Published in final edited form as:

Language (Baltim). 2015 March ; 91(1): 48–72. doi:10.1353/lan.2015.0010.

LEARNING NONADJACENT DEPENDENCIES IN PHONOLOGY: TRANSPARENT VOWELS IN VOWEL HARMONY

Sara Finley

Pacific Lutheran University

Abstract

Nonadjacent dependencies are an important part of the structure of language. While the majority of syntactic and phonological processes occur at a local domain, there are several processes that appear to apply at a distance, posing a challenge for theories of linguistic structure. This article addresses one of the most common nonadjacent phenomena in phonology: transparent vowels in vowel harmony. Vowel harmony occurs when adjacent vowels are required to share the same phonological feature value (e.g. $V_{+F} C V_{+F}$). However, transparent vowels create a second-order nonadjacent pattern because agreement between two vowels can ‘skip’ the transparent neutral vowel in addition to consonants (e.g. $V_{+F} C V_{-F}^T C V_{+F}$). Adults are shown to display initial learning biases against second-order nonadjacency in experiments that use an artificial grammar learning paradigm. Experiments 1–3 show that adult learners fail to learn the second-order long-distance dependency created by the transparent vowel (as compared to a control condition). In experiments 4–5, training in terms of overall exposure as well as the frequency of relevant transparent items was increased. With adequate exposure, learners reliably generalize to novel words containing transparent vowels. The experiments suggest that learners are sensitive to the structure of phonological representations, even when learning occurs at a relatively rapid pace.*

Keywords

nonadjacent dependencies; phonology; artificial grammar learning; vowel harmony; transparent vowels

1. Introduction

Long-distance phenomena in language have posed challenges for the representation of both syntactic and phonological processes in generative linguistics. The learnability of nonadjacent dependencies in language has received a large amount of attention in the literature, but this attention has largely focused on phrase-level dependencies and word segmentation (Gómez 2002, Gómez & Maye 2005, Misyak & Christiansen 2007, Misyak et

*This article would not have been possible without the helpful discussion, advice, and suggestions from the following people: Ariel Goldberg, Matt Goldrick, Iris Berent, Colin Wilson, Paul Smolensky, Neil Bard-han, Patricia Reeder, Anna States, Lily Schieber, Carrie Miller, Emily Kasman, Elissa Newport, Michael Tanenhaus, helpful and anonymous referees, and members of the Aslin–Newport and the M-Tan labs at the University of Rochester. Portions of this work were presented at the 2009 annual meeting of the Linguistic Society of America, the UCLA Department of Linguistics, and the E. G. G. Summer School in Poznań, Poland, 2009. Funding was provided by NIH grants DC00167 and T32DC000035. All errors are my own.

al. 2009, Newport & Aslin 2004), with few studies focusing on the learnability of long-distance phenomena in phonological processes. The goal of the present article is to use experimental approaches to expand our understanding of how nonadjacent dependencies in phonology are learned. Learners are shown to be biased against SECOND-ORDER NONADJACENT DEPENDENCIES (defined below) in phonological processes (specifically transparent vowels in vowel harmony), but this bias can be overcome with sufficient training.

The vast majority of linguistic processes are constrained by locality principles (Chomsky 1981, Culicover & Wilkins 1984). Locality principles require reference to adjacent elements along a specific domain (e.g. vowels in adjacent syllables). At the surface, long-distance dependencies violate this tendency (e.g. in vowel harmony, agreement between vowels often ‘skips’ intervening consonants), posing the question of how patterns that can apply at a distance are constrained in language. One solution to this challenge is to posit that long-distance phenomena are covertly local. For example, vowel harmony often involves the spreading (or sharing) of feature values between vowels in adjacent syllables, ignoring consonants (for the most part). This adjacency is created by placing vowels and consonants on separate ‘tiers’ of representation (Clements 1976, Goldsmith 1975). With tier-based representations, segments can be nonadjacent on the surface, but adjacent according to a tier. For example, in the English word /ɛmə/ ‘Emma’, /ɛ/ and /ə/ are nonadjacent on the surface (/m/ is in between), but if /ɛ/ and /ə/ are placed on a separate tier where only vowels reside, /ɛ/ becomes adjacent to /ə/ because /m/ is now on a separate (consonant) tier. Tier-based representations allow vowel harmony processes to apply locally even when the harmony is not adjacent on the surface. For the purposes of this article, the kind of long-distance dependency found in basic vowel harmony is referred to as FIRST-ORDER NONADJACENT DEPENDENCIES (because only consonants are skipped).

In addition to first-order nonadjacent dependencies, there are second-order nonadjacent dependencies. In vowel harmony, these second-order nonadjacent patterns appear as cases of transparent neutral vowels. A NEUTRAL VOWEL refers to any vowel that does not undergo harmony. TRANSPARENT NEUTRAL VOWELS (denoted here as V^T) (i) do not undergo harmony and (ii) allow the feature-agreement process to ‘skip’ that vowel. For example, in a language in which the feature value of the first vowel in the word determines all feature values of the word (a left-to-right harmony system), a disharmonic sequence such as $[V_+V^T_-V_+]$ is possible only when the disharmonic vowel is transparent, intervening between the first and final vowels.

Hungarian provides a classic example of transparent vowels in vowel harmony. In Hungarian,¹ the first vowel of the word determines the back feature of all following vowels, including the vowel features of suffixes that attach to the stem. If the vowel feature of the initial vowel is front, all suffix vowels must also be front; if the initial vowel is back, suffix vowels must also be back (Vago 1980). However, not all vowels alternate to agree with the harmonic feature value of the initial vowel. The front vowels [i] and [e] are transparent, since the presence of these vowels has no bearing on the harmony of surrounding vowels. Transparent vowels can create instances of second-order nonadjacent dependencies. For example, the final vowel of the form $[pɔ_+lle^T_-r-nɔ_+k]$ ‘foreman-DAT’ agrees with the feature

¹Note that this example is simplified for purposes of explication.

value of the initial vowel. This occurs despite the fact that the adjacent vowel is front; spreading has ‘skipped’ the medial transparent vowel (Hayes & Londe 2006, Hayes et al. 2009, Ringen 1980).

Languages may also have vowels that are opaque to harmony, which fall into the category of first-order nonadjacent patterns. OPAQUE NEUTRAL VOWELS (denoted here as V^O) are vowels that (i) fail to undergo harmony, but (ii) do not allow ‘spreading’ to pass through them (they ‘block’ harmony). In languages with opaque vowels, the vowel that follows the opaque vowel agrees with it (e.g. $[V_+V^O_-V_-]$, where spreading is left to right, and the medial vowel is opaque). In Turkish, for example, nonhigh vowels are opaque to round harmony (Bakovi 2000, Clements & Sezer 1982, Polgardi 1999). If a nonhigh vowel follows a high round vowel, it will not undergo spreading of the rounding feature, but any high vowels that follow the opaque vowel will remain nonround (e.g. $V_+V^O_-V_-$; $[pu_+l-la^O_-r-i_-n]$ ‘stamp’). Opaque vowels can be accounted for with first-order nonadjacent dependencies; no vowels are ‘skipped’.

The second-order nonadjacency of transparent vowels may create ambiguity for the learner. For example, if V_1 determines the feature value of V_2 in the sequence $V_1 V^T V_2$, the transmission of the feature value of V_1 must ‘skip’ through V^T . When the transparent vowel has the same feature value as the initial vowel, however, the following vowel agrees with both the initial vowel and the transparent vowel (e.g. $V_1_- V^T_- V_2_-$). In addition, when the transparent vowel is not present, the final vowel is determined by the adjacent vowel (e.g. $V_1 V_2$). This means that there is only a subset of instances that provide evidence to the learner that the transparent vowel does not initiate a harmonic domain. This ambiguity may create challenges for the learner.

In addition to ambiguity creating challenges for the learner, accounting for transparent vowels typically involves additional assumptions, rule applications, or increased complexity of constraints or representations (Bakovi & Wilson 2000, Finley 2008, Goldsmith 1985, Hayes & Londe 2006, Kiparsky & Pajusalu 2003, Ringen & Heinamaki 1999, Smolensky 2006). For example, Hayes and Londe (2006) propose two constraints in order to account for the behavior of transparent vowels in Hungarian: a LOCAL and a DISTAL harmony constraint.² The local constraint Local*[+Back][−Back] is violated whenever the vowel immediately following a [+back] vowel (ignoring consonants) is a [−back] vowel. The constraint is ‘local’ because only consonants may intervene between the two vowels, and it can account for vowel harmony patterns in which all vowels participate in the harmony, or when the neutral vowel is opaque. If not for transparent vowels, only local constraints would be necessary.

The distal constraint Distal*[+Back]X[−Back] is violated whenever a [+back] vowel is followed by a [−back] vowel, regardless of the number of consonants or vowels in between. This distal constraint is required to account for second-order nonadjacent dependencies because a pair of vowels could incur a violation of the distal constraint even if the vowels were nonadjacent (e.g. a neutral vowel intervened). The distal constraint is more complex

²Note that this summary of Hayes and Londe’s (2006) analysis is simplified for ease of discussion.

than the local constraint for two reasons. First, the constraint requires additional structure (X). Second, the constraint applies to all vowels in a word, not just adjacent vowels. This demonstrates that accounting for transparent vowels in vowel harmony requires a more complex constraint than accounting for opaque vowels in harmony.

In the present study, adult learners were tested using a two-alternative forced-choice test, choosing between a [+back] suffix and a [–back] suffix. In order to better understand how participants will respond on this task, we can make use of previous formal analyses of transparent vowels in vowel harmony. These analyses not only can help us to understand the types of representations that speakers may hold, but also can help us to form predictions about how these representations might be learned. Table 1 illustrates the basics of Hayes and Londe’s (2006) local and distal harmony constraints, as well as how learners may respond in an artificial grammar learning task. The table contains three hypothetical candidate sets, each with two representative candidates. In classic OPTIMALITY THEORY, the candidate set is infinite (Prince & Smolensky 2004 [1993]), but participants in the present study were forced to choose between two alternatives that were identical except for the final vowel, which was either [e] or [o]. If learners correctly apply the harmony constraint, they will choose [e] or [o] appropriately depending on the back and round features of the previous vowels. In examples 1/2 in the table, the choice is between harmonic [pu₊to₊ko₊]/[pe₊te₊ke₊] and disharmonic *[pu₊to₊ke₊]/*[pe₊te₊ko₊]. The harmonic candidates do not violate either harmony constraint, and will be selected regardless of which constraint is dominant. In example 3, the first two (‘stem’) vowels disagree for the back feature, which incurs violations of the harmony constraints. The transparent candidate ([pu₊te₊ko₊]) and the opaque candidate ([pu₊te₊ke₊]) incur one violation each of both Local*[+Back][–Back] and Distal*[+Back]X[–Back] for /u₊/ followed by /ε₊/ (the stem disharmony). The opaque candidate incurs a second violation of Distal*[+Back]X[–Back] for /u₊/ followed by /e₊/. The transparent candidate also violates Local*[–Back][+Back] and Distal*[–Back]X[+Back] for /ε₊/ followed by /o₊/, while the opaque candidate does not violate either constraint for /ε₊/ followed by /e₊/. Thus, the transparent candidate fares worse than the opaque candidate on every harmony constraint except Local*[+Back][–Back] (where both candidates incur a single violation) and Distal*[+Back]X[–Back], in bold (where the opaque candidate incurs two violations and the transparent candidate incurs just one). Because the opaque candidate will ‘win’ under more rankings of the four harmony-inducing constraints than the transparent candidate, one may expect that learners will be biased toward the rankings that produce the opaque candidate, predicting that opaque neutral vowels will be easier to learn than transparent neutral vowels.

The above analysis, along with a basic theory of learning, can be used to understand how learners might behave after exposure to either an opaque neutral vowel or a transparent neutral vowel. Assuming a model in which learners induce constraints from the input (Hayes & Wilson 2008), adults exposed to vowel harmony for the first time (e.g. in a vowel harmony learning experiment) should start with no harmony constraints. In a two-alternative forced-choice test between a harmonic and a disharmonic item, participants in the initial state (or a control condition) will select the harmonic item about 50% of the time. Following sufficient exposure to the harmony pattern, learners will induce a harmony constraint and will select the harmony pattern at a rate significantly greater than chance.

There are several reasonable scenarios for how learners might infer a harmony constraint following minimal exposure to a vowel harmony pattern with a neutral vowel. First, learners may infer a harmony constraint (from the four in Table 1) at random. Because two of the four constraints prefer the opaque candidate, and only one constraint prefers the transparent candidate (the fourth constraint, Local*[+Back][−Back], shows no preference), learners exposed to opaque vowels will correctly infer the correct constraint 50% of the time, while learners exposed to transparent vowels will infer the correct constraint 25% of the time. This means that learners exposed to the transparent neutral vowel will be more likely to require additional exposure to infer the correct constraints and rankings. A second possibility is that learners infer constraints in order of complexity, with the simplest first. Because the only constraint that prefers transparent neutral vowels is a more complex distal constraint, learners will infer the appropriate constraint for opaque neutral vowels before the one for transparent neutral vowels. Under both scenarios, learners should require less training to acquire the behavior of the opaque than the transparent neutral vowel.

A final possibility is that learners will, without any exposure, prefer either opaque or transparent neutral vowels. In this case, even participants in the control condition will show a preference for opaque or transparent neutral vowels. This would be equivalent to learners inferring the appropriate harmony constraints without any input or training. This possibility is unlikely, however, since participants in control conditions in previous vowel harmony learning experiments (especially the ‘no training’ control in Finley & Badecker 2009a) showed no preference for vowel harmony.

The present article uses an artificial grammar learning paradigm to test the hypothesis that opaque neutral vowels (first-order nonadjacency) should be easier to learn than transparent neutral vowels (second-order nonadjacency). An artificial grammar learning paradigm provides a forum to test differences between transparent and opaque neutral vowels without confounding the language and the lexicon. One can make two minimally different artificial languages that differ only in whether the neutral vowel is opaque or transparent. In addition, naturalistic studies on the acquisition of neutral vowels in harmony are difficult because vowel harmony production data are relatively error-free (Leiwo et al. 2006, MacWhinney 1978, Slobin 1997). Further complicating matters, when children do make production errors, it is unclear whether the error is a misapplication of harmony or a mispronunciation of the intended vowel (Leiwo et al. 2006). It is also possible that if biases against transparent neutral vowels exist, these biases may only be applicable in the very early stages of learning, making them difficult to study in a naturalistic setting. An artificial grammar learning paradigm makes it possible to control several aspects of exposure, allowing for detection of short-lived biases.

There is evidence for a bias toward first-order nonadjacent dependencies in phonology from consonant harmony (Finley 2011); given training data that is ambiguous between a first-order and a second-order nonadjacent consonant harmony pattern, learners were more likely to infer the first-order pattern. Additional experimental evidence has suggested biases against nonadjacent patterns in word segmentation (Newport & Aslin 2004, Onnis et al. 2005) and against nonadjacent dependencies at the phrasal level (Gómez 2002, Gómez & Maye 2005, Misyak & Christiansen 2007, Misyak et al. 2009, Saffran 2001). Newport and

Aslin (2004) found that adult learners in a speech segmentation task were unable to segment dependencies between nonadjacent syllables (e.g. [dikitae] and [digutae]), but were able to segment the speech stream when the nonadjacent dependencies were at the segment level (e.g. [dokibae] and [dakube]). Learners therefore appear to compute statistics across nonadjacent consonants, ‘skipping’ vowels. This suggests that the transitional probabilities used to segment speech are calculated at a level of representation that includes consonants and vowels on separate tiers. This supports the hypothesis that learners make use of the same representations proposed in the theoretical phonology literature (Ettlinger et al. 2012, Onnis et al. 2005), thereby supporting the links made earlier that connect adult learning to theoretical phonology.

Previous studies have shown that adult learners are adept at acquiring the first-order nonadjacent dependencies in vowel harmony (Finley & Badecker 2007, 2009a,b, 2012, Moreton 2008, Pycha et al. 2003). For example, learners make correct judgments regarding morphophonological alternations based on the back/round features of the stem vowel (e.g. correctly choosing back/round harmonic /bede-mi/ over disharmonic */bede-mu/). The present study extends these findings by testing for biases against second-order nonadjacent dependencies in transparent neutral vowels in vowel harmony. In experiment 1, learning the behavior of neutral vowels with a first-order nonadjacent pattern (the opaque vowel) is compared to learning a second-order nonadjacent pattern (the transparent vowel). If learners show a bias against second-order nonadjacent patterns in vowel harmony, the behavior of the opaque neutral vowel should be easier to learn than that of the transparent neutral vowel. In experiments 2–5, the properties of the transparent vowel and the properties of the training set that allow adults to reliably learn the behavior of the transparent vowel in a disharmonic context are explored. A summary of the experiments and their results and predictions is given in Table 2.

2. Experiment 1

Participants were trained on a back/round vowel harmony pattern with either a transparent or an opaque neutral vowel and tested on their learning compared to a control condition that was not exposed to a vowel harmony pattern.

2.1. Method

Participants—Fifty-six participants were recruited from the University of Rochester community and were paid \$10 for their participation. All participants were adult monolingual native English speakers with no knowledge of a vowel harmony language. Final analyses included seventeen participants in each of three conditions: a Transparent condition, an Opaque condition, and a Control condition (which consisted of training on a mixture of harmonic and disharmonic stems, without exposure to the harmony pattern). The data for two participants were excluded because they were not monolingual native English speakers. The data from three additional participants were discarded due to experimenter error.

Stimuli and design—Participants were randomly assigned to one of three training conditions: Transparent, Opaque, and Control. The Transparent and Opaque conditions

contained training stimuli that represented a back/round vowel harmony pattern with a neutral vowel (transparent in the Transparent condition and opaque in the Opaque condition). The Control condition contained items that did not represent any phonological pattern, and served as a baseline for comparison between the two critical conditions.

An adult male speaker of American English who had no explicit knowledge of the experimental design or the experimental hypothesis produced each stimulus in the training and test phases of the experiment for all conditions. The speaker was told to produce all stimuli items without reducing vowels, but to place the main stress on the initial syllable. Note that several vowels were produced as diphthongs, since English does not readily contain pure vowels.

Training stimuli: transparent and opaque conditions—Participants in the Transparent and Opaque training conditions were presented with pairs of items that conformed to a back/round vowel harmony pattern. The first item in each pair was a ‘stem’ item of the form CVCVC (e.g. [gipek], [budok]), and the second item contained a ‘suffixed’ form with either [–e] (following front, unrounded vowels [i, e], e.g. [gipeke]) or [–o] (following back, round vowels [u, o], e.g. [budoko]).³ All consonants were chosen from the set [p, t, k, b, d, g, m, n] and all vowels from the set [i, e, u, o, ε]. The vowels [i, e, u, o] have harmonic counterparts in the inventory, but the front vowel [ε] has no harmonic counterpart and may appear in disharmonic contexts within the stem. Examples of training stimuli are given in Table 3.

The initial vowel for all twenty-four stem forms was drawn from the harmonizing set of vowels [i, e, u, o], and these were evenly distributed. Thus, there were six words that contained [i] as the initial vowel, six that contained [e] as the initial vowel, and so forth. The second stem vowel contained the same harmonizing vowels in sixteen of the stem items (i.e. if the first vowel was [i], the second vowel was either [i] or [e]), and these were evenly distributed with respect to front and back vowels (half of the stems induced a back vowel suffix and half induced a front vowel suffix). The second stem vowel was the neutral vowel [ε] in the remaining eight stem forms. When the stem contained [ε] in second position, the behavior of the suffix varied based on the training condition. In the Opaque condition, the suffix vowel was always [–e] after [ε] (e.g. [dotεb-dotεbe]). In the Transparent condition, the suffix vowel varied between [–e] and [–o] after [ε], depending on the initial vowel. When the initial vowel was front ([i, e]), the suffix vowel was [–e] (e.g. [mepen-mepene]); when the initial vowel was back ([o, u]), the suffix vowel was [–o] (e.g. [dotεb-dotεbo]). The difference between the Transparent and Opaque conditions was apparent only in the four items that contained a back vowel and the neutral vowel.

While many of the vowels in the stem were identical (e.g. [netep]), only four of the items contained identical vowels in the stem and the suffix. Further, Finley and Badecker (2009a) showed that learners of a vowel harmony language do not treat stems with identical vowels differently from stems that contain different vowels, suggesting that identical vowels in a

³Note that the morphological terminology ‘stem’ and ‘suffix’ are used throughout the description of the experiments to conform to the fact that vowel harmony processes are typically instantiated as morphophonological alternations of affixes. No explicit mention of morphology or semantics was provided to participants, nor were participants tested for morphological awareness.

training set will not lead learners to infer an identity relation between the stem and the suffix. Note that the neutral vowels never appeared in words of the form [CɛCɛC], an identical vowel context.

Training stimuli: control condition—The Control condition served as a baseline for response selections at test. The exposure phase was based on the control condition in Finley & Badecker 2009a. Participants were exposed to stem items, but not affixed forms. The forty-two stem items contained all twenty-four stems that appeared in the critical conditions, plus sixteen additional disharmonic stems that contained only nonneutral vowels. Because stems containing the neutral vowel were already evenly split between harmonic and disharmonic items, no new disharmonic items of this type were created. In the final set of items, half of the stems were harmonic and half of the stems were disharmonic. A full list of the Control training items can be found in the appendices.

Participants in the Control condition were exposed to the same test items as participants in the Transparent and Opaque conditions (described below). While all items were technically new, since no participant in the Control condition heard a suffixed form, the stimuli were matched to the critical conditions. The ‘correct’ items in the old stem and new harmonic stem test items were identical for both critical conditions. Finley and Badecker (2009a, 2012) showed that the present method for controls (harmonic and disharmonic stem items) was not significantly different from a ‘no training’ control condition in which participants were given test items only, without an exposure phase. The present control condition allows for identical instructions to be given to critical and control participants, keeping the symmetry between control and critical conditions as close as possible.

Test stimuli: all conditions—Following training, participants in all conditions were given a two-alternative forced-choice test to assess learning of the vowel harmony pattern. One item followed the harmony pattern, and the other item was disharmonic. Both choices were identical except for the final vowel, which varied between [e] and [o] (e.g. [kukopo, *kukope]). Participants were asked to choose which item was most likely to belong to the language that they were trained on. There were three different test conditions with ten items in each.⁴ OLD STEM items were taken directly from the training set: four containing the neutral vowel, six containing harmonizing vowels. Table 4 contains examples of test stimuli.

NEW HARMONIC STEM items tested for knowledge of the harmony rule without the neutral vowel, and NEW DISHARMONIC STEM items tested for knowledge of the harmony rule when the neutral vowel was present in a disharmonic context. Stems in new disharmonic stem items contained a back vowel followed by the neutral vowel. If the participant chose the front [e] vowel suffix, the neutral vowel was treated as opaque; if the participant chose the back [o] vowel suffix, the neutral vowel was treated as transparent.

Procedure—Participants were told that they would be listening to words from a language they had never heard before and that their task was to listen to the way the novel language

⁴Due to a programming error, some participants only responded to nine new disharmonic stem items. In addition, there were ten filler items that were not analyzed.

sounded, but that they need not try to memorize the forms. The training consisted of five repetitions of the twenty-four items in the Transparent and Opaque conditions and the forty-two items in the Control condition. The exposure phase was immediately followed by the two-alternative forced-choice test described above.

2.2. Results

The proportion of correct responses was recorded for each participant, and these proportions are presented in Figure 1 in terms of means and standard errors. A correct response was indicated by a harmonic response for new harmonic stem items and old items containing the harmonizing vowels. A correct response for new disharmonic stem items (and old disharmonic stem items⁵) in the Transparent condition was a response conforming to the initial vowel of the stem, while a correct response in the Opaque condition was a front vowel [e], since the neutral vowel was front. Correct responses in the Control varied depending on the comparison, but are depicted in Fig. 1 to match the Transparent condition. If learners are able to understand the role of the neutral vowel in a disharmonic context, the proportion of correct responses to new disharmonic stem items will be significantly greater than chance as compared to the Control condition. If learners are biased toward first-order locality, participants will successfully learn the behavior of the neutral vowel in disharmonic contexts in the Opaque condition but not the Transparent condition.

A generalized linear mixed model fit by the Laplace approximation was performed using the lme4 (Bates et al. 2013) and language (Baayen 2008) packages in R (R Development Core Team 2011). A single model was created with random intercepts for both items and subjects. Random slopes were included where appropriate for items (over training condition) and subjects (over test condition). The Opaque and Transparent conditions were each compared to the baseline (Control) condition with contrasts comparing new harmonic and new disharmonic stem items to the baseline (old stem) items. Interactions were performed between both contrast comparisons and subject comparisons (e.g. an interaction between the Control and Opaque comparisons with old stem and new harmonic stem comparisons). The summary of results can be found in Tables 5a, b below. Due to space constraints, only significant results are reported in the text. There was a significant difference between both the Opaque and Control conditions (0.69 vs. 0.48, $\beta = 0.90$, $z = 2.70$, $p = 0.0070$) and the Transparent and Control conditions (0.65 vs. 0.48, $\beta = 1.31$, $z = 3.80$, $p < 0.001$). There was also a significant interaction between the Transparent and Control conditions for new disharmonic stem vs. old stem items ($\beta = -1.41$, $z = -2.48$, $p = 0.013$). These results demonstrate that participants in both the Transparent and Opaque conditions learned the overall harmony pattern, but that in the Transparent condition, responses to new disharmonic stem items were significantly less likely to be correct than responses to old stem items.

A second model was run to compare responses to new disharmonic stem items between the critical (Opaque and Transparent) conditions and the baseline (Control) condition, with

⁵Old and new harmonic items were analyzed with items that contained the neutral vowel. Because there were so few items containing the neutral vowel in these conditions, it was not possible to separate the neutral vowel items to perform meaningful inferential statistics.

random intercepts for both subjects and items and random slopes for subjects. There was a significant difference between Opaque and Control for new disharmonic stem items (0.71 vs. 0.51, $\beta = 1.01$, $z = 2.32$, $p = 0.020$). There was no significant difference between Transparent and Control for new disharmonic stem items (0.49 vs. 0.51, $\beta = 0.10$, $z = 0.25$, $p = 0.80$). These results suggest that participants in the Opaque condition successfully learned the behavior of the neutral vowel in a disharmonic context, but participants in the Transparent condition failed to do so.

A third model was run to compare responses to new disharmonic stem items between the Opaque and Transparent conditions, which showed a significant difference (0.71 vs. 0.49, $\beta = 1.15$, $z = 2.44$, $p = 0.014$). These results suggest that participants in the Opaque condition were more likely to respond correctly to new disharmonic stem items than participants in the Transparent condition, supporting the hypothesis that first-order patterns are easier to learn than second-order patterns.

2.3. Discussion

Participants in both critical conditions of experiment 1 learned the overall vowel harmony pattern, but only participants in the Opaque condition successfully learned the behavior of the neutral vowel in a disharmonic stem. This suggests that adult native English speakers are biased toward opaque neutral vowels, in the sense that learners are more likely to detect the behavior of the neutral vowel in an opaque context than in a transparent context. Further, it suggests that local (first-order) harmony constraints may be the default constraints postulated by learners.

It is possible that the failure to learn the behavior of the transparent neutral vowel in experiment 1 was due to the ambiguous distribution of the suffix vowel following the neutral vowel. The suffix vowel was always [e] following the opaque vowel, but varied between [e] and [o] following the transparent vowel, depending on whether the initial vowel was [e] or [o]. This difference between the Opaque and Transparent conditions may have increased the difficulty of learning the behavior of the transparent neutral vowel.

3. Experiment 2

Experiment 2 tests the hypothesis that the failure of generalization to novel items containing transparent items in experiment 1 was due to the fact that the vowel that followed the opaque vowel was consistently [e], but the vowel that followed the transparent vowel was inconsistent ([e] or [o], depending on the initial vowel). In experiment 2, this inconsistency was removed by including only disharmonic stems (stems with a back initial vowel) when the neutral vowel [ɛ] was present. If learners use general association mechanisms between the neutral vowel and the suffix vowel, then learners should be able to learn the transparent neutral vowel when the suffix vowel is always a back vowel (as opposed to experiment 1, where the suffix vowel alternated based on the stem vowel). However, if learners use more than simple association mechanisms (akin to creating constraints such as those presented in Table 1), only displaying the neutral vowel in a back context will not affect performance, and learners in experiment 2 will perform similarly to participants in the Transparent condition of experiment 1.

3.1. Method

Participants—Twenty adult monolingual native speakers of English were recruited for this study at the University of Rochester and were paid \$10 for their participation. All participants had normal hearing and had no previous exposure to a vowel harmony language, including participation in a vowel harmony learning experiment.

Design—The design of the stimuli was identical to the Transparent condition of experiment 1, except that the four harmonic instances of the transparent vowel [ɛ] (e.g. [kenɛpe]) were replaced with disharmonic instances ([konɛp-konɛpo, kunɛm-kunɛmo, tomen-tomɛno, tukɛd-tukɛdo]), so the suffix vowel was always back following a neutral vowel. This created consistency in the suffix vowel when the neutral vowel was present.

Participants in experiment 2 were compared to the same participants in the Control condition as in experiment 1.

Procedure—The procedure was identical to that of experiment 1.

3.2. Results

Proportions of harmonic, transparent responses were recorded for each participant. The mean and standard errors were presented in Fig. 1 above.

A generalized linear mixed model fit by the Laplace approximation was performed using the lme4 package in R. A single model was created with random intercepts for both items and subjects. Random slopes were included where appropriate for items (over training condition) and subjects (over test condition). The Transparent condition was compared to the baseline (Control) condition with contrasts comparing new and new disharmonic stem items to the baseline (old stem) items. Interactions were performed between both contrast comparisons and the between-subjects comparison. The summary of results can be found in Tables 5a, b above. Due to space constraints, only significant results are reported in the text. There was a marginally significant difference between the Transparent and Control conditions (0.56 vs. 0.48, $\beta = 0.51$, $z = 1.67$, $p = 0.095$), suggesting that the participants in experiment 2 learned the overall harmony pattern, but with only marginal reliability.

A second model was run to compare responses to new disharmonic stem items between the critical (Transparent) and the baseline (Control) conditions, with random intercepts for both subjects and items and random slopes for subjects. There was no significant difference between the Transparent and Control conditions for new disharmonic stem items (0.55 vs. 0.51, $\beta = 0.14$, $z = 0.34$, $p = 0.74$). These results suggest that participants in the Transparent condition failed to learn the behavior of the neutral vowel in a disharmonic context.

A third model was run to compare responses between the Transparent conditions in experiment 1 and experiment 2 with random intercepts for both subjects and items and random slopes for subjects. There was a marginally significant difference between experiments 1 and 2 (0.65 vs. 0.55, $\beta = 0.57$, $z = 1.72$, $p = 0.085$). There was also a marginally significant interaction between the difference between new disharmonic stem and old stem items and the difference between experiments 1 and 2 ($\beta = -0.85$, $z = -1.71$, p

= 0.088). These results suggest that overall learning decreased in the Transparent condition from experiment 1 to experiment 2.

A fourth model was run to compare responses to new disharmonic stem items between the Transparent conditions from experiments 1 and 2, with random intercepts for both subjects and items and random slopes for subjects. This model showed no significant differences between experiments 1 and 2 for new disharmonic stem items (0.49 vs. 0.51, $\beta = -0.28$, $z = -0.93$, $p = 0.35$). These results suggest that simply increasing the number of neutral items in a disharmonic context does not increase performance on novel transparent items. In addition, increasing the number of neutral items in a disharmonic context may actually impede learning the general harmony pattern.

3.3. Discussion

The results of experiment 2 demonstrate that the difference between the Transparent and Opaque conditions in experiment 1 was not due to ambiguous training items in the Transparent condition. If learners in the Opaque condition in experiment 1 were simply associating the neutral vowel with the suffix [e], then participants in experiment 2 should have performed as well as those in the Opaque condition in experiment 1 when all items containing the neutral vowel ended in the [o] suffix. Further, if it were simply the proportion of items that fit a particular suffix, one should have expected higher performance in experiment 2, which was not shown.

Interestingly, overall performance in experiment 2 decreased, as experiment 2 was only marginally different from the Control condition, and performance in the Transparent condition marginally decreased from experiment 1 to experiment 2. It is possible that removing the harmonic items containing the neutral vowel created a conflict of cues for harmony. The majority of the items in experiment 1 involved harmonic items, but the neutral items were always disharmonic. This conflicting cue provides insight into the challenge of learning a nonadjacent harmony dependency. While both opaque and transparent vowels can result in disharmony, transparent vowels are more likely to result in increased disharmony. This increase in overall disharmony may make learning vowel harmony more difficult overall (Pycha et al. 2003).

While the present experiments demonstrated a bias toward first-order nonadjacent harmony patterns over second-order nonadjacent harmony patterns, the question remains as to what is necessary for learning transparent vowels. Given that transparent vowels can be found in a variety of languages, one should expect such languages to be learnable. The learnability of transparent vowels may be affected by the degree of coarticulation shared between the trigger and target vowels. Benus and Gafos (2007) showed that transparent vowels in Hungarian are pronounced differently depending on the harmonic contexts; in back (disharmonic) contexts, the transparent vowels are pronounced more back than in front (harmonic) contexts. In addition, Gordon (1999) measured coarticulation between trigger and target vowels by comparing the F2 values of neutral vowels in harmonic contexts vs. disharmonic contexts. If the F2 value of a neutral vowel increases in a harmonic context compared to a disharmonic context, it suggests that the backness of the neutral vowel is affected by the harmonic context, and therefore coarticulation. Gordon (1999) showed that

F2 values for neutral vowels in Finnish differed significantly depending on the harmonic context, but only in the direction of harmony. These results suggest the possibility that the failure to learn the transparent neutral vowels was due to insufficient coarticulatory cues in the Transparent conditions.

One way of increasing the coarticulatory cues for the transparent vowel is to change the height of the neutral vowel. In experiments 1 and 2, the transparent neutral vowel was mid. Mid vowels are typologically less likely to behave as a transparent vowel for back vowel harmony than high vowels (Anderson 1980, Benus & Gafos 2007, Kimper & Ylitalo 2012), which may be due to differences in coarticulation (Benus & Gafos 2007), as well as compensation for coarticulation (Ohala 1994a,b). Experiment 3 addresses the possibility that learners in experiments 1 and 2 were biased against [ɛ] as a neutral vowel by increasing the coarticulation of the neutral vowel in disharmonic contexts through the use of a high front vowel ([ɪ]) as the neutral vowel.

4. Experiment 3

Experiments 3–5 explore the factors that increase learnability of the transparent neutral vowel. If the failure to learn the transparent behavior of the neutral vowel is due to a bias against [ɛ] as a transparent vowel, replacing [ɛ] with [ɪ] should induce learning of the behavior of the transparent vowel.

4.1. Method

Participants—Forty adult monolingual native speakers of English were recruited for this study at the University of Rochester and were paid \$10 cash for participation. All participants had normal hearing and had no previous exposure to a vowel harmony language, including participation in a vowel harmony learning experiment.

Design—The design of the stimuli was identical to that of experiment 1, with the following changes. All instances of the neutral vowel were changed from the mid front vowel [ɛ] to the high front lax vowel [ɪ]. All stimulus tokens were recorded by an adult female native English speaker (different from the speaker in experiment 1, but still naive to the purpose of the experiment). The training and test items were parallel to those in experiments 1 and 2, except that, due to experimenter error, only one disharmonic stem item contained a nonneutral vowel ([tukede-tukedo]) and was thrown out of the analyses. Final analyses included nine new disharmonic stem items. Lists of stimuli for experiment 3 can be found in Appendix B.

To test whether the change of the transparent vowel from [ɛ] to [ɪ] resulted in increased coarticulation, the F2 of all instances of [ɛ] and [ɪ] in both harmonic and disharmonic contexts was measured. If the decrease in F2 from a harmonic (front) context to a disharmonic (back) context is greater for [ɪ] than for [ɛ], it suggests that [ɪ] undergoes greater coarticulation than [ɛ] as a transparent vowel; this will be found if there is a significant interaction between vowel height and vowel backness in an ANOVA comparing F2 measurements for high and mid neutral vowels in both front and back contexts. To assess the level of coarticulation, the F2 of the transparent vowel was measured for all new

disharmonic stem items in both the harmonic and disharmonic contexts from both experiments 1 and 2, and experiment 3.

A 2×2 (Backness $\times$ Height) ANOVA was performed. We found a main effect of Backness on F2 (2,128 vs. 2,037, $CI = 34$; $F(1,18) = 18.68$, $p < 0.001$), a main effect of Height (2,339 vs. 1,826, $CI = 54$; $F(1,18) = 209.30$, $p < 0.001$), and a significant interaction ($F(1,18) = 7.77$, $p < 0.05$). This is due to the fact that the mean differences between F2 for high vowels is greater than the difference in F2 for mid vowels (132 vs. 40, $CI = 39$; $t(18) = 2.79$, $p < 0.05$). This suggests that using a high transparent vowel resulted in an increase in coarticulation in transparent contexts.

The design of the Control condition was identical to that of experiment 1 (except for the change in the neutral vowel and the voice used in the stimuli). Participants were exposed only to stems (both harmonic and disharmonic) and were tested on the same items as participants in the critical condition.

Procedure—The procedure was identical to that of experiments 1–2.

4.2. Results

Proportions of correct responses (in the experimental conditions, a correct response was a harmonic response for items containing harmonizing vowels, and a transparent response for items containing the neutral vowel) were recorded for each participant for each of the training conditions. The means and standard errors for experiments 3–5 are presented in Figure 2.

A generalized linear mixed model fit by the Laplace approximation was performed using the lme4 package in R. A single model was created with random intercepts for both items and subjects. Random slopes were included only for items (over test condition), since the model that included random slopes for both subjects and items failed to converge. The Transparent condition was compared to the baseline (Control) condition with contrasts comparing new harmonic and new disharmonic stem items to the baseline (old stem) items. Interactions were performed between both contrast comparisons and the between-subjects comparison. The summary of results can be found in Table 6a, b. Due to space constraints, only significant results are reported in the text. There was a significant difference between the Transparent and Control conditions ($\beta = 0.81$, $z = 3.12$, $p = 0.018$). There was also a significant interaction between the Transparent and Control conditions for new disharmonic stem vs. old stem items ($\beta = -0.62$, $z = -2.09$, $p = 0.037$). These results demonstrate that participants in the Transparent condition learned the overall harmony pattern, but that there was a relative decrease between old and new disharmonic stem items.

A second model was run to compare responses to new disharmonic stem items between the critical (Transparent) and baseline (Control) conditions, with random intercepts for both subjects and items and random slopes for items. There was no significant difference between the Transparent and Control conditions for new disharmonic stem items (0.59 vs. 0.47, $\beta = 0.20$, $z = 0.59$, $p = 0.56$). These results suggest that participants in the Transparent condition failed to learn the behavior of the transparent neutral vowel in a disharmonic context.

4.3. Discussion

The results of experiment 3 demonstrate that increasing the coarticulation on the neutral vowel did not have a significant effect on learners' ability to acquire the nature of transparent vowels in harmony, since a model comparing the new disharmonic stem items in the Transparent condition of experiment 3 to the Transparent condition of experiment 1, with random intercepts for both subjects and items and random slopes for items, was not significant (0.54 vs. 0.49, $\beta = 0.25$, $z = 0.56$, $p = 0.58$). Participants in experiment 3 were unable to generalize to novel items containing the neutral vowel in a disharmonic context, suggesting that increased coarticulation may not be a significant factor in adult learning of transparent vowels. It is possible that coarticulation was not increased 'enough' from experiment 1 to experiment 3. While future research might be able to discern how much coarticulation might be necessary to significantly increase learning, there is no way of knowing, from the present data, how much coarticulation is sufficient to induce learning.

Another possibility is that a greater amount of exposure is required to learn the vowel harmony patterns with transparent vowels (as compared to opaque vowels). We address this possibility in experiment 4 by increasing the amount of exposure to the neutral vowel in a disharmonic context. In experiment 4a, the exposure from experiments 1–3 was doubled, and in experiment 4b, exposure to the neutral vowel was further increased by expanding the proportion of items containing the neutral vowel from 1/6 to 1/3 (in addition to the increased exposure time). In experiment 4c, the frequency of disharmonic neutral transparent tokens was also increased from experiment 4a (in addition to the increased exposure time).

5. Experiment 4

In experiment 4, the role of increased exposure in learning the behavior of the transparent neutral vowel was tested. Experiment 4 was divided into three conditions, each with a different amount of exposure to the transparent vowel. Increasing the amount of exposure to the transparent vowel should result in both (i) a significant difference between the critical and control conditions for new disharmonic stem items and (ii) a significant increase in correct transparent responses to new disharmonic stem items from experiment 3.

5.1. Method

Participants—Sixty adult monolingual native speakers of English were recruited for this study at the University of Rochester and paid \$10 for their participation. All participants had normal hearing and had no previous exposure to a vowel harmony language, including participation in a vowel harmony learning experiment.

Design—Participants in experiment 4a heard the same exposure set as in experiment 3, for twice the number of repetitions (increased from five to ten). In experiment 4b, the number of disharmonic items containing the neutral vowel was increased from four to ten, thereby doubling the proportion of disharmonic stem items (from 4/24 (1/6) to 10/30 (1/3)). In this case, each item, whether it contained the neutral vowel or not, was heard ten times. In experiment 4c, the four disharmonic items containing the neutral vowel found in experiments 3 and 4a were played two times per iteration of the training set, for a total of

twenty times. This repetition was based on the hypothesis that if a few pertinent items were heard more frequently, then these items would be more salient and thus more likely to be remembered. These remembered items could be useful in generalizing to novel items of a similar type.

The test items in experiment 4 were identical to those in experiment 3. The Control condition in experiment 3 was used as the baseline in experiment 4.

Stimuli—The stimuli were identical to those used in experiment 3, except that in experiment 4b, six additional disharmonic items were added to the training set. These additional items were /motɪp-motɪpo, dugɪb-dugɪbo, topɪm-topɪmo, nukɪt-nukɪto, konɪk-konɪko, bumɪg-bumɪgo/.

Procedure—The procedure was identical to that of experiments 1–3.

5.2. Results

A generalized linear mixed model fit by the Laplace approximation was performed using the lme4 package in R. A single model was created with random intercepts for both items and subjects. Random slopes were included where appropriate for items (over training condition) and subjects (over test condition). The Transparent condition was compared to the baseline (Control) condition with contrasts comparing new harmonic and new disharmonic stem items to the baseline (old stem) items. Interactions were performed between both contrast comparisons and the between-subjects comparison. The summary of results can be found in Tables 6a, b above. Due to space constraints, only significant results and results that differed from experiment 3 are reported in the text.

Experiment 4a: doubled exposure from experiment 3—There was a significant difference between experiment 4a and the Control condition (0.72 vs. 0.46, $\beta = 1.89$, $z = 4.78$, $p < 0.001$). Unlike experiment 3, there was a significant interaction between experiment 4a and the Control condition for new disharmonic vs. old stem items ($\beta = -1.05$, $z = -2.84$, $p < 0.01$). These results demonstrate that participants in experiment 4a learned the overall harmony pattern, but that there was a decrease in performance in new disharmonic stem items compared to old stem items.

A second model was run to compare responses to new disharmonic stem items between experiment 4a and the baseline (Control) condition, with random intercepts for both subjects and items and random slopes for items. There was a marginally significant difference between experiment 4a and the Control condition for new disharmonic stem items (0.64 vs. 0.49, $\beta = 0.83$, $z = 1.91$, $p = 0.056$). Because a reliable difference was not obtained, it is not clear whether increased training reliably caused participants to learn the behavior of the neutral vowel in a disharmonic context.

A third model was run to compare responses to new disharmonic stem items between experiment 4a and experiment 3, with random intercepts for both subjects and items and random slopes for items. There was no significant difference between experiment 4a and experiment 3 for new disharmonic stem items (0.64 vs. 0.54, $\beta = 0.63$, $z = 1.25$, $p = 0.21$).

Because a reliable difference was not obtained, it is not clear whether doubling exposure created an effective increase in performance on new disharmonic stem items.

Experiment 4b: increased disharmonic stem items from experiment 4a—There was a significant difference between experiment 4b and the Control condition (0.66 vs. 0.46, $\beta = 1.04$, $z = 3.59$, $p < 0.001$). There was no significant interaction between experiment 4b and the Control condition for new disharmonic stem vs. old stem items ($\beta = -0.42$, $z = -1.05$, $p = 0.29$). These results demonstrate that participants in experiment 4b learned the overall harmony pattern.

A second model was run to compare responses to new disharmonic stem items between experiment 4b and the baseline (Control) condition, with random intercepts for both subjects and items and random slopes for items. There was a marginally significant difference between experiment 4b and the Control condition for new disharmonic stem items (0.63 vs. 0.49, $\beta = 0.65$, $z = 1.90$, $p = 0.058$). Because a reliable difference was not obtained, it is not clear whether increased training reliably caused participants to learn the behavior of the neutral vowel in a disharmonic context.

A third model was run to compare responses to new disharmonic stem items between experiment 4a and experiment 3, with random intercepts for both subjects and items and random slopes for items. There was no significant difference between experiments 3 and 4a for new disharmonic stem items (0.63 vs. 0.59, $\beta = 0.50$, $z = 0.84$, $p = 0.40$). Because a reliable difference was not obtained, it is not clear whether the increase in exposure in experiment 4b created an effective increase in performance on new disharmonic stem items.

Experiment 4c: doubled repetitions of disharmonic stem items from experiment 4a—The model for experiment 4c included only random slopes for subjects (over test condition), since the model that included random slopes for both subjects and items failed to converge. There was a significant difference between experiment 4c and the Control condition (0.81 vs. 0.46, $\beta = 2.17$, $z = 6.68$, $p < 0.001$). There was no significant interaction between experiment 4c and the Control condition for new disharmonic stem vs. old stem items ($\beta = -0.66$, $z = -1.36$, $p = 0.17$). These results demonstrate that participants in the Transparent condition learned the overall harmony pattern without a significant relative decrease in performance in new disharmonic stem items compared to old stem items.

A second model was run to compare responses to new disharmonic stem items between the critical (Transparent) and the baseline (Control) conditions, with random intercepts for both subjects and items and random slopes for items. There was a significant difference between the Transparent and Control conditions for new disharmonic stem items (0.72 vs. 0.49, $\beta = 1.38$, $z = 2.98$, $p = 0.0029$). These results suggest that the increase in exposure in experiment 4c was sufficient for learning the behavior of the transparent neutral vowel.

A third model was run to compare responses to new disharmonic stem items between experiment 3 and experiments 4a and 4b, with experiment 4c as the baseline. There was a significant difference between experiments 3 and 4c for new disharmonic stem items (0.54

vs. 0.72, $\beta = 1.18$, $z = 2.29$, $p = 0.022$), but not between experiments 4a and 4c (0.72 vs. 0.64, $\beta = 0.55$, $z = 1.04$, $p = 0.30$) or experiments 4b and 4c (0.72 vs. 0.63, $\beta = 0.68$, $z = 1.22$, $p = 0.22$). While experiment 4c was the only condition to show a reliable increase in correct responses to new disharmonic stem items from experiment 3, it was not different from the other increases in exposure in experiments 4a and 4b.

5.3. Discussion

Participants in experiment 4c reliably learned the role of the transparent vowel. This suggests that learning transparent vowels in a vowel harmony language is possible with sufficient exposure. While participants in experiments 4a and 4b numerically increased harmonic performance for new disharmonic stem items, there was only a marginally significant difference from the Control condition, and no significant increase from experiment 3. One explanation for this lack of statistical significance is that there is a large amount of variability in the responses to new disharmonic stem items, shown in Figure 3.

If we assume that selecting the transparent (correct) item in new disharmonic stem test items is theoretically equivalent to ranking the distal harmony constraint above the local harmony constraint, one can roughly infer the proportion of participants relying on the distal harmony constraint. Any participant selecting the transparent item 100% of the time (or on all nine trials) can be considered to be relying primarily (if not exclusively) on the distal harmony constraint. This number jumps from three participants in experiment 3 to nine participants in experiment 4c. The range of responses in experiments 4a and 4b is more variable, suggesting that participant variability and strong individual differences resulted in a failure to reliably increase transparent responses in experiments 4a and 4b.

Overall, it appears that increasing the number of repetitions of the same set of transparent items (experiment 4c) created the highest level of harmonic responses to new disharmonic stem test items. Because there were no significant differences between the three experiments, however, it is unclear whether this numeric difference is due to random variation in participants or whether a high proportion of the same transparent item systematically produces less overall variation. This is a question for future research.

The neutral vowel in experiments 3–4 was [ɪ], rather than the [ɛ] used in experiments 1–2. As noted above, the neutral vowel was changed from mid to high because high vowels are typologically more likely to be transparent for back harmony than mid vowels. In addition, the change in height also created an increase in coarticulation, which may aid learning. While simply changing the quality of the neutral vowel in experiment 3 did not yield reliable learning for new disharmonic stem items, the possibility remains that the combination of a change in the vowel along with increased training (in experiment 4c) yielded successful learning of the behavior of the neutral vowel. Experiment 5 tests the hypothesis that increased training yields successful learning of the transparent neutral vowel, even when it is mid.

6. Experiment 5

6.1. Method

Participants—Seventeen adult monolingual native speakers of English were recruited for this study at the University of Rochester. All participants had normal hearing and had no previous exposure to a vowel harmony language, including participation in a vowel harmony learning experiment. All participants were paid \$10 for their participation.

Design—The experimental design was identical to that of experiment 4c, except that the neutral vowel item was [ε]. Participants in experiment 5 were compared to the Control participants in experiment 1.

Stimuli—The stimuli were identical to those used in experiment 4c, except that the neutral vowel was changed from [ɪ] to [ε].

Procedure—The procedure was identical to that of experiments 1–4.

6.2. Results

A generalized linear mixed model fit by the Laplace approximation was performed using the lme4 package in R. A single model was created with random intercepts for both items and subjects. Random slopes were included where appropriate for items (over training condition) and subjects (over test condition). The Transparent condition was compared to the baseline (Control) condition with contrasts comparing new harmonic and new disharmonic stem items to the baseline (old stem) items. Interactions were performed between both contrast comparisons and the between-subjects comparison. The summary of results can be found in Table 7 below. Due to space constraints, only significant results are reported in the text. There was a significant difference between the Transparent and the Control conditions (0.83 vs. 0.48, $\beta = 2.77$, $z = 6.12$, $p < 0.001$). There was also a marginally significant interaction between new disharmonic stem vs. old stem items and the training conditions ($\beta = -0.98$, $z = -1.73$, $p = 0.084$). These results demonstrate that participants in the Transparent condition learned the overall harmony pattern without a significant relative decrease in performance in new disharmonic stem items compared to old stem items.

A second model was run to compare responses to new disharmonic stem items between the critical (Transparent) and the baseline (Control) conditions, with random intercepts for both subjects and items and random slopes for items. There was a significant difference between the Transparent and Control conditions for new disharmonic stem items (0.76 vs. 0.51, $\beta = 1.57$, $z = 2.51$, $p = 0.012$). These results suggest that learning the behavior of a neutral vowel in a disharmonic context is not dependent on the neutral vowel.

A third model was run to compare responses to new disharmonic stem items between the Transparent conditions in experiments 5 and 1, with random intercepts for both subjects and items and random slopes for items. There was a significant difference between the Transparent conditions in experiment 1 vs. experiment 5 for new disharmonic stem items (0.76 vs. 0.49, $\beta = 1.86$, $z = 2.96$, $p = 0.0030$). These results suggest that there was a

significant improvement on learning the behavior of the transparent vowel from experiment 1 to experiment 5.

A fourth model was run to compare responses to new disharmonic stem items between the Transparent conditions in experiments 5 and 4c, with random intercepts for both subjects and items and random slopes for items. There was no significant difference between the Transparent conditions in experiment 4c vs. experiment 5 for new disharmonic stem items (0.76 vs. 0.72, $\beta = 0.18$, $z = 0.21$, $p = 0.84$). These results suggest that there was no significant improvement on learning the behavior of the transparent vowel from experiment 4c to experiment 5.

6.3. Discussion

The results of experiment 5 mirror the results of experiment 4c: participants were able to learn the behavior of the transparent neutral vowel with increased training. This suggests that the failure to learn the behavior of the transparent vowel was not due to the vowel quality of the transparent vowel.

7. General discussion

The present study included five experiments addressing the learnability of transparent vowels in vowel harmony. Experiment 1 compared learning of two types of neutral vowels: opaque (first-order nonadjacent) vowels and transparent (second-order nonadjacent) vowels. The results from this experiment indicated that learners were able to acquire the opaque but not the transparent neutral vowel pattern. Experiment 2 demonstrated that the failure to learn the behavior of the transparent neutral vowel was not due to inconsistency between final (affix) vowel and the neutral vowel, which is front when the transparent vowel is in a harmonic stem, and back when the transparent vowel is in a disharmonic stem. Experiment 3 demonstrated that the failure to learn the behavior of the transparent neutral vowel was not simply due to its vowel quality; raising it from [ɛ] to [ɪ] did not induce learning of the transparent neutral vowel. Experiments 4 and 5 showed that it is possible for learners to internalize the transparent harmony rule given sufficient training. While simply doubling the overall exposure time was not sufficient (experiment 4a), doubling the exposure to disharmonic transparent vowels (in addition to doubling overall exposure) produced reliable improvement from experiment 3 (experiment 4c).⁶

It is possible that the reason that experiments 4c and 5 showed reliable learning of the transparent items was due not just to increased training, but also to the fact that the training set included a higher proportion of disharmonic stem items compared to experiments 1–4b. While more research is needed to fully understand the relative role of the distribution of the training set in learning, it supports a view in which lexical statistics play a role in the behavior of phonological patterns (Bybee 2002, Pierrehumbert 2001), and a combination of type and token frequency is required to learn novel phonological patterns (Richtsmeier et al.

⁶It is unlikely that the increased learning of transparent vowels found in experiment 4c is solely due to increasing the frequency of transparent vowel trials. In a pilot experiment with eight participants, the frequency of transparent items was doubled, but the total number of iterations was kept to five (as in experiments 1 and 2). The results showed a level of variability similar to experiment 4a; of the eight participants, only three were above 50% chance; three participants were below 50%.

2011). A goal for future research is to test the role of lexical statistics in the relative learnability of linguistic patterns.

The present study demonstrated that with sufficient training, participants can learn the behavior of a neutral vowel in a disharmonic context. One potential issue is the possibility that rather than learning that the neutral vowel is transparent, participants simply learned that a back vowel always followed the transparent vowel. Because the new disharmonic stem items contained only back vowels in the stems, establishing learning at test was contingent upon recognizing the behavior of the neutral vowel in the presence of a back vowel stem. In order to test this hypothesis, new harmonic stem items that contained a neutral vowel were analyzed separately. Because there were few such items, only descriptive information could be obtained. If participants did not learn the behavior of the neutral vowel when in a front vowel stem, one would expect that correct responses to old stem and new harmonic stem items that contained a neutral vowel would be below chance (50%). Table 8 lists the proportion of correct (harmonic, front vowel) responses when the old stem and new harmonic stem items contained a neutral vowel. Participants in experiments 1 and 4(a–c) showed average responses above 60% for both old stem and new harmonic stem items. In experiment 3, participants showed a 48% correct response rate for new harmonic stem items. One possibility for this difference is that, rather than learning that the quality of the suffix following the neutral vowel depends on the initial vowel, participants in experiment 3 learned that the quality of the suffix vowel following the neutral vowel is ‘random’. In many languages with transparent neutral vowels (e.g. Hungarian), the neutral vowel will usually select a suffix that agrees with it, but will sometimes select a disharmonic suffix (Hayes & Londe 2006). It is possible that participants in experiment 3 inferred a harmony pattern in which the neutral vowel selects both front vowel suffixes.

While the learnability of nonadjacent dependencies in language has received a large amount of attention in the literature, this attention has largely been focused on phrase-level dependencies and word segmentation (Gómez 2002, Gómez & Maye 2005, Misyak & Christiansen 2007, Misyak et al. 2009, Newport & Aslin 2004). Important insights can be gained into how nonadjacent patterns are learned when we expand the repertoire of study to nonadjacent dependencies in phonological processes. The present results support a view that learning second-order nonadjacent dependencies is constrained compared to first-order nonadjacent patterns, but this learning is possible given appropriate exposure to relevant training items.

We have shown that learners are biased against transparent vowels, which require second-order nonadjacent dependencies and therefore more complex linguistic structure. This finding supports models that require additional structure or rule ordering in order to account for transparent vowels (Bakovi & Wilson 2000, Finley 2008, Goldsmith 1985, Hayes & Londe 2006, Kiparsky & Pajusalu 2003, Ringen & Heinamaki 1999). One might expect more complex patterns to be typologically less frequent. In a preliminary survey of fifty-one languages across thirteen language families, Rhodes (2010) notes that there are more vowel harmony systems with opaque vowels than with transparent vowels. There were roughly twenty-five harmony systems with opaque vowels (from eight language families) and twelve harmony systems with transparent vowels (from five language families), including four with

both transparent and opaque vowels. The high overlap in language families makes it difficult to discern the stability of the differences found in the survey. Even if transparent vowels are not as frequent as opaque vowels, they are typologically robust; transparent vowels can be found across a wide range of language families (e.g. Indo-European, Altaic, Niger-Congo). While typologically infrequent patterns may be more difficult to learn, it is not necessarily the case that learning difficulties will directly translate into a typological difference (Rafferty et al. 2013). This suggests that learning may play a role in the relative typological difference between transparent and opaque vowels, but additional factors may keep the existence of transparent vowels stable crosslinguistically.

The present article is part of a larger research program that incorporates experimental results into generative models of phonological processes, and is working to develop an interdisciplinary, cognitively structured model of language (Smolensky & Legendre 2006). While the research presented here is only a small step toward this ultimate goal, these experiments provide a rich foundation for future work to pursue a model of language that incorporates both formal representations and cognitive processes.

Appendix A

Experiments 1 and 2 stimuli

TRAINING ITEMS.

Harmonic stems

budok-budoko, degib-degibe, dupob-dupobo, gemit-gemite, gitek-giteke, kimet-kimete, kukop-kukopo, mekin-mekine, midik-midike, monuk-monuko, netep-netepe, nopub-nopubo, puduk-puduko, tikep-tikepe, todup-todupo, tokot-tokoto

Harmonic stems with neutral vowel (exp. 1)

dimek-dimeke, mepen-mepene, tedet-tedete, piken-pikene

Opaque stems with neutral vowel

doteb-dotebe, pokeg-pokege, gupek-gupeke, bupet-bupete

Transparent stems with neutral vowel

doteb-dotebo, pokeg-pokego, gupek-gupeko, bupet-bupeto

Additional transparent stems with neutral vowel (exp. 2)

konep-konepo, kunem-kunemo, tomen-tomeno, tuked-tukedo

Control stems

gitek, gitok, gupek, kimet, kimot, kukep, kukop, mekin, mepen, midik, mikot, modeb, monik, monuk, netep, nopib, nopub, nuged, piken, piton, pokeg, pudeg, puduk, tedet, tikep, tikop, todup, tomen, tuked, tokot

TEST ITEMS (incorrect/disharmonic items are starred (*); correct opaque items are underlined; correct transparent items are in **bold**).

Old stems

bupete-**bupeto**, gemite-*gemito, giteke-*giteko, kukopo-*kukope, mepene-*mepeno,
monuko-*monuke, nopubo-*nopube, pikene-*pikeno, pokege-**pokego**, tikepe-*tikepo

New harmonic stems

bedite-*bedito, kenete-*keneto, godomo-*godome, nupuko-*nupuke, pedebe-*pedebo,
toguko-*toguke, butoko-*butoke, mukobo-*mukobe, tidipe-*tidipo, mitene-*miteno

New disharmonic stems

gobeke-**gobeke**, mokene-**mokeno**, nuteme-**nutemo**, kogeme-**kogemo**, kupege-**kupego**,
pomeke-**pomeko**, pukene-**pukeno**, todepe-**todepo**, tudepe-**tudepo**

Appendix B

Experiments 3–4 stimuli

Training items (items containing the neutral vowel in bold)

bupit-**bupito**, **dimik**-**dimike**, **dotib**-**dotibo**, **gupik**-**gupiko**, **mepim**-**mepime**, **pikm**-**pikme**,
pokig-**pokigo**, **tedi**-**tedite**, budok-budoko, degib-degibe, dupob-dupobo, gemit-gemite,
gitek-giteke, kimet-kimete, kukop-kukopo, mekin-mekine, midik-midike, monuk-monuko,
netep-netepe, nopub-nopubo, puduk-puduko, tikep-tikepe, todup-todupo, tokot-tokoto

Additional transparent stems with neutral vowel for experiment 4b

motip-motipo, dugib-dugibo, topim-topimo, nukit-nukito, konik-koniko, bumig-bumigo

Control stems

bidok, bipit, budok, bupit, degib, detib, dipob, dogib, dotib, dumik, dupob, gemit, gipik,
gitek, gomit, gupik, gitek, kikop, kimet, kukop, kumet, medik, mekin, menuk, mokin,
monuk, mopin, nepub, netep, nopub, notep, pekig, piduk, pikin, pokin, puduk, pugin, tedup,
tekot, tikep, tedit, todup, tokot, tukep

Test items (incorrect/disharmonic items are starred (*)).

Old stems

bupite-*bupito, gemite-*gemito, giteke-*giteko, kukopo-*kukope, mepine-*mepino,
monuko-*monuke, nopubo-*nopube, pikine-*pikino, pokigo-*pokige, tikepe-*tikepo

New harmonic stems

bedite-*bedito, gudomo-*gudome, kenite-*kenito, nupuko-*nupuke, pedebe-*pedebo,
toguko-*toguke, butoko-*butoke, mukobo-*mukobe, tidipe-*tidipo, mitne-*mitno

New disharmonic stems

*gobike-gobiko, *mokine-mokino, *nutime-nutimo, *kogime-kogimo, *kupige-kupigo,
*pomike-pomiko, *pukine-pukino, *todipe-todipo, *tudipe-tudipo

REFERENCES

- Anderson, Lloyd B. Using asymmetrical and gradient data in the study of vowel harmony. In: Vago, Robert M., editor. *Issues in vowel harmony*. Amsterdam: John Benjamins; 1980. p. 271-340.
- Baayen, R. Harald *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press; 2008.
- Bakovi, Eric. *Harmony, dominance and control*. New Brunswick, NJ: Rutgers University dissertation; 2000.
- Bakovi, Eric; Wilson, Colin. Transparency, strict locality and targeted constraints. *West Coast Conference on Formal Linguistics (WCCFL)*. 2000; 19:43–56.
- Bates, Douglas; Maechler, Martin; Bolker, Ben. lme4: Linear mixed-effects models using Eigen and Eigen. R package version 0.999999-2.999999. 2013. Online: <http://cran.r-project.org/web/packages/lme4/index.html>
- Benus, Stephan; Gafos, Adamantios I. Articulatory characteristics of Hungarian ‘transparent’ vowels. *Journal of Phonetics*. 2007; 35:271–300. [PubMed: 18389086]
- Bybee, Joan. Word frequency and the context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change*. 2002; 14:261–90.
- Chomsky, Noam. *Lectures on government and binding*. Dordrecht: Foris; 1981.
- Clements, George N. The autosegmental treatment of vowel harmony. *Phonologica*. 1976; 19:111–119.
- Clements, George N.; Sezer, Engin. Vowel and consonant disharmony in Turkish. In: van der Hulst, Harry; Smith, Norval, editors. *The structure of phonological representations*. Dordrecht: Foris; 1982. p. 213-255.
- Culicover, Peter W.; Wilkins, Wendy. *Locality in linguistic theory*. New York: Academic Press; 1984.
- Ettlinger, Marc; Finn, Amy; Kam, Carla Hudson. The effect of sonority on word segmentation: Evidence for a phonological universal. *Cognitive Science*. 2012; 36:655–673. [PubMed: 22050005]
- Finley, Sara. *Formal and cognitive restrictions on vowel harmony*. Baltimore, MD: Johns Hopkins University dissertation; 2008.
- Finley, Sara. The privileged status of locality in consonant harmony. *Journal of Memory and Language*. 2011; 65:74–83. [PubMed: 21686094]
- Finley, Sara; Badecker, William. Towards a substantively biased theory of learning. *Berkeley Linguistics Society*. 2007; 33:142–154.
- Finley, Sara; Badecker, William. Artificial grammar learning and feature-based generalization. *Journal of Memory and Language*. 2009a; 61:423–437.
- Finley, Sara; Badecker, William. Right-to-left biases for vowel harmony: Evidence from artificial grammar. *North East Linguistic Society (NELS)*. 2009b; 38:269–282.
- Finley, Sara; Badecker, William. Learning biases for vowel height harmony. *Journal of Cognitive Science*. 2012; 13:287–327.
- Goldsmith, John. Tone melodies and the autosegment; *Ohio State Working Papers in Linguistics (Proceedings of the 6th Conference on African Linguistics)*; 1975. p. 135-147.
- Goldsmith, John. Vowel harmony in Khalka Mongolian, Yaka, Finnish and Hungarian. *Phonology Yearbook*. 1985; 2:253–275.
- Gómez, Rebecca L. Variability and detection of invariant structure. *Psychological Science*. 2002; 13:431–436. [PubMed: 12219809]
- Gómez, Rebecca L.; Maye, Jessica. The developmental trajectory of nonadjacent dependency learning. *Infancy*. 2005; 7:183–206.

- Gordon, Matthew K. The 'neutral' vowels of Finnish: How neutral are they? *Linguistica Uralica*. 1999; 35:17–21.
- Hayes, Bruce; Londe, Zsuzsa. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology*. 2006; 23:59–104.
- Hayes, Bruce; Wilson, Colin. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*. 2008; 39:379–440.
- Hayes, Bruce; Zuraw, Kie; Siptár, Péter; Londe, Zsuzsa. Natural and unnatural constraints in Hungarian vowel harmony. *Language*. 2009; 85:822–863.
- Kimper, Wendell; Ylitalo, Riikka. Variability and trigger competition in Finnish disharmonic loanwords. *West Coast Conference on Formal Linguistics (WCCFL)*. 2012; 30:195–204. Online: <http://www.lingref.com/cpp/wccfl/30/paper2817.pdf>.
- Kiparsky, Paul; Pajusalu, Karl. Towards a typology of disharmony. *The Linguistic Review*. 2003; 20:217–241.
- Leiwo, Matti; Kulju, Pirjo; Aoyama, Katsura. The acquisition of Finnish vowel harmony. *SKY Journal of Linguistics*. 2006; 19:149–161.
- MacWhinney, Brian. *The acquisition of morphophonology*. Oxford: Blackwell; 1978.
- Misyak, Jennifer B.; Christiansen, Morten H. Extending statistical learning farther and further: Long-distance dependencies, and individual differences in statistical language learning; *Proceedings of the 29th annual conference of the Cognitive Science Society*; 2007. p. 1307-1312.
- Misyak, Jennifer B.; Christiansen, Morten H.; Tomblin, Bruce. Statistical learning of nonadjacent dependencies predicts on-line processing of long-distance dependencies in natural language; *Proceedings of the 31st annual conference of the Cognitive Science Society*; 2009. p. 177-182.
- Moreton, Elliott. Analytic bias and phonological typology. *Phonology*. 2008; 25:83–127.
- Newport, Elissa; Aslin, Richard N. Learning at a distance I: Statistical learning of non-adjacent dependencies. *Cognitive Psychology*. 2004; 48:127–162. [PubMed: 14732409]
- Ohala, John. Hierarchies of environments for sound variation; plus implications for 'neutral' vowels in vowel harmony. *Acta Linguistica Hafniensia*. 1994a; 27:371–382.
- Ohala, John. Towards a universal, phonetically-based theory of vowel harmony; Yokohama. *Proceedings of the 3rd International Conference on Spoken Language Processing (ICSLP 94)*; 1994b. p. 491-494.
- Onnis, Luca; Monaghan, Padraic; Richmond, Korin; Chater, Nick. Phonology impacts segmentation in online speech perception. *Journal of Memory and Language*. 2005; 53:225–237.
- Pierrehumbert, Janet. Exemplar dynamics: Word frequency, lenition and contrast. In: Bybee, Joan L.; Hopper, Paul J., editors. *Frequency effects and emergent grammar*. Amsterdam: John Benjamins; 2001. p. 137-157.
- Polgardi, Krisztina. Vowel harmony and disharmony in Turkish. *The Linguistic Review*. 1999; 16:187–204.
- Prince, Alan; Smolensky, Paul. *Optimality theory: Constraint interaction in generative grammar*. Oxford: Blackwell; 2004 [1993].
- Pycha, Anne; Nowak, Pawel; Shin, Eurie; Shosted, Ryan. Phonological rule-learning and its implications for a theory of vowel harmony. *West Coast Conference on Formal Linguistics (WCCFL)*. 2003; 22:101–113.
- R Development Core Team. *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing; 2011. Online: <http://www.r-project.org>
- Rafferty, Anna N.; Griffiths, Thomas L.; Ettlinger, Marc. Greater learn-ability is not sufficient to produce cultural universals. *Cognition*. 2013; 129:70–87. [PubMed: 23831566]
- Rhodes, Russell. A survey of vowel harmony with special focus on neutral segments, glides, and diphthongs. Berkeley: University of California, ms; 2010. Online: <http://linguistics.berkeley.edu/~russellrhodes/pdfs/Prospectus.pdf>
- Richtsmeier, Peter; Gerken, LouAnn; Ohala, Diane. Contributions of phonetic token variability and word-type frequency to phonological representations. *Journal of Child Language*. 2011; 38:951–978. [PubMed: 21126387]

- Ringen, Catherine O. A concrete analysis of Hungarian vowel harmony. In: Vago, Robert M., editor. *Issues in vowel harmony*. Amsterdam: John Benjamins; 1980. p. 134-154.
- Ringen, Catherine O.; Heinamaki, Orvokki. Variation in Finnish vowel harmony: An OT account. *Natural Language and Linguistic Theory*. 1999; 17:303–337.
- Saffran, Jenny R. The use of predictive dependencies in language learning. *Journal of Memory and Language*. 2001; 44:493–515.
- Slobin, Dan I. *The cross-linguistic study of language acquisition*. Mahwah, NJ: Lawrence Erlbaum; 1997.
- Smolensky, Paul. Optimality in phonology II: Harmonic completeness, local constraint conjunction, and feature-domain markedness. In: Smolensky, Paul; Legendre, Geraldine, editors. *The harmonic mind: From neural computation to optimality-theoretic grammar*. Vol. 2. Cambridge, MA: MIT Press; 2006. p. 27-160.
- Smolensky, Paul; Legendre, Geraldine. *The harmonic mind*, vol. 1: Cognitive architecture. Cambridge, MA: MIT Press; 2006.
- Vago, Robert M. *The sound pattern of Hungarian*. Washington, DC: Georgetown University Press; 1980.

Figure 1.

Experiment 1 and 2 results (means and standard errors). For new disharmonic stem items, a front vowel response is considered harmonic in the Opaque condition, and a back vowel response in the Transparent condition. The Control condition was matched to the conditions for statistical tests, but is shown matched for the Transparent condition.

Figure 2.

Experiments 3–5 results (means and standard errors). A back vowel response was considered correct for new disharmonic stem items.

Figure 3.
Experiments 3–4: individual variation for new disharmonic stem items.

Table 1

Local vs. distal harmony constraints.

EXAMPLE	OUTCOME	PHONETIC FORM	LOCAL CONSTRAINT: *[+Back] [-Back]	LOCAL CONSTRAINT: *[-Back] [+Back]	DISTAL CONSTRAINT: *[+Back]X [-Back]	DISTAL CONSTRAINT: *[-Back]X [+Back]
(1)	harmonic	pu ₊ to ₊ ko ₊				
	disharmonic	pu ₊ to ₊ ke ₋	*		**	
(2)	disharmonic	pe ₋ te ₋ ko ₊		*		**
	harmonic	pe ₋ te ₋ ke ₋				
(3)	transparent	pu ₊ te ₋ ko ₊	*	*	*	*
	opaque	pu ₊ te ₋ ke ₋	*		**	

Table 2

Summary of experimental manipulations, predictions, and results.

	TRAINING MANIPULATIONS	HYPOTHESIS	NEUTRAL VOWEL LEARNED?
EXP 1	Transparent [ɛ]; Opaque [ɛ]; 24 items 5×	Learn opaque neutral vowels but not transparent neutral vowels.	only opaque
EXP 2	Transparent [ɛ]; All stems with [ɛ] disharmonic	Learn neutral vowels with consistent suffix following neutral vowel.	no
EXP 3	Transparent [ɪ]; 24 items 5×	Learn transparent vowel if bias is against [ɛ].	no
EXP 4a	Transparent [ɪ]; 24 items 10×	Learn transparent vowel with double exposure of all training items.	no
EXP 4b	Transparent [ɪ]; 24 + 6 (additional neutral vowel) items 10×	Learn transparent vowel with double exposure of all training items; increased proportion of disharmonic neutral vowel items.	no
EXP 4c	Transparent [ɪ]; 20 items 10×; 4 (neutral vowel) items 20×	Learn transparent vowel with double training of disharmonic neutral vowel items.	yes
EXP 5	Replicate Exp 4c with transparent [ɛ]		

Table 3
Examples of training stimuli: experiment 1 (underlined vowels differ across conditions).

	STEM VOWELS	N	EXAMPLES STEM + SUFFIX	
			TRANSPARENT	OPAQUE
Harmonic stems: harmonizing vowels (i, e, o, u)	ee	1	nete <u>p</u> e	nete <u>p</u> e
	ei	3	gemit <u>e</u>	gemit <u>e</u>
	ii	1	midik <u>e</u>	midik <u>e</u>
	ie	3	gitek <u>e</u>	gitek <u>e</u>
	oo	1	tokot <u>o</u>	tokot <u>o</u>
	ou	3	monuk <u>o</u>	monuk <u>o</u>
Harmonic stems: neutral vowel (i, e, ε)	uu	1	puduk <u>o</u>	puduk <u>o</u>
	uo	3	kukop <u>o</u>	kukop <u>o</u>
	eε	2	tedet <u>e</u>	tedet <u>e</u>
	iε	2	dimek <u>e</u>	dimek <u>e</u>
	oε	2	dotε <u>b</u> o	dotε <u>b</u> o
	uε	2	gupek <u>o</u>	gupek <u>o</u>

Table 4

Examples of test stimuli, experiment 1.

OLD STEM (disharmonic items starred)	[gemite, *gemito] [kukopo, *kukope]
NEW HARMONIC STEM (disharmonic items starred)	[bedite, *bedito] [mukobo, *mukobe]
NEW DISHARMONIC STEM (correct opaque items are underlined; correct transparent items are in bold)	[mokene, mokeno] [nuteme, nutemo]

Table 5

a. Summary of mixed-effect models: experiments 1 and 2.

	EXP 1			EXP 2			EXP 1 vs.2		
	β	SE	z	β	SE	z	β	SE	z
Opaque	0.90	0.34	2.70**	—	—	—	—	—	—
Transp	1.31	0.34	3.80***	-0.510	0.31	1.67 [†]	0.57	0.33	1.72 [†]
New harmonic	-0.05	0.25	-0.20	-0.015	0.24	-0.062	-0.23	0.32	-0.72
New disharm	0.14	0.35	0.41	0.17	0.32	0.53	0.290	0.39	0.73
Transp × New harmonic	-0.18	0.45	-0.40	-0.210	0.36	-0.58	0.093	0.37	0.25
Transp × New disharm	-1.41	0.57	-2.48*	-0.360	0.46	-0.77	-0.85	0.50	-1.71 [†]
Opaque × New harmonic	0.17	0.44	0.38	—	—	—	—	—	—
Opaque × New disharm	0.22	0.57	0.39	—	—	—	—	—	—

b. Summary of mixed-effect models comparing new disharmonic items: experiments 1 and 2.

	EXP 1			EXP 2			EXP 1 vs.2		
	β	SE	z	β	SE	z	β	SE	z
Opaque vs. Control, New disharm	1.01	0.43	2.32*	—	—	—	—	—	—
Transp vs. Control, New disharm	0.10	0.41	0.25	0.14	0.40	0.34	—	—	—
Opaque vs. Transp, New disharm	1.15	0.47	2.44*	—	—	—	—	—	—
Exp1 vs.2, New disharm	—	—	—	—	—	—	-0.28	0.30	-0.93

[†] 0.05 < p < 0.10,* p < 0.05,** p < 0.01,*** p < 0.001

Table 6

a. Summary of mixed-effect models: experiments 3–4.

	EXP 3			EXP 4a			EXP 4b			EXP 4c		
	β	SE	z	β	SE	z	β	SE	z	β	SE	z
Transp	0.81	0.26	3.12**	1.89	0.39	4.78***	1.04	0.29	3.59***	2.17	0.32	6.68***
New harm	-0.32	0.20	-1.55	-0.34	0.21	-1.63	-0.31	0.22	-1.43	-0.32	0.22	-1.47
New disharm	0.039	0.21	0.19	-0.016	0.22	-0.74	0.040	0.26	0.16	0.040	0.30	0.13
Transp × New harm	-0.09	0.29	-0.31	-0.29	0.36	-0.80	-0.13	0.35	-0.37	-0.014	0.37	-0.037
Transp × New disharm	-0.62	0.30	-2.09*	-1.05	0.37	-2.84**	-0.42	0.40	-1.05	-0.66	0.48	-1.36

b. Summary of mixed-effect models comparing new disharmonic items: experiments 3–4.

	EXP 3			EXP 4a			EXP 4b			EXP 4c		
	β	SE	z	β	SE	z	β	SE	z	β	SE	z
Transp vs. Control, New disharm	0.20	0.34	0.59	0.83	0.43	1.91 [†]	0.65	0.34	1.90 [†]	1.38	0.47	2.98**
Exp 3 vs. 4, New disharm	—	—	—	0.63	0.51	1.25	0.50	0.59	0.84	1.18	0.51	2.29*
Exp 4c vs. 4a, 4b, New disharm	—	—	—	0.55	0.53	1.04	0.68	0.56	1.22	—	—	—

[†] 0.05 < *p* < 0.10,
* *p* < 0.05,
** *p* < 0.01,
*** *p* < 0.001)

Table 7

Summary of mixed-effect models: experiment 5.

		β	SE	z
OVERALL MODEL	Transp	2.77	0.42	6.12***
	New harm	0.073	0.30	0.24
	New disharm	0.37	0.37	1.00
	Transp $\times$ New harm	-0.28	0.47	-0.59
	Transp $\times$ New disharm	-0.98	0.56	-1.73 [†]
MODELS COMPARING NEW DISHARM ITEMS	Transp vs. Control, New disharm	1.57	0.62	2.51*
	Exp 1 vs. Exp 5, New disharm	1.86	0.63	2.96**
	Exp 4c vs. Exp 5, New disharm	0.18	0.89	0.21

[†] $0.05 < p < 0.10$,* $p < 0.05$,** $p < 0.01$,*** $p < 0.001$

Table 8

Proportion , correct' (front vowel suffix) responses for items containing a neutral stem vowel and front-vowel (harmonic) stems containing a neutral vowel.

	OPAQUE	TRANSF	EXP 1	EXP 3	EXP 4A	EXP 4B	EXP 4C	EXP 5
Old stem	0.88	0.94	0.68	0.85	0.95	0.68	0.91	0.59
New harmonic stem	0.91	0.88	0.48	0.80	0.80	0.68	0.59	