A BMRS analysis of Trojan (and transparent) vowels in Hungarian vowel harmony

Scott Nelson & Eric Baković*

Abstract. A 'Trojan' vowel in a vowel harmony system is a vowel that spreads an underlyingly specified value of the spreading feature even though the Trojan vowel itself redundantly surfaces with the opposite value. We propose an analysis of Trojan vowels with Boolean Monadic Recursive Schemes (BMRS; Chandlee & Jardine 2021), a computational formalism describing phonological maps using conditional `IF ... THEN ... ELSE ...` terms. We show how a BMRS analysis of Trojan vowel behavior in Hungarian follows straightforwardly when coupled with freely-specified input strings. Our analysis also accounts for the fact that redundantly-specified Trojan vowels in Hungarian are transparent in other contexts, and we briefly contrast this with the situation in Yoruba where Trojan vowels are otherwise opaque.

Keywords. vowel harmony; Hungarian; Yoruba; Trojan vowels; transparent vowels; opaque vowels; Boolean Monadic Recursive Schemes

1. Introduction. Vowel harmony can be succinctly defined as a phonological pattern by which all vowels in a domain (typically the word) agree with each other in terms of some feature. This agreement often results in regular morphophonological alternations, in which case the origin of the spreading feature value can be attributed to (vowels in) morphemes that do not alternate. These are often stems, as opposed to affixes, in which case affix vowels generally alternate to agree with the vowel(s) in the stems to which they are attached.

Full spreading is illustrated in (1), with [±back] harmony examples from Turkish (Clements & Sezer 1982). Stems in these examples are initial, and indicated with '√'. The stem vowel in (1a) is underlyingly specified as [+back] while the stem vowel in (1b) is underlyingly specified as [−back]. Following suffix vowels undergo spreading of [±back] from the stem and thus alternate depending on the stem to which the suffixes are attached. (Whether the suffix vowels are themselves underlyingly specified for the spreading feature is a question we will address below.)

(1) Full spreading of [±back] from stem to suffixes in Turkish

a. √a**tʃ** + dɯr + saj + dɯ
[+back]
'come + CAUS + COND + PAST'

b. √gel + dir + sej + di
[−back]
'open + CAUS + COND + PAST'

Like most phonological patterns, the situation is often if not always more complex and interesting than this oversimplified picture. For instance, some vowels in what would otherwise be expected to be alternating positions resist alternation, often for phonologically predictable reasons — e.g., they are redundantly specified for one value of the spreading feature given their value for some other feature(s). This alternation resistance may result in *opaque* vowels, which block the spread of one value of the spreading feature and spread their own redundant value further.

* We thank the audiences at a UC San Diego PhonCo meeting (11/10/2025), at the LSA Annual Meeting in New Orleans (1/11/2026), at a UIUC Phonetics and Phonology Forum meeting (3/5/2026), and at a UC San Diego LIGN 211 class meeting (3/11/2026); any remaining errors are ours. Authors: Scott Nelson, University of Illinois Urbana-Champaign (sjnelson@illinois.edu) & Eric Baković, University of California San Diego (ebakovic@ucsd.edu).

Vowel opacity is illustrated in (2) with [±round] harmony in Turkish (Clements & Sezer 1982), where [−high] vowels are redundantly [−round]. The value [−round] can thus spread to and through [−high] vowels, as shown in (2a), but [−high] vowels resist the spread of [+round] and spread their redundant [−round] value instead, as shown in (2b).

(2) Opacity of [−high] vowels to [+round] spreading in Turkish

Alternation resistance may instead result in *transparent* vowels, which appear to allow the opposite value of the spreading feature to 'skip over' them.[1] Vowel transparency is illustrated in (3) with [±back] harmony in Finnish (Kiparsky 1981), where [−low, −round] vowels are redundantly [−back]. The value [−back] can thus spread to and through these vowels, as in (2a); when they resist the spread of [+back], they do so by allowing [+back] to spread past them, as in (2b).

(3) Transparency of [−low, −round] vowels to [+back] spreading in Finnish

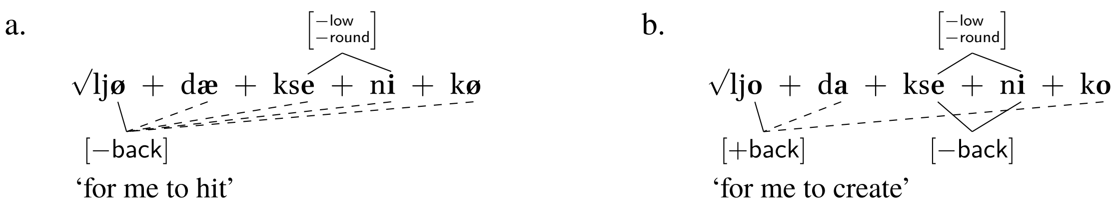

Alternation-resistant vowels are relatively commonplace, and a great deal of attention is paid to them in the literature. Somewhat less common are instances of stem vowels that appear to spread the opposite value of the spreading feature than the value that they themselves surface with. One analysis of such vowels, due originally to Vago (1973), is that they are *underlyingly* specified for the value of the feature that they spread, but that for independent reasons they redundantly surface with the opposite value themselves, much as opaque and transparent vowels do. Krämer (2003:34) calls these 'Trojan' vowels, and we adopt this evocative term here.

There are stems with Trojan vowels in the [±back] harmony pattern of Hungarian (Vago 1973, 1976; Ringen & Vago 1998; Siptár & Törkenczy 2000). A near-minimal contrast is shown in (4), between a stem [−low, −round] vowel that spreads its redundant [−back] value (4a) and a redundantly identical Trojan stem vowel that spreads its underlying [+back] value (4b).

(4) Spreading of [+back] from Trojan [−low, −round] vowels in Hungarian

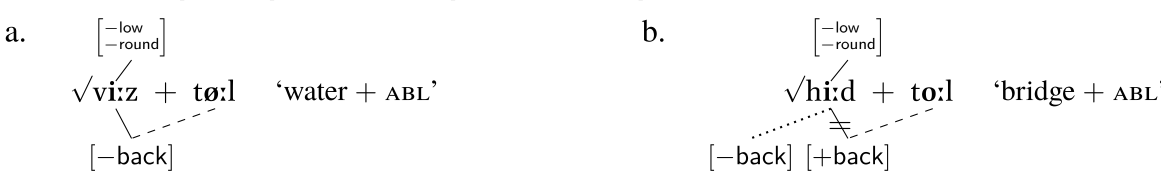

[1] Whether transparent vowels are actually specified with the redundant value of the spreading feature is debated in the literature; see the review of approaches to transparency in Baković (2025:§4). We assume here that they are.

We present an analysis of Trojan vowel behavior in Hungarian with Boolean Monadic Recursive Schemes (BMRS; Bhaskar et al. 2020, 2023; Chandlee & Jardine 2021). The relevant vowel harmony facts of Hungarian are summarized in §2. The BMRS framework and our analysis are elaborated in §3, where we show how Trojan vowel behavior follows straightforwardly from the BMRS framework when coupled with the assumption that inputs are freely specified (akin to the *Richness of the Base* hypothesis of Optimality Theory; Prince & Smolensky 1993/2004).

The set of Trojan vowel types in Hungarian is also the set that is transparent to [±back] harmony. Our analysis thus also accommodates vowel transparency, building on our prior work mapping out the typology of other forms of alternation resistance behavior in BMRS (Nelson & Baković 2025). The broader aim of our work is to identify the computational structures and basic predicate types that underlie different types of phonological spreading patterns, like other recent work that relates different types of phonological generalizations to specific BMRS structures (Chandlee & Jardine 2021; Jardine & Oakden 2023; Nelson & Baković 2024). In §4 we address the fact that Trojan vowels can otherwise be *opaque* to harmony, as in Yoruba (Bamgboṣe 1967; Archangeli & Pulleyblank 1989), in §5 we address a couple of challenges, and §6 concludes.

2. Hungarian. The vowel inventory of Hungarian is as shown in (5), brutally summarizing Siptár & Törkenczy (2000) and setting aside the fact that all seven vowel qualities contrast in length.

(5) Hungarian vowel inventory

		[−back]		[+back]	
		[−round]	[+round]	[−round]	[+round]
[−low]	[+high]	**i**	**ü**		**u**
[−low]	[−high]	**e**	**ö**		**o**
[+low]	[−high]			**a**	

Note both the [±back] symmetries and asymmetries in this inventory. The only 'perfectly paired' vowels are the [+round] ones: **ü** is paired with **u** and **ö** with **o**. In general, then, suffixes with these vowels alternate to agree with stems in terms of [±back]: √haːz + unk 'our house' ∼ √kert + ünk 'our garden'; √laːb + uː '-legged' ∼ √fej + üː '-headed'; √haːz + hoz 'to (the) house' ∼ √föld + höz 'to (the) land';[2] √vaːr + oː 'waiting (adj.)' ∼ √keːr + öː 'asking (adj.)'.

The [−round] vowels, on the other hand, are not perfectly paired, with different consequences for the [+low] vowel **a** and the [−low] vowels **i**, **e**. The [−low] vowels are alternation-resistant, surfacing redundantly as [−back] in suffixes even when preceding stem vowels are [+back]: √fut + ni 'to run', √tan + iːt 'teach', √nyoma + teːk 'emphasis'. Further evidence shows that these alternation-resistant vowels are transparent to [±back] harmony, as we'll see below.[3]

The [+low] vowel **a**, on the other hand, does alternate, but with [−low] **e**: √haːz + ban 'in (the) house' ∼ √kert + ben 'in (the) garden'; √vaːr + naː 'he would wait for it' ∼ √keːr + neː 'he would ask for it'. This is due to *re-pairing* (Baković 2000): assuming that the vowels in these suffixes are underlyingly [+low], harmonic pressure to surface as [−back] results in an additional

[2] There is an additional [−round] alternant of short mid vowel suffixes like this one (e.g. √kert + hez 'to (the) garden') that we do not discuss here; see e.g. Ringen & Vago (1998:406ff) and Siptár & Törkenczy (2000:72ff).

[3] There are also some suffixes with *unpredictably* alternation-resistant [+back] vowels that are opaque, suggesting that they are external to the harmony domain; see Siptár & Törkenczy (2000:65–66) for examples and discussion.

change from [+low] to [−low]. But re-pairing is asymmetrical: underlying [−low] vowels are not re-paired and instead resist alternation and are transparent, as just discussed in the previous paragraph. We address this asymmetry in re-pairing at the end of our analysis in §3.

Putting aside all alternation-resistant vowels, then, the examples tabulated in (6) illustrate the basic evidence for [±back] harmony in Hungarian. The unsuffixed stem and gloss are in the first two columns, followed by the dative and ablative forms. Each type of form in the pattern is schematized in the last column, where we refer to the [−back] (that is, 'front') vowels as 'F' and to the [+back] (that is, 'back') vowels as 'B', following Siptár & Törkenczy (2000). Stems with only F vowels take suffixes with F vowels (6a), and stems with only B vowels take suffixes with B vowels (6b). Stems may also be disharmonic, containing both F and B vowels in any order. In such cases, suffix vowels systematically agree with the last stem vowel, whether F (6c) or B (6d).

(6) Evidence for [±back] harmony in Hungarian

a.	√tü:z	'fire'	√tü:z + nek	√tü:z + tö:l	√FF+F
	√tükör	'mirror'	√tükör + nek	√tü:kör + tö:l	
b.	√ha:z	'house'	√ha:z + nak	√ha:z + to:l	√BB+B
	√kupa	'goblet'	√kupa + nak	√kupa + to:l	
c.	√sofö:r	'driver'	√sofö:r + nek	√sofö:r + tö:l	√BF+F
	√allü:r	'mannerism'	√allü:r + nek	√allü:r + tö:l	
d.	√nüansz	'nuance'	√nüansz + nak	√nüansz + to:l	√FB+B
	√amö:ba	'amoeba'	√amö:ba + nak	√amö:ba + to:l	

Stems containing only the [−low, −round] vowels **i**, **e**, referred to as 'N' (for 'neutral'), are of two types: those that take suffixes with F vowels, as would be expected given that both N and F vowels are [−back] (7a), and those that take unexpectedly but nevertheless systematically take suffixes with B vowels (7b). These latter stems thus contain Trojan vowels.

(7) Evidence for [−low, −round] Trojan vowels in Hungarian

a.	√vi:z	'water'	√vi:z + nek	√vi:z + tö:l	√NN+F
	√szege:ny	'poor'	√szege:ny + nek	√szege:ny + tö:l	
b.	√hi:d	'bridge'	√hi:d + nak	√hi:d + to:l	√NN+B
	√dere:k	'waist'	√dere:k + nak	√dere:k + to:l	

Last, but not least: polysyllabic stems in which N precedes F (8a) or N precedes B (8b) behave as expected: suffix vowels systematically agree with the last stem vowel, just like (6c–d). Stems in which N *follows* F (8c) also behave as expected: all stem vowels are [−back], and so are the vowels of following suffixes. Stems in which N follows B provide evidence for the transparency of N: following suffix vowels are [+back], in agreement with B and skipping over N.

(8) Evidence for [−low, −round] transparent vowels in Hungarian

a.	√szemölcs	'wart'	√szemölcs + nek	√szemölcs + tö:l	√NF+F
	√re:zsü:	'slope'	√re:zsü: + nek	√re:zsü: + tö:l	
b.	√bika	'bull'	√bika + nak	√bika + to:l	√NB+B
	√hernyo:	'caterpillar'	√hernyo: + nak	√hernyo: + to:l	

c.	√**ü:veg**	'glass'	√**ü:veg** + **nek**	√**ü:veg** + **tö:l**	√FN+F
	√**rövid**	'short'	√**rövid** + **nek**	√**rövid** + **tö:l**	
d.	√**papi:r**	'paper'	√**papi:r** + **nak**	√**papi:r** + **to:l**	√BN+B
	√**korde:**	'cart'	√**korde:** + **nak**	√**korde:** + **to:l**	

This concludes the empirical overview of [±back] harmony in Hungarian.[4] We continue in §3 with our overview of the BMRS formalism and our analysis within it.

3. BMRS analysis. From a technical perspective, BMRS provide logical characterizations of subclasses of rational functions (Bhaskar et al. 2020, 2023; Yolyan 2025b). As a phonological and typological framework, BMRS is a formalism uniting computational universals discovered in work rooted in formal language theory (e.g. Johnson 1972; Kaplan & Kay 1994; Heinz 2018) and substantive universals emerging from descriptive and theoretical phonological analysis (e.g. Greenberg 1966; Hyman 2008). We provide a brief introduction to phonological analysis using BMRS and then use them to characterize Hungarian backness harmony as described above.

3.1. Essential elements. We follow Chandlee & Jardine (2021) in defining BMRS for a general audience.[5] A BMRS program defines string-to-string mappings via functions $\phi_f(x)$ which determine whether an output string position x has property f based on properties of an input string. They are *boolean* because all functions within the program must have an output type of `Boolean` $\in \{\top, \bot\}$, *monadic* because all functions take as their input a singular domain element, and *recursive* because functions may reference their own output properties. An inductive definition of BMRS functions is provided in (9). At the root of every BMRS function is the evaluation of one or more predicates, $P(t)$, (9b). P could be a phonological feature (e.g. f; $[\![f(x) = \top]\!]$ = [+f], $[\![f(x) = \bot]\!]$ = [−f]), or a boundary ($\rtimes$ = left, $\ltimes$ = right), and t is a string position, the position under consideration (x), its predecessor ($p(x)$), or its successor ($s(x)$).

(9) a. $\top$ (= true) and $\bot$ (= false) are expressions
b. any predicate $P(t)$ is an expression
c. if E_1, E_2, and E_3 are expressions, then so is '`IF` E_1 `THEN` E_2 `ELSE` E_3'
d. nothing else is an expression

The structures over which our functions operate are called *monadic string models* (Chandlee & Lindell, to appear; Yolyan 2025a), defined by denoting how a set of relations and functions $M = \langle \{R_f \mid f \in \mathcal{F}\}, \rtimes, \ltimes, s, p \rangle$ apply to a finite domain $\mathcal{D}$, where $\mathcal{F}$ is the set of all features. The domain of a given structure is the set of natural numbers from zero to the length of the word $|w|$ that the structure represents, plus one. The zero position must be labeled as the left word boundary ($\rtimes$) and the $w + 1$ position must be labeled as the right word boundary ($\ltimes$). One consequence of these axioms is that string-initial and string-final positions are easily computed: they are positions x such that $p(x) = \rtimes$ or $s(x) = \ltimes$, respectively, as we'll see in (12) below.

In the examples in (10), we show that a simple change in E_2 can lead to noteworthy phonological generalizations. Each function in (10) computes an output value for the feature [±nasal]

[4] Note that there are also 'mixed vacillating' √BN stems after which suffix vowels vary between F and B, and 'mixed disharmonic' √BN stems after which suffix vowels are consistently F. See §5.1 for further discussion of these.
[5] Readers interested in a more technical introduction are referred to Bhaskar et al. (2020) and Bhaskar et al. (2023).

and is conditioned on whether the target element is [+syllabic] (i.e., a vowel). In (10a), vowels are redundantly [−nasal] because $E_2 = \bot$. In (10b), [±nasal] spreads to vowels from the preceding *input* segment because $E_2 = \texttt{nas}(p(x))$. This is non-recursive and equivalent to non-iterative assimilation. In (10c), [±nasal] spreads to vowels from the preceding *output* segment because $E_2 = \phi_{\texttt{nas}}(p(x))$. This is a recursive call and equivalent to iterative spreading. In each case, $E_3 = \texttt{nas}(x)$ ensures that [−syllabic] segments (i.e., consonants) remain faithful to their underlying specification for [±nasal] regardless of what happens to vowels.

(10) Given $\phi_{\texttt{nas}}(x) := \texttt{IF syll}(x)\ \texttt{THEN}\ E_2\ \texttt{ELSE nas}(x)\ldots$

if E_2 is…		then /NVVVCV/ maps to…	and the overall pattern is…
a.	$\bot$	[NVVVCV]	vowels are redundantly [−nasal]
b.	$\texttt{nas}(p(x))$	[NṼVVCV]	non-iterative progressive [±nasal] spreading
c.	$\phi_{\texttt{nas}}(p(x))$	[NṼṼṼCV]	iterative progressive [±nasal] spreading

A function of the form in (10c) is not the only way to describe iterative spreading. In Nelson & Baković (2025), we show that BMRS provide a straightforward way to define a typology of different spreading patterns, adapting definitions from Archangeli & Pulleyblank (1994) to describe the basic elements of spreading processes and their associated BMRS function subparts. The four basic elements are **content**, the phonological feature (value) that spreads; **source**, the position that initiates spreading; **orientation**, the direction in which the feature spreads relative to the source, and **redundancy**, any defined conditions on the association of the spreading feature.

The content is determined by an output condition of the form $\phi_{\texttt{spf}}(x) := E$, which says that [±spf] is the *spreading feature* determined for every element x by the expression E. The source is determined by a *source faithfulness* condition of the form $\texttt{IF src}(x)\ \texttt{THEN spf}(x)\ \texttt{ELSE}\ldots$, which says that when the positional predicate src (= initial, final, stem, …) holds for a domain element x, then it remains faithful to its input specification for [±spf]. The orientation is determined by an *orientation setting* condition based on whether the predecessor function ($p(x)$; left-to-right) or successor function ($s(x)$; right-to-left) is used in the spreading call, and the iterativity of spreading depends on whether the spreading call refers to an input value, leading to non-iterative local assimilation ($\ldots\texttt{ELSE spf}(\{^p_s\}(x))$), or to an output value, leading to iterative spreading ($\ldots\texttt{ELSE}\ \phi_{\texttt{spf}}(\{^p_s\}(x))$). Finally, redundancies are determined by *redundancy conditions* of the form $\texttt{IF rdn}(x)\ \texttt{THEN}\ \left\{^\top_\bot\right\}\ \texttt{ELSE}\ldots$, ensuring that a domain element x matching the properties defined by $\texttt{rdn}(x)$ is redundantly specified as [+spf] (= $\top$) or as [−spf] (= $\bot$).

In our previous work we showed that altering the presence and scope of different types of redundancy conditions describes a formal typological space for spreading and blocking patterns.All of those patterns were iterative, which from a BMRS perspective means that the spreading feature is transmitted through the string recursively based on adjacent output specifications. One of the things that makes the Hungarian pattern interesting from our perspective is the fact that the spreading is iterative, but only after non-iterative initial transmission from the source vowel to all following targets of spreading. This initial non-iterative spreading is necessary in order to account for the Trojan behavior of certain stem vowels. In the next subsection, we outline our formal BMRS analysis of the pattern based on this high-level characterization.

3.2. Hungarian backness harmony. Recall from §2 that the neutral vowels of Hungarian are all and only the [−low, −round] vowels. We thus define the function $\texttt{neut}(x)$ in (11) identifying the class of neutral vowels; only vowels that are neither [+low] nor [+round] will evaluate to $\top$.[6]

(11) $\texttt{neut}(x) := \texttt{IF syll}(x)\ \texttt{THEN} \begin{bmatrix}\texttt{IF low}(x)\ \texttt{THEN}\ \bot\ \texttt{ELSE} \\ \quad \texttt{IF round}(x)\ \texttt{THEN}\ \bot\ \texttt{ELSE}\ \top\end{bmatrix} \texttt{ELSE}\ \bot$

We now identify the first stem vowel, the potential source of Trojan vowel spreading. We adopt a method introduced by Bhaskar et al. (2023) and that was shown by Nelson & King (submitted) to be directly applicable to computing tier-based functions. This method requires that we define two functions: one to identify the first vowel of the string (which is also the first vowel of the stem), $\texttt{init}_v(x)$, and another, helper function $g(x)$ called from within the body of $\texttt{init}_v(x)$.

(12) a. $\texttt{init}_v(x) := \texttt{IF syll}(x)\ \texttt{THEN}\ g(x)\ \texttt{ELSE}\ \bot$

b. $g(x) := \texttt{IF}\ \rtimes(p(x))\ \texttt{THEN}\ \top\ \texttt{ELSE}\ \big[\texttt{IF syll}(p(x))\ \texttt{THEN}\ \bot\ \texttt{ELSE}\ g(p(x))\big]$

The functions in (12) work together as follows. The $\texttt{init}_v(x)$ function (12a) states that for any domain element that is a vowel ($\texttt{IF syll}(x)$), whether it is the *first* vowel is determined by the helper function ($\texttt{THEN}\ g(x)$); otherwise, any non-vowel is perforce not the first vowel ($\texttt{ELSE}\ \bot$). The $g(x)$ helper function (12b) positively identifies the string-initial element ($\texttt{IF}\ \rtimes(p(x))\ \texttt{THEN}\ \top$), transmitting its 'initialness' up to and including the first vowel, but not propogating it any further. This is accomplished by the first part of the embedded clause, which states that if an element is preceded by a vowel ($\texttt{IF syll}(p(x))$), then it is redundantly false for $g(x)$ ($\texttt{THEN}\ \bot$); since the first vowel is a vowel, there is thus no way for the initialness to propogate beyond it. These results are then lifted to the $\texttt{init}_v(x)$ function ($\texttt{THEN}\ g(x)$). Despite the fact that some non-vowels may evaluate to $\top$ for $g(x)$, it is only the first vowel's value that will end up being true for $\texttt{init}_v(x)$.

Table 1 illustrates how the functions defined thus far evaluate the input string /√**papi:r** + **tö:l**/. Dashed arrows show how non-spreading ($g(x)$) and non-redundant ($\texttt{init}_v(x)$) values are determined and the solid arrows show how transmission of initialness occurs in the $g(x)$ function.

Input:	$\rtimes$	p	**a**	p	**i:**	r	t	**ö:**	l	$\ltimes$
x	0	1	2	3	4	5	6	7	8	9
$\texttt{syll}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\top$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{neut}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$g(x)$	$\bot$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\texttt{init}_v(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$

Table 1. Evaluation table determining which domain element is the first vowel

The last thing before getting to spreading proper is a function which determines membership on the *harmonic tier*, $\tau_h(x)$, consisting of the first stem vowel (identified by $\texttt{init}_v(x)$) and all remaining non-neutral vowels (the complement of $\texttt{neut}(x)$). A segment's membership on the

[6] This function and others we define below are not atomic properties of our monadic string models, but can instead be thought of as "user-defined predicates" that help make the analysis more interpretable without any worry of expanding the expressive capabilities of the formalism (Strother-Garcia 2018, 2019).

harmonic tier is thus first conditioned on being a vowel. The first stem vowel is automatically on the tier; remaining neutral vowels are not, and any non-neutral vowels remaining after that are. Given that tier membership is determined in part on string position rather than labeled properties of an element, this is an example of *structure-sensitive* tier projection, which has been explored in relation to formal languages (sets of strings) but not for function classes (De Santo & Graf 2019).

(13) $$\tau_h(x) := \texttt{IF}\ \texttt{syll}(x)\ \texttt{THEN} \begin{bmatrix} \texttt{IF}\ \texttt{init}_v(x)\ \texttt{THEN}\ \top\ \texttt{ELSE} \\ \quad \texttt{IF}\ \texttt{neut}(x)\ \texttt{THEN}\ \bot\ \texttt{ELSE}\ \top \end{bmatrix} \texttt{ELSE}\ \bot$$

This brings us to the two functions required for spreading [$\pm$back] over a tier (Bhaskar et al. 2023; Nelson & King, submitted). One of these, called $\phi^*_{\texttt{back}}(x)$ and shown in (14), is the one that determines the output specification for the spreading feature on all domain elements. This function first specifies that neutral vowels are redundantly [$-$back] ($\texttt{IF}\ \texttt{neut}(x)\ \texttt{THEN}\ \bot$) and otherwise restricts attention to vowels on the harmonic tier ($\texttt{IF}\ \tau_h(x)\ \texttt{THEN} \ldots$). Stem vowels are then protected by a source faithfulness condition ($\texttt{IF}\ \texttt{stem}(x)\ \texttt{THEN}\ \texttt{back}(x)$) and iterative spreading proceeds with reference to the second function in (15) below ($\texttt{ELSE}\ \phi^{\tau_h}_{\texttt{back}}(p(x))$). We assume that remaining non-tier elements (i.e., consonants) are redundantly specified as [$-$back] ($\texttt{ELSE}\ \bot$).

(14) $$\phi^*_{\texttt{back}}(x) := \texttt{IF}\ \texttt{neut}(x)\ \texttt{THEN}\ \bot\ \texttt{ELSE}\ \texttt{IF}\ \tau_h(x)\ \texttt{THEN} \begin{bmatrix} \texttt{IF}\ \texttt{stem}(x)\ \texttt{THEN}\ \texttt{back}(x) \\ \quad \texttt{ELSE}\ \phi^{\tau_h}_{\texttt{back}}(p(x)) \end{bmatrix} \texttt{ELSE}\ \bot$$

The limitation of iterative spreading to harmonic tier elements is due to the second function, $\phi^{\tau_h}_{\texttt{back}}(x)$, shown in (15). This function recursively skips over non-tier elements, resulting in tier-relativized local transmission of spreading feature information. It accomplishes this by first restricting attention to the harmonic tier ($\texttt{IF}\ \tau_h(x)\ \texttt{THEN} \ldots$). Then, the initial element of that tier is identified and its input specification for [$\pm$back] is transmitted ($\texttt{IF}\ \texttt{init}_v(x)\ \texttt{THEN}\ \texttt{back}(x)$), thus capturing the potential Trojan behavior of the first stem vowel. Otherwise, the output specification for [$\pm$back] calculated by the function in (14) is transmitted ($\texttt{ELSE}\ \phi^*_{\texttt{back}}(x)$). Remaining, non-tier elements recursively transmit information from their predecessors ($\texttt{ELSE}\ \phi^{\tau_h}_{\texttt{back}}(p(x))$).

(15) $$\phi^{\tau_h}_{\texttt{back}}(x) := \texttt{IF}\ \tau_h(x)\ \texttt{THEN} \begin{bmatrix} \texttt{IF}\ \texttt{init}_v(x)\ \texttt{THEN}\ \texttt{back}(x) \\ \quad \texttt{ELSE}\ \phi^*_{\texttt{back}}(x) \end{bmatrix} \texttt{ELSE}\ \phi^{\tau_h}_{\texttt{back}}(p(x))$$

While the $\phi^*_{\texttt{back}}(x)$ spreading function in (14) uses the $\tau_h(x)$ function in (13) to remove non-tier elements from the set of targets, the $\phi^{\tau_h}_{\texttt{back}}(x)$ transmission function in (15) uses $\tau_h(x)$ to remove them from the set of triggers. Instead of redundantly specifying non-tier elements, $\phi^{\tau_h}_{\texttt{back}}(x)$ actively skips over them. This is why the outer conditional is a recursive call on elements preceding non-tier elements; transmission succeeds only once a tier element is identified.

3.3. Evaluations. In Table 2 we illustrate how [$\pm$back] is computed for √BN+B and √FN+F forms, from hypothetical inputs for the examples √**papiːr** + **toːl** and √**röːvid** + **nek**. To show that (and how) the functions really work, we assume in each case that the underlying [$\pm$back] values of the second, neutral vowel and the suffix vowel are the opposite of the (underlying and surface) [$\pm$back] value of the first stem vowel: /√**papiːr** + **töːl**/ and /√**röːvɯd** + **nak**/.[7]

[7] It is often assumed that such completely predictable surface feature values are (or *must be*) underlyingly underspecified. See Nelson (2022) for discussion of the problem of representational underspecification in logical transductions,

Input:	⋊	p	**a**	p	**i:**	r	t	**ö:**	l	⋉
x	0	1	2	3	4	5	6	7	8	9
$\texttt{neut}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\tau_h(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{stem}(x)$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$
$\texttt{back}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{init}_v(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{\tau_h}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$
$\phi^{*}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
Output:	⋊	p	**a**	p	**i:**	r	t	**o:**	l	⋉

Input:	⋊	r	**ö:**	v	**ɯ**	d	n	**a**	k	⋉
x	0	1	2	3	4	5	6	7	8	9
$\texttt{neut}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\tau_h(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{stem}(x)$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$
$\texttt{back}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{init}_v(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{\tau_h}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{*}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
Output:	⋊	r	**ö:**	v	**i**	d	n	**e**	k	⋉

Table 2. Evaluations determining [±back] values for √BN+B (transparent) and √FN+F forms

Dashed arrows show how [±back] is computed and solid arrows show the transmission effect. Positions 2, 4, and 7 — the vowels — are highlighted. The vowel in position 2 in both cases is the first stem vowel ($\texttt{stem} = \top$, $\texttt{init}_v(x) = \top$), is on the harmonic tier ($\tau_h(x) = \top$), and is not neutral ($\texttt{neut}(x) = \bot$), so its output value $\phi^{*}_{\texttt{back}}(x)$ is the same as its input value $\texttt{back}(x)$. The vowel in position 4 in both cases is a neutral vowel ($\texttt{neut}(x) = \top$) and in the stem ($\texttt{stem}(x) = \top$), but not first ($\texttt{init}_v(x) = \bot$), so it is not on the harmonic tier ($\tau_h(x) = \bot$) and is redundantly [−back] ($\phi^{*}_{\texttt{back}}(x) = \bot$). The vowel in position 7 is not in the stem ($\texttt{stem}(x) = \bot$) but is not neutral ($\texttt{neut}(x) = \bot$), so it is on the harmonic tier ($\tau_h(x) = \top$) and its value is thus determined by the preceding element on the tier (position 2). This value is covertly transmitted by $\phi^{\tau_h}_{\texttt{back}}(x)$ from position 2 through positions 3–6 and ultimately to position 7.

In Table 3, [±back] is computed for √NN+B and √NN+F forms using the examples √d**e**r**e:**k + t**o:**l and √s**e**g**e:**ny + n**e**k. The Trojan vowel analysis requires that the first, neutral stem vowel be underlyingly [+back] in the √NN+B case and [−back] in the √NN+F case, so the hypothetical inputs are /√d**ɤ**r**e:**k + t**ö:**l/ and /√s**e**g**ɤ:**ny + n**a**k/. The computation paths for positions 4 and 7 are identical to Table 2 above. The primary difference now is the effect of the first stem vowel in position 2 being neutral, meaning that its output value is redundantly [−back] ($\phi^{*}_{\texttt{back}}(x) = \bot$), but because of how the transmission function $\phi^{\tau_h}_{\texttt{back}}(x)$ works, the input [±back] value of this first neutral stem vowel is ultimately what determines the [±back] value for position 7.

Input:	⋊	d	**ɤ**	r	**e:**	k	t	**ö:**	l	⋉
x	0	1	2	3	4	5	6	7	8	9
$\texttt{neut}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\tau_h(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{stem}(x)$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$
$\texttt{back}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{init}_v(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{\tau_h}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$
$\phi^{*}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
Output:	⋊	d	**e**	r	**e:**	k	t	**o:**	l	⋉

Input:	⋊	s	**e**	g	**ɤ:**	ny	n	**a**	k	⋉
x	0	1	2	3	4	5	6	7	8	9
$\texttt{neut}(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\tau_h(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{stem}(x)$	$\bot$	$\top$	$\top$	$\top$	$\top$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$
$\texttt{back}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$
$\texttt{init}_v(x)$	$\bot$	$\bot$	$\top$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{\tau_h}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
$\phi^{*}_{\texttt{back}}(x)$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
Output:	⋊	s	**e**	g	**e:**	ny	n	**e**	k	⋉

Table 3. Evaluations determining [±back] values for √NN+B (Trojan) and √NN+F forms

Nelson & Baković (2024) for a computational solution, and Chandlee & Jardine (2021:514) for a proposal on how to modify BMRS to include underspecified representations.

3.4. Re-pairing. We return now to the fact noted in §2 that some harmonic suffix vowels alternate between [+low, +back] **a** and [−low, −back] **e**, while other, neutral suffix vowels always surface as **e**. There must of course be an underlying distinction between each of these types of suffix vowels. Standard methods of phonological analysis point us to the following conclusion: alternation-resistant **e** is underlyingly /**e**/, while alternating **a** ∼ **e** must *not* be /**e**/, and so is /**a**/.

In the context of our analysis, this is only half right. Recall that both the alternation resistance of neutral **e** and the re-pairing alternation **a** ∼ **e** are due to [±back] asymmetries in the Hungarian vowel inventory: there is no [−low, +back, −round] vowel ***ɤ** and no [+low, –back] vowel ***æ**. The underlying representation of the non-alternating neutral **e** could thus be *either* [−back] /**e**/ *or* [+back] /**ɤ**/, surfacing redundantly as [−back] as our analysis thus far already states, while the underlying representation of alternating **a** ∼ **e** can be *either* [+back] /**a**/ *or* [−back] /**æ**/, surfacing redundantly as [−low] when it is harmonically compelled to surface as [−back]. The $\phi_{\texttt{low}}(x)$ function in (16) accomplishes this result: output [+back] segments are faithful to their underlying value of [±low] ($\texttt{IF}\ \phi^{*}_{\texttt{back}}(x)\ \texttt{THEN}\ \texttt{low}(x)$); otherwise, [−back] segments are [−low] ($\texttt{ELSE}\ \bot$).

(16) $\phi_{\texttt{low}}(x) := \texttt{IF}\ \phi^{*}_{\texttt{back}}(x)\ \texttt{THEN}\ \texttt{low}(x)\ \texttt{ELSE}\ \bot$

A brief comparison with Turkish (Clements & Sezer 1982) is worthwhile here. As was shown in (1), [±back] harmony in Turkish also results in a re-pairing alternation between [+low] **a** and [−low] **e**. In this case, however, **e** is not independently neutral and the re-pairing is fully symmetrical: the underlying representation of **a** ∼ **e** suffix vowels in Turkish could be any of /**a**/, /**æ**/, /**e**/, or /**ɤ**/.[8] Putting aside the (axiomatic) fact that [+high] vowels are redundantly [−low], the $\phi_{\texttt{low}}(x)$ function for Turkish would thus differ from (16) by replacing $\texttt{THEN}\ \texttt{low}(x)$ with $\texttt{THEN}\ \top$.

3.5. Free input specification. Note that our analysis of [±back] harmony in Hungarian is structured in such a way that the observed surface patterns hold regardless of the [±back] specifications of any of the vowels in the input. This result is rooted in the assumption that our BMRS program represents a *total function* over the set of possible phonological structures that can be built from the set of representational primes (e.g., the feature set $\mathcal{F}$), an assumption related (but not identical) to the Richness of the Base hypothesis of OT (Prince & Smolensky 1993/2004).

Gaps in the surface inventory are thus controlled directly by the grammar rather than by limiting the set of possible inputs. Thus, when a neutral vowel is not the first stem vowel, the redundancy condition at the start of the $\phi^{*}_{\texttt{back}}(x)$ function in (14), $\texttt{IF}\ \texttt{neut}(x)\ \texttt{THEN}\ \bot$, ensures that the vowel surfaces as [−back], and the same condition embedded in the $\tau_h(x)$ function in (13) ensures that it is not on the harmonic tier and that it is thus transparent to harmony. On the other hand, being the first stem vowel ensures that a neutral vowel *is* on the harmonic tier, allowing a Trojan vowel to transmit its underlying [+back] value despite redundantly surfacing as [−back].

4. Trojan vowels can be opaque. Whether a [−low, −round] (i.e. neutral) vowel in Hungarian is transparent or Trojan depends on its position in the string. In our analysis, the structure-senstive $\tau_h(x)$ function in (13) places the first stem vowel (neutral or not) and all non-neutral vowels on the harmonic tier. If the first stem vowel is neutral, it is a potential Trojan, depending on its underlying value of [±back]; other neutral vowels are not on the harmonic tier and so are potentially transparent, depending on the value of [±back] of the first stem vowel. It may appear, then, that Trojan and transparent vowel behaviors are intimately intertwined in our analysis. This is not the

[8] Or indeed any [+round] variant of these, given the opacity of these [−high] suffix vowels to [±round] harmony (2).

case, nor should it be, because there exist vowel harmony systems in which the class of vowels exhibiting Trojan behavior in the source position are otherwise opaque, not transparent.

One such case is Yoruba (Bamgboṣe 1967; Archangeli & Pulleyblank 1989; see also Baković 2000:§3.6 and Krämer 2003:§6.2), where the content is [±ATR], the source is the final vowel (sometimes, but not always, identifiable as a stem vowel, *pace* Baković 2000:§3), the orientation is leftward, and there are two classes of redundantly-specified vowels: [+low] vowels are redundantly [−ATR], [+high] vowels are redundantly [+ATR], and each of these classes of vowels is opaque to the spread of the opposite value of [±ATR] in the manner shown in (2). However, some final [+high] vowels are unexpectedly preceded by [−ATR] vowels, a systematic fact that is most apparent when the final vowel is a stem vowel and preceding vowels are prefix vowels.

(17) a. ò + √kú 'corpse' *cf.* √kú 'to die' < /√kú/ (presumed 'normal')
b. ɔ̀ + √mu 'drinker' *cf.* √mu 'to drink' < /√mʊ/ (presumed Trojan)

A more complete analysis of the pattern in Yoruba is provided in (18), following the structure of the analysis of Hungarian in §3.2. The final vowel in Yoruba is identified in the same way that the initial vowel is in Hungarian, with the $\texttt{final}_v(x)$ function in (18a) and the associated, now right-oriented helper function $g(x)$ in (18b). Unlike Hungarian, the harmonic tier in Yoruba includes all vowels, regardless of position; the $\tau_h(x)$ function in (18c) is thus superfluously co-extensive with $\texttt{syll}(x)$. This is the critical difference that makes the relevant classes of vowels opaque in Yoruba rather than transparent: the corresponding Hungarian $\tau_h(x)$ function achieves transparency because it excludes all neutral vowels that are not also the first vowel.

(18) a. $\texttt{final}_v(x) := \texttt{IF } \texttt{syll}(x) \texttt{ THEN } g(x) \texttt{ ELSE } \bot$

b. $g(x) := \texttt{IF } \ltimes(s(x)) \texttt{ THEN } \top \texttt{ ELSE } \big[\texttt{IF } \texttt{syll}(s(x)) \texttt{ THEN } \bot \texttt{ ELSE } g(s(x))\big]$

c. $\tau_h(x) := \texttt{IF } \texttt{syll}(x) \texttt{ THEN } \top \texttt{ ELSE } \bot$

d. $$\phi^{\tau_h}_{\texttt{ATR}}(x) := \texttt{IF } \tau_h(x) \texttt{ THEN} \begin{bmatrix} \texttt{IF } \texttt{final}_v(x) \texttt{ THEN} \\ \quad \texttt{IF } \texttt{low}(x) \texttt{ THEN } \bot \texttt{ ELSE } \texttt{ATR}(x) \\ \texttt{ELSE } \phi^{*}_{\texttt{ATR}}(x) \end{bmatrix} \texttt{ELSE } \phi^{\tau_h}_{\texttt{ATR}}(s(x))$$

e. $$\phi^{*}_{\texttt{ATR}}(x) := \begin{bmatrix} \texttt{IF } \texttt{low}(x) \texttt{ THEN } \bot \texttt{ ELSE} \\ \quad \texttt{IF } \texttt{high}(x) \texttt{ THEN } \top \texttt{ ELSE} \end{bmatrix} \texttt{IF } \tau_h(x) \texttt{ THEN} \begin{bmatrix} \texttt{IF } \texttt{final}_v(x) \texttt{ THEN } \texttt{ATR}(x) \\ \quad \texttt{ELSE } \phi^{\tau_h}_{\texttt{ATR}}(s(x)) \end{bmatrix} \texttt{ELSE } \bot$$

The main action is in the remaining two functions. The transmission function in (18d) is the mirror image of the corresponding function for Hungarian in (15), referencing the final vowel and successors rather than the initial vowel and predecessors. Given that $\tau_h(x)$ includes all vowels in Yoruba, this transmission function only skips consonants. However, like Hungarian, it provides the ultimate source (here, final) vowel the opportunity to transmit its underlying [±ATR] value, but only if it is also [−low] — a necessary exclusion, given that both [+low] and [+high] vowels surface with redundant values of [±ATR] but only [+high] vowels are potential Trojans.

Last but not least, the harmony function in (18e) determines output values of [±ATR]. This function is likewise the mirror image of the corresponding function for Hungarian in (14). The other notable difference is that there is only one class of redundantly-specified vowels in Hungarian, and there are two such classes in Yoruba: [+low] vowels are redundantly specified as [−ATR] (IF low(x) THEN ⊥) and [+high] vowels as [+ATR] (IF high(x) THEN ⊤). Again, the fact that only the [+high] vowels exhibit Trojan behavior is due to the exclusion of [+low] vowels in (18d).

5. Empirical challenges. Here we address a couple of empirical challenges for the characterization of neutral vowels as transparent or abstractly Trojan, as we've done in our analysis.

5.1. Mixed stems. In addition to the √BN stems that consistently take B suffix alternants (8d) and covered by our analysis, Hungarian also has √BN stems that consistently take F suffix alternants (√**o**kt**o**:b**e**r + t**ö**:l 'from October', √**o**xig**e**:n + n**e**l 'with oxygen') and others that variably take B and F suffix alternants (√k**o**nkr**e**:t + **a**n/**e**n 'concretely', √kl**a**rin**e**:t + t**a**l/t**e**l 'with a clarinet'). The former are referred to as 'mixed disharmonic' and the latter as 'mixed vacillating' by Siptár & Törkenczy (2000:69ff); see also (Ringen & Kontra 1989; Hayes & Londe 2006; Hayes et al. 2009) for relevant experimental and corpus evidence. Our current BMRS analysis does not account for these mixed stems, but we briefly suggest one way it could be adapted to do so.

The key to this solution is that non-initial neutral vowels may be *probabilistically* placed on the harmonic tier (*cf.* Mayer 2021). The evaluation of non-initial neutral vowels by the $\tau_h(x)$ function in (13) would be changed from simply ⊥ to a function that samples from a Bernoulli distribution over the set $\{\top, \bot\}$. Since non-initial harmonic tier elements spread their output value for [±back], a non-initial neutral vowel that ends up on the harmonic tier would cause following suffix vowels to take on its redundant [−back] value, despite the first vowel of the stem potentially being [+back] — that is, some √BN stems will (sometimes) behave like √BF stems (6c).

Hayes & Londe (2006) demonstrate that bisyllabic √BN stems are more likely to take [+back] affix vowels than longer (√BNN) stems are. The probabilistic approach to tier membership just sketched accounts for this observation. Suppose the probability mass on ⊤ is some non-zero value $x \in (0, 1]$. The probability of a √BN stem taking a [+back] stem is then x. Assuming independence, the probability of a √BNN stem taking a [+back] stem is $x + x - (xx) = 2x - x^2$. Since x is a value between zero and one, the probability of one of the two neutral vowels being on the harmonic tier is guaranteed to be greater than x, therefore leading to there being more probability for [−back] suffix vowels with √BNN stems than with √BN stems, exactly as observed.

One issue with this approach is that it makes our system non-deterministic and therefore non-functional. The technical properties of BMRS ensure that they only compute string-to-string *functions*, which by definition must be deterministic. There are ways to get around this by using weighted monadic second-order logic (Droste & Gastin 2009) to encode *semi-deterministic* functions that map a single input string to a set of strings (Beros & Higuera 2014). This has been briefly explored for phonological analysis using predicate logic (Heinz, forthcoming), and we hope to pursue this approach as it pertains to BMRS analyses in more detail in future work.

5.2. Phonetics-phonology interface. Benus & Gafos (2007) examined the articulatory properties of stem-final neutral vowels in Hungarian and found that neutral vowels followed by a [+back] suffix vowel are articulated with a more rectracted tongue body gesture than neutral vowels followed by a [−back] suffix vowel. Operating under the assumption that harmony is strictly local, they argue that the suffix vowel alternation is sensitive to subtle differences in the articulatory properties of neutral vowels and therefore that the phonological grammar needs to be structured in a way that can integrate both macro- and microscopic properties of speech sounds. They implement this idea with a phonological grammar defined using nonlinear dynamics (Gafos & Benus 2006), which mathematically integrates continuous and discrete information.

It is challenging to directly relate these results to our approach since Benus & Gafos assume a completely different characterization of the harmony pattern. The biggest distinction between the two approaches arguably has to do with the acceptance of abstract vowels. We readily accept abstract underlying representations that necessarily absolutely neutralize on the surface while Benus & Gafos do not. One consequence of our acceptance of abstractness is that we are able to provide a different type of explanation for the apparent articulatory facts. Recently, Nelson & Heinz (2025) propose a model of the phonetics-phonology interface called the *Blueprint Model of Production* (BMP). The BMP states that all phonetic implementation is a combination of the corresponding phonetic properties for both underlying and surface specifications of phonological segmental structure. Typically this manifests as small traces of the underlying form in the phonetic signal (e.g., incomplete neutralization), but notably is scaled by external factors. Our speculative explanation for Benus & Gafos's articulatory observations is thus that the neutral vowels articulated with a more retracted tongue gesture are those which are also underlyingly [+back]. The subtle phonetic differences are traces of the underlying form rather than the output target.

There are ways in which this (again, speculative) hypothesis could be empirically tested. For example, Zeng et al. (2025) show that an incomplete neutralization effect in Mandarin third tone sandhi disappeared when participants were required to do a simultaneous task that affected verbal working memory. We envision a similar set up for testing Hungarian speakers' productions of neutral vowels. If the subtle articulatory differences are specified output targets, then the effect should persist regardless of additional pressures on working memory. On the other hand, if the BMP claim that both underlying and surface specifications affect phonetic implementation is correct, we might expect the observed retraction of neutral vowels in [+back] contexts to disappear given additional pressures on working memory much as it did in Zeng et al.'s study.

Paramore & Bennett (2025) also argue for abstract and covert underlying representations based on acoustic analysis of nasalization in Panjabi. In this context, we hope to explore the phonetic properties of Trojan and other types of absolutely neutralized vowels in future work.

6. Conclusion. In this paper we have provided an analysis of the [±back] harmony pattern in Hungarian using BMRS. The analysis is designed such that the observed surface pattern emerges from the structure of the individual functions within the BMRS program without placing any restrictions on the [±back] specifications of vowels in input forms. Allowing for this type of free input specification ultimately makes the potential for Trojan vowel behavior an inevitability.

References

Archangeli, Diana & Douglas Pulleyblank. 1989. Yoruba Vowel Harmony. *Linguistic Inquiry* 20(2). 173–217. https://www.jstor.org/stable/4178624.

Archangeli, Diana & Douglas Pulleyblank. 1994. *Grounded Phonology*. MIT Press.

Baković, Eric. 2000. *Harmony, dominance and control*: Rutgers University dissertation.

Baković, Eric. 2025. When harmony fails, markedness prevails. In Kathleen Currie Hall, Gunnar Ó. Hansson & Anne-Michelle Tessier (eds.), *DougSchrift: A Collection of Squibs and Puzzles presented to Doug Pulleyblank*, https://arts-pulleyblank-2024.sites.olt.ubc.ca/dougschrift/.

Bamgboṣe, Ayọ. 1967. Vowel harmony in Yoruba. *Journal of African Languages* 6. 268–273.

Benus, Stefan & Adamantios Gafos. 2007. Articulatory characteristics of Hungarian 'transparent' vowels. *Journal of Phonetics* 35(3). 271–300. https://doi.org/10.1016/j.wocn.2006.11.002.

Beros, Achilles & Colin Higuera. 2014. A canonical semi-deterministic transducer. In Alexander Clark, Makoto Kanazawa & Ryo Yoshinaka (eds.), *The 12th International Conference on Grammatical Inference*, vol. 34 Proceedings of Machine Learning Research, 33–48. Kyoto, Japan: PMLR. https://proceedings.mlr.press/v34/beros14a.html.

Bhaskar, Siddharth, Jane Chandlee & Adam Jardine. 2023. Rational functions via recursive schemes. ArXiv preprint arXiv:2302.03074.

Bhaskar, Siddharth, Jane Chandlee, Adam Jardine & Christopher Oakden. 2020. Boolean Monadic Recursive Schemes as a Logical Characterization of the Subsequential Functions. In Alberto Leporati, Carlos Martín-Vide, Dana Shapira & Claudio Zandron (eds.), *Language and Automata Theory and Applications* Lecture Notes in Computer Science, 157–169. Springer. https://doi.org/10.1007/978-3-030-40608-0_10.

Chandlee, Jane & Adam Jardine. 2021. Computational universals in linguistic theory: Using recursive programs for phonological analysis. *Language* 97(3). 485–519. https://doi.org/10.1353/lan.2021.0045.

Chandlee, Jane & Steve Lindell. To appear. Logical perspectives on strictly local transformations. In Jeffrey Heinz (ed.), *Doing Computational Phonology*, Oxford University Press.

Clements, George N. & Engin Sezer. 1982. Vowel and consonant disharmony in Turkish. In Harry van der Hulst & Norval Smith (eds.), *The Structure of Phonological Representations, Pt. 2*, 213–255. Foris.

De Santo, Aniello & Thomas Graf. 2019. Structure sensitive tier projection: Applications and formal properties. In *Formal Grammar: 24th International Conference Proceedings*, 35–50. https://doi.org/10.1007/978-3-662-59648-7_3.

Droste, Manfred & Paul Gastin. 2009. Weighted Automata and Weighted Logics. In Manfred Droste, Werner Kuich & Heiko Vogler (eds.), *Handbook of Weighted Automata*, 175–211. Berlin, Heidelberg: Springer. https://doi.org/10.1007/978-3-642-01492-5_5.

Gafos, Adamantios & Stefan Benus. 2006. Dynamics of phonological cognition. *Cognitive Science* 30(5). 905–943. https://doi.org/10.1207/s15516709cog0000_80.

Greenberg, Joseph H. 1966. Synchronic and diachronic universals in phonology. *Language* 42(2). 508–517. https://www.jstor.org/stable/411706.

Hayes, Bruce & Zsuzsa Londe. 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology* 23(1). 59–104. https://doi.org/10.1017/S0952675706000765.

Hayes, Bruce, Kie Zuraw, Péter Siptár & Zsuzsa Londe. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85(4). 822–863. https://muse.jhu.edu/pub/24/article/369841.

Heinz, Jeffrey. 2018. The computational nature of phonological generalizations. In Larry M. Hyman & Frans Plank (eds.), *Phonological Typology*, 126–195. De Gruyter Mouton. https://doi.org/10.1515/9783110451931-005.

Heinz, Jeffrey. Forthcoming. *Doing Computational Phonology*. Oxford University Press.

Hyman, Larry. 2008. Universals in phonology. *The Linguistic Review* 25. 83–137. https://doi.org/10.1515/TLIR.2008.003.

Jardine, Adam & Christopher Oakden. 2023. Computing process-specific constraints. *Linguistic Inquiry* 1–9. https://doi.org/10.1162/ling_a_00510.

Johnson, C. Douglas. 1972. *Formal Aspects of Phonological Description*. De Gruyter Mouton.

Kaplan, Ronald M & Martin Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics* 20(3). 331–378. https://aclanthology.org/J94-3001/.

Kiparsky, Paul. 1981. Vowel harmony. Unpublished manuscript, MIT.

Krämer, Martin. 2003. *Vowel Harmony and Correspondence Theory*. Mouton de Gruyter.

Mayer, Connor. 2021. Capturing gradience in long-distance phonology using probabilistic tier-based strictly local grammars. In *Proceedings of the Society for Computation in Linguistics 2021*, 39–50. https://aclanthology.org/2021.scil-1.4/.

Nelson, Scott. 2022. A model theoretic perspective on phonological feature systems. *Proceedings of the Society for Computation in Linguistics* 5(1). 1–10. https://doi.org/10.7275/g77n-za52.

Nelson, Scott & Eric Baković. 2024. Underspecification without underspecified representations. *Proceedings of the Society for Computation in Linguistics* 7(1). 352–356. https://doi.org/10.7275/scil.2227.

Nelson, Scott & Eric Baković. 2025. Feature spreading, redundancy, and blocking behavior. Manuscript, UIUC and UC San Diego. https://ling.auf.net/lingbuzz/009295.

Nelson, Scott & Jeffrey Heinz. 2025. The blueprint model of production. *Phonology* 42. e12. https://doi.org/10.1017/S0952675725100055.

Nelson, Scott & Matthew King. Submitted. Computing tier-based functions using Boolean Monadic Recursive Schemes. Manuscript, UIUC.

Paramore, Jonathan C. & Ryan Bennett. 2025. Covert URs: evidence from nasalization in Western Panjabi. Manuscript, UC Santa Cruz. https://ling.auf.net/lingbuzz/009318.

Prince, Alan & Paul Smolensky. 1993/2004. *Optimality Theory: Constraint Interaction in Generative Grammar*. John Wiley & Sons.

Ringen, Catherine O. & Miklós Kontra. 1989. Hungarian neutral vowels. *Lingua* 78(2-3). 181–191. https://doi.org/10.1016/0024-3841(89)90052-1.

Ringen, Catherine O. & Robert M. Vago. 1998. Hungarian vowel harmony in Optimality Theory. *Phonology* 15(3). 393–416. https://www.jstor.org/stable/4420136.

Siptár, Péter & Miklós Törkenczy. 2000. *The Phonology of Hungarian*. Oxford University Press.

Strother-Garcia, Kristina. 2018. Imdlawn Tashlhiyt Berber syllabification is quantifier-free. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2018*, 145–153. https://aclanthology.org/W18-0315/.

Strother-Garcia, Kristina. 2019. *Using model theory in phonology: a novel characterization of syllable structure and syllabification*: University of Delaware dissertation.

Vago, Robert M. 1973. Abstract Vowel Harmony Systems in Uralic and Altaic Languages. *Language* 49. 579–605. https://www.jstor.org/stable/412352.

Vago, Robert M. 1976. Theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry* 7(2). 243–263. https://www.jstor.org/stable/4177921.

Yolyan, Tatevik. 2025a. A Framework for Learning Phonological Maps as Logical Transductions. In *Proceedings of the 18th Meeting on the Mathematics of Language*, 44–58. https://aclanthology.org/2025.mol-1.4/.

Yolyan, Tatevik. 2025b. A Logical Characterization of Weak Determinism as Simultaneous Application. *Journal of Logic, Language and Information* 34. 89–153. https://doi.org/10.1007/s10849-025-09429-9.

Zeng, Yuyu, Chang Wang & Jie Zhang. 2025. Cascading activation in spoken word production drives incomplete neutralization: An internet-based study of Mandarin 3rd tone sandhi. *Journal of Phonetics* 112. 101428. https://doi.org/10.1016/j.wocn.2025.101428.