & Haspelmath 2013 and many others), although this does pose questions for languages which rely on the right edge of the stem for word recognition instead of the left edge. A more thorough investigation into language-specific differences is needed to understand how the position of the prefix would interact with a language with right-edge preferences for word recognition.

on all other conditions and was excluded for not performing the task properly. This resulted in responses from 41 native English-speaking participants being included in data analysis.

3.2. DESIGN AND MATERIALS. This study tested whether the ability to learn phonological processes is the same across both the prefix-stem and suffix-stem boundaries using a frontness/backness vowel harmony pattern (which is attested in human language but does not exist in English). A frontness/backness pattern was chosen because it has been used successfully by previous artificial language learning tasks (White et al. 2018), and because it has been shown that English-speaking infants have an easier time segmenting words using frontness/backness harmony as opposed to using a height harmony pattern (Mintz et al. 2017). We used the same vowels as Mintz et al. (2017): [i] and [e] for front vowels and [o] and [u] for the back vowels. A within-subjects design was used, meaning that each participant learned the harmony pattern with one prefix, and one suffix. If there is a learning bias which favors learning across the suffix-stem boundary, participants should show better performance when their learning is tested on forms with suffixes rather than with prefixes. In the rest of this section, we discuss the properties of our stimuli. In Section 3.3.3 below, after presenting the procedure, we present the 4 experimental conditions that we tested.

3.2.1. STIMULI: LINGUISTIC AND VISUAL. Our study used both linguistic and visual stimuli. The linguistic stimuli consisted of nonce word stems (e.g. [bibi], [gige]) and nonce affixes (e.g. [ʃe-], [-mu]). All stems are CVCV forms and affixes are CV forms to avoid any potential confounds from effects of syllable structure on learnability. We used a stem-controlled harmony pattern because it is more common cross-linguistically (Rose & Walker 2011), and has been used in similar studies (White et al. 2018). The stem-controlled pattern means that vowel features of the stem condition changes in the vowels of the affixes. For example, if a suffix of the form [-mu] combines with a stem like [bibe], the combined stem+affix form would be [bibe-mi], not [bubo-mu]. In other words, the suffix changes, not the stem. Within stems, vowels always harmonized with one another in order to avoid any interactions with harmony locality (White et al, 2018). This meant stems like [biti] and [vugo] are consistent with the vowel harmony pattern, since both vowels in each stem match in their frontness/backness feature. Stems that do not match in their frontness/backness features, like [bitu] and [vego], were not included. Stems with front and back vowels were also balanced so that participants saw an equal number of front and back vowel stems in the stimuli. There were 120 unique stems, 2 unique affixes, and 36 fillers total in the entire task. Participants saw all 120 unique stems over the course of the study.

As White et al. (2018) report, stress location may interact with learning in vowel harmony processes. This study intends to remain neutral with respect to the influence of stress on phonological learnability, so each syllable in test and filler items was recorded separately as a monosyllabic word (recorded on an iPhone X using the voice memos application (iOS version 16.0.2)). Then, each stimulus was constructed by splicing the individual syllables together in Praat (Boersma & Weenink 2016). This was repeated to create the affixed forms of each stimulus as well. All stem stimuli had consonants that are either stops or affricates which made splicing consistent between all items, and each affix used a nasal consonant and a fricative consonant respectively. Different consonants were used in the affixes than in the stems to encourage participants to recognize the affixes as new elements being added to the stem, avoiding a possible interpretation that that the affixes were simply another syllable in the stem.

In addition to linguistic stimuli, the study also includes a visual component. We used a picture-based paradigm to associate all stems (target and filler stems) with everyday objects and recognizable animals. Specifically, in each trial in the learning phase, a picture of an animal or an everyday object appeared on the left side of the screen, followed by a pre-recorded sound file of the stem starting 500 ms later. 1000 ms later, the same picture appeared on the right side of the screen with a blue orb above or below the object (indicating the meaning of the affix, see below). The audio for the affixed form also played 500 ms after that. These images both stayed on the screen for 1000 ms more, and then disappeared from the screen until a new trial appeared automatically. Figure 1 shows an example of an exposure phase. This presentation method is very similar to the method used by White et al. (2018).

Figure 1. Example of an exposure trial with a suffix. The stem image is shown on the left side along with the stem audio [gobo]. Then, the affixed image (object + blue orb) appears on the right side, along with the affix audio [gobo-mu]. (This example shows a situation where [mi/mu] is associated with the meaning ‘below’, as shown by the blue circle being below the object.)

To help participants to understand the nonce word affixes as affixes (as opposed to longer stems or compounds), consistent meanings – namely location information – were visually associated with each nonce affix, modifying the meanings of the s tems. There were two affixes that could attach to stems, each with a front and back vowel form: [mi/mu] and [ʃe/ʃo]. Crucially, as indicated by the location of the blue orb on the screen, relative to the object, each affix was associated with the meaning ‘above’ or ‘below’. An example of this can be found on the right side of Figure 1. This was counterbalanced between participants: some participants learned that [mi/mu] meant above and [ʃe/ʃo] meant below. Others learned the opposite association.[2]

3.2.2. FILLER TRIALS. The test phase of the experiment also included filler t rials. No fillers were included in the exposure phase where participants were being trained on the meaning and harmony patterns. Filler trials were included to balance between ‘yes’ and ‘no’ responses across the experiment and also allowed us to check whether participants were paying attention to the task. Filler trials where we expected yes responses had a stem presented on the left side of the

[2] The meanings ‘above’ and ‘below’ were chosen over more commonly used meanings in artificial language learning tasks, like plural and diminutive, in order to avoid meanings that are already associated with existing affix locations in English. For example, plurality is already indicated with a suffix ([-s]) in English, so it is possible that participants could be better at learning in suffix position, not because of any phonological learning bias, but because they are already more familiar with the concept of plurality in suffix position. The preposition-type meanings were chosen to avoid this possibility, since although English prepositions occur before the stem, they are not affixes.

screen with an image and audio to communicate the meaning, and the right side of the screen simply repeated the exact same image and audio (with no affix added). Filler trials where we expected no responses were of a few different types. One type was just like the yes response fillers, except that instead of having the same stem twice, they had the same image on both sides of the screen, but with different audio. Participants should have easily said ‘no’ to these fillers since the meaning represented on both sides of the screen was the same (as indicated by the image), but the word they heard was different (as indicated by the audio). This filler used bare stem forms in some cases, and affixed forms in other cases. Other types of no response fillers included items which had a completely new affix form not present in the exposure phase, and fillers that had the same audio on both sides of the screen, but with mismatched images. The practice trials also included these kinds of trials, so participants were familiar with them.

3.3. PROCEDURE. This task was built and run online using PCIbex (Zehr & Schwarz 2018). The procedure for artificial language learning tasks typically includes both an exposure phase, and a test phase.

3.3.1. EXPOSURE PHASE. Participants first participated in a short training session for each phase which used different phonological patterns and meanings from the test stimuli. This was done to help the participants understand what to do once the actual experiment started. After this, the exposure phase began. In the exposure phase, each stimulus stem and affix were associated with a meaning via picture and audio presentation (discussed in Section 3.2.1). All stimuli in the exposure phase were “correct”, meaning that the correct images were always shown with the correct meaning, and all stem+affix combinations were always harmonic. Participants were implicitly tasked with two things in this phase of the experiment, namely to learn the meaning of each affix (*above/below*) and also to learn how the harmony pattern works. Participants were subsequently tested on what they learned in the test phase.

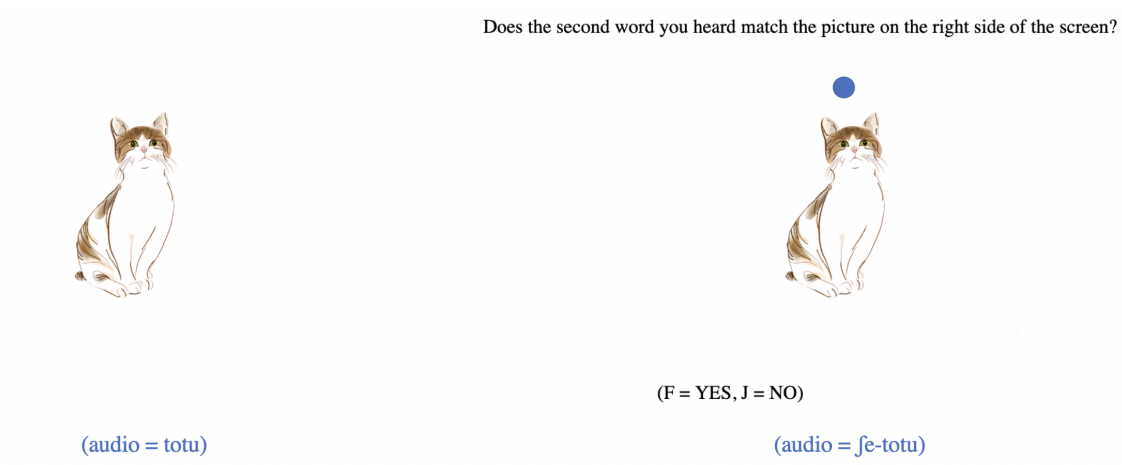

Figure 2. Example of a test trial. The stem image was presented on the left side along with the stem audio [totu]. Then, the affixed image (object + blue orb) appeared along with the affix audio [ʃe-totu]. This is an example of a condition where the harmony pattern is incorrect, but the meaning is correct (HiMc). (This example shows a situation where [ʃe-] is associated with the meaning ‘above’, as shown by the blue circle being above the object).

3.3.2. Test phase. The test phase introduced new stimuli to the participant to see how well they could generalize the patterns from the exposure phase to new data. Presentation of images and audio worked very similarly to the exposure trials: there was a new stem introduced on the left side of the screen using an image and audio presentation to communicate the stem meaning. However, instead of all target stimuli being "correct", there were some trials which had the incorrect meaning for an affix, or an incorrect (dis-harmonic) harmony pattern between the stem and affix. The right side of the screen also contained a prompt at the top of the screen which asked the participant the following question: *"Does the second word you heard match the picture on the right side of the screen?"*. Participants responded by pressing the 'F' key for 'yes' (the word correctly communicates the meaning shown in the image), or by pressing the 'J' key for 'no' (the word does NOT correctly communicate the meaning shown in the image). Participants were instructed to respond as quickly and accurately as possible, and a reminder about which key meant 'yes' and 'no' was always present on the screen during the test phase. An example of a test trial is shown in Figure 2.[3]

3.3.3. Experimental conditions. Test and exposure items were split into three blocks and were presented to participants in a pseudo-randomized order (randomized within each block). Each block of the exposure phase had 24 trials (no fillers), and each test phase had 36 trials (24 targets, and 12 fillers). Thus, each participant saw 180 trials in total (72 exposure trials and 108 test trials). To ensure that that each item appeared in each condition, there were 16 lists (8 lists for each meaning counterbalance measure). Each list was assigned to approximately 3 participants, and the task took about 30 minutes to complete. There was also a post-test questionnaire which was meant to ask participants more in-depth questions about the task, such as what they thought the meaning of each affix was, whether they thought the nonce affixes were affixes or separate words, and whether they noticed the harmony pattern.

Condition	Harmony & meaning correct (HcMc)	Harmony correct, meaning incorrect (HcMi)	Harmony incorrect, meaning correct (HiMc)	Harmony & meaning incorrect (HiMi)
Image				
Audio	[ʃo-totu]	[ʃo-totu]	[ʃe-totu]	[ʃe-totu]
Affix meaning	above	above	above	above

Table 1. The four conditions used in the test phase, where prefix ʃe/ʃo means 'above'.

There were four different conditions that a participant would encounter in the test phase. One condition had both the meaning and the harmony presented correctly, which we call HcMc

[3] Artificial language learning tasks also commonly have a verification phase, where participants are further trained on the exposure phase pattern to ensure that it was learned properly before moving on to the test phase. We chose to exclude this phase from our experiment since we are interested in the accuracy of the phonological learning process, which would be concealed by the use of a verification phase before the test phase.

(**H**armony-**c**orrect **M**eaning-**c**orrect). The second condition had the correct harmony pattern, but the meaning communicated by the image did not match the meaning associated with the affix, which is referred to as HcMi (**H**armony-**c**orrect **M**eaning-**i**ncorrect). The third condition had the incorrect harmony pattern between the stem and affix, but the meaning was correct, referred to as HiMc (**H**armony-**i**ncorrect **M**eaning-**c**orrect). The fourth condition had both the harmony pattern and the meaning association between the image and affix as incorrect, referred to as HiMi (**H**armony-**i**ncorrect **M**eaning-**i**ncorrect). Examples of these conditions are shown in Table 1.

3.4. PREDICTIONS. If participants learn only the harmony patterns, they should respond with 'yes' more frequently overall in conditions where the harmony pattern is correct (HcMc and HcMi) than in conditions where it is incorrect (HiMc and HiMi). If participants learn only the meaning of the nonce affixes, they should respond with 'yes' more frequently overall in conditions where the meaning is correct (HcMc and HiMc) than in conditions where it is incorrect (HcMi and HiMi). If participants learn *both* the meaning and harmony patterns, they should respond with 'yes' more frequently in conditions where both meaning and harmony are correct (HcMc), and respond with 'no' more frequently in every other condition (HcMi, HiMc, and HiMi).

To support the hypothesis that phonological learning is easier across the suffix-stem boundary than the prefix-stem boundary, we would need to see a difference in the responses between conditions for each type of affix. Specifically, we would expect that in the suffixed stimuli, participants are able to discriminate based on both the meaning AND the harmony pattern more often than for the prefixed items. This means that for suffixed items, participants are predicted to give more yes-responses to HcMc conditions, and more no-responses for all other conditions (HiMc, HcMi, HiMi). For the prefixed items, we expect participants to be attending to the meaning more often than they are attending to the harmony pattern. This means that with prefixed items, we would predict that participants would give more yes-responses to conditions with the correct meaning (HcMc and HcMi), and give more no-responses to conditions with incorrect meaning (HcMi and HiMi), paying less attention to the correctness of the harmony pattern than in the suffixed conditions.

The crucial condition for confirming this prediction is the HiMc condition. If participants are attending to only to meaning, they are expected to say 'yes' to this condition. However, if they are attending to both meaning and harmony, we predict that they would say 'no' to this condition. Essentially, to find evidence for the hypothesis we predict there to be more no responses to the HiMc condition for suffixed items than for prefixed items. In a statistical analysis, if participants are learning the meaning patterns, we expect a significant main effect of meaning. If they are learning the harmony patterns, we expect a significant main effect of harmony. If learning is easier across the suffix boundary than the prefix boundary, we expect a significant interaction between harmony and affix type.

4. Results. Before statistical analyses, trials with reaction times above 8000 ms and below 200 ms were excluded (12 data points). After removing extreme values, any additional datapoints with RTs more than 2.5 SDs from the overall mean were also removed from the data set (146 data points). All together, these measures excluded 5.3% of the total data.

For statistical analyses, generalized linear mixed effects models were used with participant response (yes, the word correctly communicates the meaning shown in the image = 1; no, the word does NOT correctly communicate the meaning shown in the image = 0) as the depen-

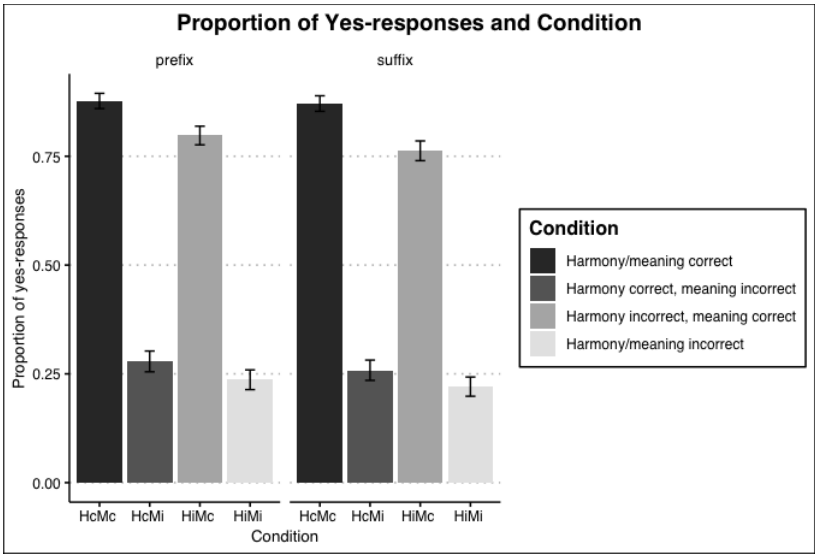

Figure 3. Summary of participant responses to each condition as a proportion of yes responses, split between prefix (left) and suffix conditions (right).

dent variable, and harmony (contrast-coded, incorrect harmony = 0.5, correct harmony = -0.5), meaning (contrast-coded, incorrect meaning = 0.5, correct meaning = -0.5), and affix type (prefix = 0.5, suffix = -0.5), and a harmony x meaning interaction as independent variables. In other words, we tested how the rate of participants saying 'yes' is influenced by (i) whether what they heard matches the spatial configuration on the screen (meaning of the affix) and on (ii) whether the stem+affix combination exhibits the vowel harmony pattern they were trained on (harmony). Random intercepts and random slopes for harmony pattern, meaning, and affix type grouped by participant were included (attempts to include any random effects for item led to either a failure to converge, or a singular model).

Models for participants' responses were estimated using the *lme4* package (version 1.1.21) (Bates et al. 2015) and *lmertest* (version 3.1.1) (Kuznetsova et al. 2017) using R studio (R Core Team 2023). For the random effects, we entered intercepts for subjects and items, as well as by-subject and by-item random slopes for the effects of harmony, meaning, affix type, and their interaction. Random effects were fully crossed and fully specified with by-subject and by-item intercepts and slopes for harmony, meaning, affix type, and their interaction. They were then reduced (starting with by-item effects) via model comparison, so that only random effects that contributed significantly to the model ($p < 0.05$) were included (Baayen et al. 2008).

The proportion of yes responses in each of the conditions is in Figure 3. At first glance, we see high rates of yes responses (i.e., participants accepting the word as correctly communicating the meaning of the image) in the conditions where meaning is correct, suggesting that participants successfully learned that the affix means 'above' or 'below' based on the exposure phase. The rates of yes responses in conditions where meaning is incorrect are much lower. What about potential differences between the prefixed vs. suffixed versions (leftmost four bars vs. rightmost four bars)? Visual assessment suggests that the crucial condition, HiMc (**H**armony **i**ncorrect, **M**eaning **c**orrect, the third bar in each set), has a numerically lower proportion of yes responses in the suffixed condition (seventh bar from the left) than it does in the prefixed condition (third bar from the left). The proportion of yes responses in other conditions appears similar between

Fixed Effects: Response

	Estimate (β)	SE	z	*p*
(intercept)	-0.6	0.48	-1.26	0.21
Harmony (incorrect = 0.5, correct = -0.5)	-0.89	0.30	-3.01	0.003**
Meaning (incorrect = 0.5, correct = -0.5)	-5.74	0.89	-6.43	< 0.001***
Affix Type (prefix = 0.5, suffix = -0.5)	0.46	0.21	2.23	0.03*
Harmony:Meaning	-0.47	0.33	-1.43	0.15
Harmony:Affix Type	0.13	0.28	0.49	0.65
Meaning:Affix Type	0.07	0.31	0.23	0.82
Meaning:Meaning:Affix Type	0.19	0.57	0.35	0.73

Table 2. Outcomes of statistical analysis (all participants). Number of obs.: 2847, group(s): participant, 41. p-values/df calculated using the Satterthwaite approximation.

prefixes and suffixes.

In the statistical analysis, the final model (Response $\sim$ harmony * meaning * affix type + (1 + harmony + meaning+ affix type | participant)) fails to show the expected interaction effect between affix type and harmony. The details of the statistical analyses are in Table 2. There are no other significant interactions in the model. However, we do see more yes responses when the harmony pattern is correct than when it is not correct (main effect of harmony, p = 0.003), and more yes responses when the image shows the correct meaning of the affix (main effect of meaning, p < .001). There are also more yes responses with prefixes than with suffixes (main effect of affix type, p = 0.03).

The main effects of harmony and of meaning indicate that participants were indeed learning the meaning of the affix and the correct harmony pattern, and responding to the stimuli accordingly. This serves as a 'sanity check' and validates the use of online data collection for this experimental method. The main effect of affix type can be interpreted as participant responses varying significantly according to whether participants are hearing a prefix or a suffix: overall, yes responses occur more with prefixes than suffixes, which may be due to an increased awareness of the incorrectness of the harmony pattern in the crucial HiMc condition for suffixes. This is noteworthy, as it shows that there are some phonological learning differences between prefixes and suffixes. However, in this dataset, affix type did not interact significantly with the predictor crucial to the hypothesis (harmony).

4.1. Post-hoc analyses and discussion. After looking at the post-test responses, as well as the pattern of responses for individual participants, it became clear that many participants – around 65% – failed to learn the harmony pattern (i.e., the fact that the backness/frontness of the stem vowels controls the backness/frontness of the affix vowels, regardless of whether they are prefixes or suffixes). Many only showed evidence of discriminating prefixed and suffixed items on the basis of their meaning, and some even failed to discern any meaning from the affixes at all. In light of this, we conducted follow-up analyses looking only at the subset of participants who showed evidence of learning that there is a harmony pattern between the stem and the affix. Figure 4 shows only those participants whose responses indicated that they had learned the harmony pattern (14 out of 41); whom we call the harmony learners. Here, focusing on the critical HiMc condition, we now see in Figure 4 that the asymmetry between the prefixed and suffixed conditions is clearer: participants are better at discriminating the harmony pattern with the suffixed items than with the prefixed items.

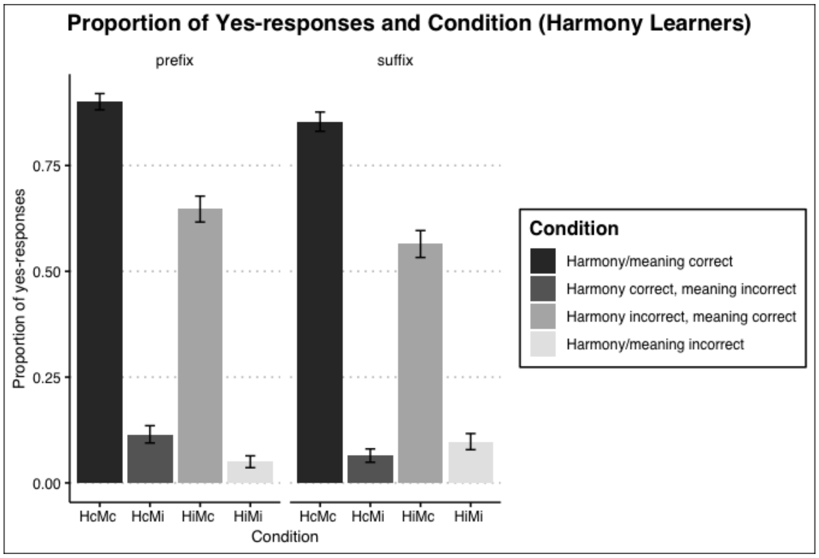

Figure 4. Subset of 14 participants who showed evidence of learning vowel harmony. Summary of responses to each condition as a proportion of yes responses, split between prefix (left) and suffix conditions (right).

For the statistical analyses with the harmony learners, a generalized linear mixed effects model was again used to test if participant response varied as a function of affix type (prefix or suffix), the harmony pattern (being correct or incorrect), meaning (correct or incorrect), and their interactions. Random intercepts and random slopes that varied by participant were also included (any attempts to include further random effects led to failure to converge, or model singularity).

The final model (Response $\sim$ harmony * meaning * affix type + (1 $|$ participant)) finds more yes responses when the affix+stem combinations are harmonic (main effect of harmony, $p < 0.001$), and more yes responses when the image matches the affix meaning (main effect of meaning, $p < 0.001$), and an interaction between meaning and harmony ($p = 0.00097$). The detailed statistics are in Table 3. These effects show evidence that the correct meaning and harmony patterns were being learned by this subset of participants. No main effect or interaction effects for affix type were found. However, the expected interaction affect between affix type and harmony is now marginally significant ($p = 0.08$), which shows some limited evidence in favor of the hypothesis that phonological learning is asymmetrically facilitated in suffix position over prefix position. Thus, with those participants who learned the harmony patterns, we see hints of them learning this pattern better with suffixes than prefixes, in line with our hypothesis.

5. Discussion and conclusion. This study set out to investigate asymmetries between prefixes and suffixes. We tested the hypothesis that a phonological learning bias makes learning easier for phonological processes applying across suffix boundaries than prefix boundaries. This hypothesis is rooted in cross-linguistic work pointing to frequent phonological exceptions in prefix position (Elkins 2020; Gribanova 2008; Poser 1990). We used an artificial language learning task to test

Fixed Effects: Response

	Estimate (β)	SE	z	*p*
(intercept)	-0.73	0.26	-2.82	0.005**
Harmony	-0.95	0.22	-4.34	$<$ 0.001***
Meaning	-4.12	0.25	-16.6	$<$ 0.001***
Affix Type	0.18	0.22	0.82	0.41
Harmony:Meaning	1.44	0.44	3.3	0.00097***
Harmony:Affix Type	-0.77	0.43	-1.8	0.08 .
Meaning:Affix Type	-0.54	0.43	-1.23	0.22
Meaning:Meaning:Affix Type	-1.16	-1.34	0.35	0.18

Table 3. Outcomes of statistical analysis (only those participants who learned harmony patterns). Number of obs.: 976, group(s): participant, 14. p-values calculated using the Satterthwaite approximation.

whether a phonological learning bias was present in English-speakers when learning a vowel harmony pattern on a novel language. Our results showed that participant responses varied depending on the type of affix found in the test item. Specifically, participants showed evidence that they were attending to the harmony pattern more closely in the suffix condition than in the prefix condition. However, the increased ability to discriminate stimuli using the harmony pattern did not show a significant increase in accuracy levels between prefixes and suffixes. Numerically, participant responses go in the expected direction, but the interaction effect (between harmony and affix type) did not reach significance.

However, additional analyses provide somewhat clearer hints in favor of the hypothesis. When we look at responses from only the small subset of participants who showed evidence of learning the harmony pattern, we find that participants are marginally more accurate in discriminating stimuli using vowel harmony in suffix conditions than in prefix conditions. This suggests that phonological learning was facilitated more in suffix position than in prefix condition.

Identifying the specific underlying cause for the suffix-prefix asymmetry requires further research. However, several accounts have been proposed in the literature thus far, including the asymmetrical processing account from Hawkins & Cutler (1988), and the importance of the beginnings of words for stem recognition from Hupp et al. (2009), both of which could be extended to account for a phonological learning asymmetry.

References

Baayen, R. Harald, Douglas J. Davidson & Douglas M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59. 390–412. https://doi.org/10.1016/j.jml.2007.12.005.

Bates, Douglas, Martin Mächler, Ben Bolker & Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1). 1–48. https://doi.org/10.18637/jss.v067.i01.

Boersma, Paul & David Weenink. 2016. Praat: Doing phonetics by computer [computer program]. version 6.0.19.

Culbertson, Jennifer. 2012. Typological universals as reflections of biased learning: Evidence from artificial language learning. *Language and Linguistics Compass* 6/5. 310–329.

Dryer, Matthew S. & Martin Haspelmath. 2013. The world atlas of language structures online. Leipzig: Max Planck Institute for Evolutionary Anthropology (Available online at http://wals.info, Accessed on 2022–11–15).

Elkins, Noah Eli. 2020. *Prefix independence: Typology and theory*. Los Angeles: University of California MA thesis.

Gribanova, Vera. 2008. Russian prefixes and prepositions in stratal OT. In Charles B. Chang & Hannah J. Haynie (eds.), *WCCFL 26*, 217–225. Somerville, MA: Cascadilla Press.

Hawkins, John A. & Anne Cutler. 1988. Psycholinguistic factors in morphological asymmetry. In John A. Hawkins (ed.), *Explaining language universals*, 280–317. Oxford: Blackwell.

Hupp, Julie M., Vladimir M. Sloutsky & Peter W Culicover. 2009. Evidence for a domain-general mechanism underlying the suffixation preference in language. *Language and Cognitive Processes* 24. 876–909. https://doi.org/10.1080/01690960902719267.

Hyman, Larry. 2002. Is there a right-to-left bias in vowel harmony? Presentation at the 9th International Phonology Meeting. Vienna, Nov. 1, 2002.

Hyman, Larry. 2008. Directional asymmetries in the morphology and phonology of words, with special reference to bantu. *Linguistics* 45(2). 309–350.

Kuznetsova, Alexandra, Per Bruun Brockhoff & Rune Haubo Bojesen Christensen. 2017. lmertest package: Tests in linear mixed effects models. *Journal of Statistical Software* 82. 1–26.

Martin, Alexander & Jennifer Culbertson. 2020. Revisiting the suffixing preference: Native-language affixation pattern influence perception of sequences. *Psychological Science* 31(9). 1107–1116.

Mintz, Toben H, Rachel L Walker, Ashlee Welday & Celeste Kidd. 2017. Infants' sensitivity to vowel harmony and its role in segmenting speech. *Cognition* 171. 95–107.

Nespor, Marina & Irene Vogel. 1986. *Prosodic phonology*. Dordrecht: Foris.

Palan, Stefan & Christian Schitter. 2018. Prolific.ac—A subject pool for on-line experiments. *Journal of Behavioral and Experimental Finance* 17. 22–27. https://doi.org/10.1016/j.jbef.2017.12.004.

Poser, William. 1990. Word-internal phrase boundary in japanese. In Sharon Inkelas & Draga Zec (eds.), *The phonology-syntax connection*, 279–287. Stanford University and University of Chicago Press.

R Core Team. 2023. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing Vienna, Austria. https://www.R-project.org/.

Rose, Sharon & Rachel Walker. 2011. *Harmony systems*. Oxford: Blackwell.

White, James, Rene Kager, Tal Linzen, Giorgos Markopoulos, Alexander Martin, Andrew Nevins, Sharon Peperkamp, Krisztina Polgar, Nina Topintzi & Ruben van de Vijver. 2018. Preference for locality is affected by the prefix/suffix asymmetry: Evidence from artificial language learning. *North East Linguistics Society (NELS)* 48.

Wynne, Hilary, Swetlana Schuster, Beinan Zhou & Aditi Lahiri. 2021. Asymmetries in the processing of affixed words in bengali. *Language* 97(3). 599–628. https://doi.org/mzv7.

Zehr, Jérémy. & Florian Schwarz. 2018. Penncontroller for internet based experiments (ibex). Software. https://doi.org/10.17605/OSF.IO/MD832.