A TARGET-ORIENTED APPROACH TO NEUTRALITY IN VOWEL HARMONY

by

Avery Ozburn

B.Math, University of Waterloo, 2012

M.A., University of Toronto, 2014

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Linguistics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

November 2019

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

A target-oriented approach to neutrality in vowel harmony

submitted by Avery Ozburn in partial fulfillment of the requirements for

the degree of Doctor of Philosophy

in Linguistics

Examining Committee:

Gunnar Hansson
Supervisor

Douglas Pulleyblank
Supervisory Committee Member

Anne-Michelle Tessier
Supervisory Committee Member

Molly Babel
University Examiner

Patrick Moore
University Examiner

Abstract

This dissertation provides a novel perspective on neutrality in vowel harmony, using evidence from multiple front/back and ATR harmony systems. While many standard accounts of harmony assume an equivalence between vowels that are neutral to harmony and those that lack a counterpart in the harmonic feature (e.g. van der Hulst 2016), this correspondence is demonstrably false in both directions. For example, in Hungarian (Chapter 3), [e:] lacks a harmonic counterpart, but is not consistently neutral to front/back harmony, in that it can alternate harmonically in some suffixes with [a:]. Conversely, in Mayak (Chapter 9), [a] has a contrastive ATR counterpart, yet is nonetheless neutral to ATR harmony.

I argue that these types of patterns force a new, target-focused approach, where participation is based on the drive of specific vowel qualities to undergo harmony; neutrality results when this drive is insufficient to force unfaithfulness. This idea is motivated by cross-linguistic and phonetic facts suggesting that vowels that are low and/or rounded are inherently better targets of front/back harmony, while higher vowels are better targets of ATR-dominant harmony.

I implement this approach formally in Harmonic Grammar; the harmony constraint is scaled by the quality of a vowel as a potential target, parallel to Kimper's (2011) trigger strength scaling. This account can capture the complexities in the relationship between contrast and neutrality in a variety of harmony systems, including the gradience of neutrality (the height effect) in Hungarian (Hayes & Londe 2006), and paired neutral vowels in Mayak (Andersen 1999), among other cases. I argue that this view of harmony is necessary: neutrality is crucially about the quality of a vowel as a potential target of harmony, where target quality is determined in a cross-linguistically consistent, phonetically motivated way.

Lay summary

The goal of this dissertation is to establish and analyze the generalizations about which sounds can be exempt from a process called HARMONY, in which sounds within a word must match in some aspect of how they are pronounced. It is often the case that certain sounds, called NEUTRAL, are not subject to this requirement, and an ongoing question in linguistics is to determine which sounds are neutral and why. One of the most common generalizations made about neutral sounds is that they are those that lack a counterpart, produced in all the same ways except the aspect in which sounds must match for harmony. I discuss a number of languages in which this equivalence is false and provide a novel analysis, based on the relative drive for particular sounds to agree in particular properties.

Preface

A version of Chapter 3 has been published, as original, sole-authored work by me. Ozburn, Avery. 2019. A target-oriented approach to neutrality in vowel harmony: Evidence from Hungarian. *Glossa: a journal of general linguistics* 4(1): 47. 1–36. DOI: https://doi.org/10.5334/gjgl.681

A version of Chapter 5 has also been published, again as original, sole-authored work by me. Ozburn, Avery. 2019. A segment-specific metric for quantifying participation in harmony. In Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose & Matthew Zaslansky (eds.), *Supplemental Proceedings of the 2018 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America.

A version of Chapter 7 is in press as a proceedings paper, again as original, sole-authored work by me. Ozburn, Avery. To appear. An analysis of ATR harmony in Alur. In *Proceedings of the 2019 Canadian Linguistics Association Conference.*

All of the aforementioned chapters have been modified from their published form. The remainder of the dissertation is unpublished, original work by me.

Table of Contents

List of Tables

List of Figures

Acknowledgements

Thank you to my supervisor, Gunnar Hansson, my committee members, Douglas Pulleyblank and Anne-Michelle Tessier, and my former committee member, Carla Hudson Kam. You were the best committee I could have asked for, and this dissertation could not have happened without all of your advice, feedback, and support. Thank you as well to my external examiner, Sharon Rose, and my university examiners, Molly Babel and Patrick Moore, for their thoughtful and constructive comments, and to Paul Tupper for his insights on how to measure segment-specific harmony participation. I would also like to acknowledge the advice of many other faculty members, the support of my many fellow graduate students, and the stimulating conversations with a number of linguists too large to enumerate. Thank you to everyone who helped make this dissertation possible.

Chapter 1: Introduction

1 Background

VOWEL HARMONY is a phonological process in which vowels within a domain, typically a phonological word, must agree in some feature. For example, in Hungarian, vowels must agree in the feature [back], such that the dative suffix alternates between [-nɔk], with a back vowel, and [-nɛk], with a front vowel, depending on whether the vowels in the root are back or front respectively; examples include [haːz-nɔk] 'house-DAT' versus [tyːz-nɛk] 'fire-DAT' (Siptár & Törkenczy 2000). In vowel harmony systems, certain vowels can be NEUTRAL, such that they are not subject to this agreement requirement; an example in Hungarian occurs with the vowel [i], which can occur with back vowels in examples like [pɔpiːr-nɔk] 'paper-DAT' (Siptár & Törkenczy 2000). As expected, [i] can also occur with front vowels, in cases like [røvid-nɛk] 'short-DAT' (Siptár & Törkenczy 2000).

Vowel harmony is an extremely well-studied phonological phenomenon, and within that, a great deal of research focuses on neutral vowels, namely those that are exempt from undergoing harmony (e.g. Vago 1973; Goldsmith 1985; van der Hulst 1985; Ringen 1988; Archangeli & Pulleyblank 1994; Pulleyblank 1996; Ringen & Vago 1998; Kiparsky & Pajusalu 2003; Krämer 2003; Pulleyblank 2004; Archangeli & Pulleyblank 2007; Nevins 2010; Gafos and Dye 2011; Kimper 2011; Törkenczy et al. 2013; etc.). A considerable amount of research has been devoted to distinguishing between neutral vowels that are opaque, interrupting harmony, versus transparent, allowing harmony to pass through them (e.g. Kiparsky & Pajusalu 2003; Kimper 2011; Walker 2012; van der Hulst 2016), but less research has examined what determines whether a vowel is neutral at all, although it is a major focus of Archangeli & Pulleyblank (1994).

In the literature, an equivalence is often drawn between neutrality and lack of contrast for the harmonic counterpart, both explicitly as well as implicitly in the analysis. In other words, the idea is that a vowel will participate in harmony except when doing so is impossible due to a lack of harmonic counterpart in the inventory (e.g. Ringen and Vago 1998; Kimper 2011; Walker 2012; van der Hulst 2016). For example, in Finnish, the back vowels [a, u, o] and the front vowels [æ, y, ø] are three respective sets of harmonic counterparts, in that they differ only in the feature [back]; in contrast, there are no back vowels in the Finnish inventory that differ from the front vowels [i] and [e] in only this feature.

This assumption about the connection between neutrality and lack of contrast has been made since the beginning studies of neutrality within generative grammar. For example, Halle (1964) uses the neutrality of [i] to Finnish front/back harmony as an example of how non-contrastiveness for a given feature can allow a segment to be exempt from phonological patterns. Similarly, in many rule-based analyses of harmony (e.g. Lightner 1965; Vago 1973; etc.), neutrality is derived by having all vowels undergo harmony, followed by a process of absolute neutralization for vowels that are not contrastive for the harmonic feature. Applying this approach to the Finnish case, /i/ would harmonize to [ɯ] in back-harmonic words, but then the absence of a high back unrounded vowel in Finnish means that an absolute neutralization rule would then change [ɯ] back to [i]. Evidently, such an analysis is possible only if the harmonic counterpart of the neutral vowel cannot surface in the language.

This idea that neutrality and lack of contrast are equivalent is not limited to earlier literature; it is still present in a great deal of modern research. For example, van der Hulst (2016) draws a direct equivalence between neutrality and unpairedness in stating that "The non-advanced low vowel /a/ in Tangale misses a harmonic counterpart. This is what we call a neutral vowel." A version of this idea has also been advanced as part of the Contrastive Hierarchy (Dresher 2009), in which contrast is linked to feature specification; this type of approach can predict that segments without counterparts can participate in a harmony system, if they are appropriately specified, but not that segments with counterparts can be neutral.

Despite this pervasive view of neutrality, the connection between neutrality and lack of harmonic counterpart is false in both directions, and some literature, such as Archangeli & Pulleyblank (1994), has suggested divorcing the concepts. Vowels with counterparts may nonetheless be neutral; this occurs in Mayak (Chapter 9), where low vowels are contrastive for ATR but generally neutral to ATR harmony (Andersen 1999). Conversely, vowels without contrastive counterparts may nonetheless be targets of harmony, either through re-pairing (changing an additional feature; Baković 2000), or by participating allophonically. For example, in Alur (Chapter 7), there is no ATR low vowel in the language, but the RTR low vowel [a] undergoes harmony in some contexts by re-pairing to the mid vowel [e], which is otherwise the ATR counterpart of [ɛ] (Kutsch Lojenga 1991). In Mayak (Chapter 9), the mid RTR vowels do not have contrastive ATR counterparts, yet ATR mid vowels occur allophonically as the result of harmony (Andersen 1999).

This dissertation considers these types of complexities in the relationship between neutrality and lack of contrast, in a variety of front/back and ATR harmony systems, most of which are problematic for a traditional view of neutrality and most existing theoretical analyses of vowel harmony. I propose a novel theoretical analysis that is a revised way of viewing participation and neutrality in harmony. Under this approach, participation is based in part on vowel-specific asymmetries in the drive to undergo harmony. I formalize this idea by extending

Kimper's (2011) idea of scaling harmony constraints by their triggers to targets, and I show that this account can predict a number of patterns that are problematic in other approaches.

2 Central questions and main point

The central question posed by this dissertation is what participates in harmony and why. I look specifically at front/back harmony and ATR-dominant harmony, as these types of vowel harmony systems are widespread and tend to have particularly interesting target asymmetries in some well-described languages.[1] Despite all the literature that draws an equivalence between lack of contrast and neutrality, it is known that these two properties are independent in at least some languages, but what is less clear is how exactly neutrality does get determined in these cases, and how we can analyze such patterns within a theoretical framework. This dissertation aims to address such questions. It is worth noting that there is a match between lack of contrast and neutrality in many languages, and this connection is not merely a coincidence, because the same markedness constraints that create an inventory gap can, under the right conditions, also result in a vowel being neutral to harmony.[2] This dissertation does not argue that such an approach is wrong for the languages for which it works, but rather that such harmony-independent mechanisms cannot be the sole way in which neutrality is established.

The equivalence that has been drawn between neutrality and lack of contrast is really two separate claims, but is false in both directions: not all neutral vowels are non-contrastive for the harmonic feature, and not all vowels that are non-contrastive for the harmonic feature are neutral. Moreover, there are patterns independent of contrast in terms of which vowels are neutral, such

---

[1] Other types of harmonies do also show interesting target asymmetries, as described for instance for rounding harmony by Kaun (1995). I simply focus on front/back and ATR-dominant harmonies for reasons of scope.

[2] Of course, there is also the possibility that reports of neutrality are wrong, as is the case for Kinande, which was often reported to have a neutral low vowel that in fact participates in regressive harmony allophonically (Gick et al. 2006).

as the height effect that has been described for front/back harmony (e.g. Hayes & Londe 2006), in which higher vowels are more likely to be neutral. Such patterns are justifiable in terms of the phonetic properties of the relevant vowels.

In this dissertation, I argue that such patterns should be incorporated into formal analyses, and in fact need to be, in order to provide adequate accounts of a number of languages. While languages can (and sometimes do) use the same mechanisms to define both inventory structure and participation in the harmony system, and indeed such an approach may be the right analysis for a majority of harmony systems, this cannot account for all cases of neutrality. Instead, there must be an additional harmony-specific mechanism that other languages use in order to establish which segments participate in harmony. In other words, the penalty for a disharmonic form should depend on which vowels are interacting.

3 Formal proposal

In this thesis, I will argue for a specific formal proposal in which the vowel quality of the target is integrated into the penalty for disharmony. This analysis is situated in Harmonic Grammar (HG; Legendre et al. 1990a; b; Smolensky & Legendre 2006; Pater 2009), which is a version of Optimality Theory (OT; Prince & Smolensky 1993) in which constraints are weighted rather than ranked; the 'harmony score' of the candidate is determined by summing penalties across its constraint violations, and the winning candidate is the one with the highest harmonic score. Although any type of harmony constraint could work with the target-scaling approach, I adopt harmony-driving constraints adapted from Pulleyblank (2004), which penalize disharmonic V-V combinations, and which identify one vowel as the focus or 'target' and the other as the context or 'trigger'; violations are counted based on the focus. An example is *[RTR]∞[ATR], which assigns a violation for every RTR vowel (the focus/target, noted with an underline) followed at any distance by an ATR vowel (the context/trigger); I adopt this constraint over other possibilities because it has a relatively natural way of identifying which vowels are subject to trigger scaling versus target scaling. While I refer to the focus as the ‘target’ and the context as the ‘trigger’, this constraint can equally well be satisfied by harmony in the opposite direction; in other words, these are standard markedness constraints that penalize a configuration without requiring a particular solution.

In order to incorporate the target-based generalizations that are the focus of this dissertation, I adopt and extend the notion of ‘scaling factor’ from previous works like Kimper (2011). Scaling factors, in this usage, are constants based on specific properties of the violation; the base weight of a constraint is multiplied by any relevant scaling factors to determine the entire penalty for the violation. Kimper (2011) uses such scaling factors on the harmony constraint for trigger ‘strength’ (i.e. propensity of a particular vowel to trigger harmony) and trigger-target distance. In this dissertation, I extend this use to target ‘goodness’, or the susceptibility of a particular vowel quality to undergo harmony. In this way, the penalty incurred for disharmony depends on the particular vowel that could have undergone harmony but did not, such that the vowel-quality-based difference in the propensity of different vowels to undergo harmony is incorporated into the penalty assigned by the harmony constraint.

Considering the constraint *[RTR]∞[ATR], we can scale it for the combination of distance, trigger, and target, as will be done in Chapter 8 for Lango. Looking at targets specifically, suppose that the mid RTR vowel [ɛ] is a worse target than the high RTR vowel [ɪ], and is scaled as such in a given language like Lango. Then, for a given distance and trigger, disharmonic sequences in which [ɛ] is the potential target will score a lower penalty on

*[RTR]∞[ATR] than sequences in which [ɪ] is the potential target. In the analysis of Lango in Chapter 8, for instance, the weight of *[RTR]∞[ATR] is 2, the target scaling factor for [ɛ] is 2 and that for [ɪ] is 3; additionally, the trigger scaling of [u] is 4 and the distance scaling across multiple consonants is 0.5. Thus, an ɛCCu sequence, like in the Lango example [dɛ̀k-wú] 'stew-2PL', would be given a penalty of 2 (weight) * 2 (target scaling) * 4 (trigger scaling) * 0.5 (distance scaling) = 8 on *[RTR]∞[ATR]. In contrast, an ɪCCu sequence, as in the Lango example /nìŋ-wú/, which surfaces as [nìŋ-wú] 'name-2PL', would receive the higher penalty of 2 (weight) * 3 (target scaling) * 4 (trigger scaling) * 0.5 (distance scaling) = 12. These penalties are negative, such that if the weight of the faithfulness constraint is between 8 and 12, we expect harmony in the latter case but not the former. This is precisely what happens in Lango with a high back trigger across multiple consonants, and this scaling is therefore how I derive the high versus non-high target effect in such contexts in Chapter 8. With different scaling factors, it could alternatively be possible to expect harmony in both cases or in neither.

4 Typological predictions

In a theory without target quality scaling, a variety of possible harmony systems can be predicted by a combination of harmony (H) constraints, general markedness (GM) constraints against particular segment types (feature combinations), and faithfulness (F) constraints. Specifically, we can predict no harmony (when F>>H), harmony to contrastive vowels (when GM>>H>>F), or harmony to all vowels, either through allophonic participation or re-pairing of non-contrastive vowels (depending on the ranking of faithfulness to other features). Pulleyblank et al. (1995) provide a typology of this sort, in considering what happens with low vowels in a variety of African ATR harmony systems.

Once target scaling is added to the grammar, more possibilities are predicted. Specifically, target scaling is similar to using multiple harmony constraints, each with specified targets, with a fixed relationship among these harmony constraints. For example, scaling high targets above mid targets above low targets for an ATR harmony constraint would be effectively equivalent to having three ATR harmony constraints, one for high targets, one for mid, and one for low, with a fixed ranking in that order. The faithfulness and general markedness constraints can then be ranked/weighted anywhere within this harmony hierarchy. Specifically, then, target quality allows for asymmetries that are otherwise not possible, critically including the possibility for non-participation of contrastive vowels, if harmony targeting a given vowel is weighted lower than faithfulness. For example, if ATR faithfulness is weighted between harmony targeting mid vowels and harmony targeting low vowels, then low vowels will not harmonize for ATR, regardless of whether they are contrastive. This possibility does not exist in the framework that does not separate harmony by potential target.

5 Structure

This work is a manuscript-based thesis, in that each of the eight content chapters is designed to be able to stand alone as a separate paper; two have already been published as such, and a third is in press. The theoretical tools will be presented for each language, along with the data and other relevant argumentation. Most content chapters are about individual languages, and these are divided into two parts: one on front/back harmony, and the other on ATR-dominant harmony. However, the same style of analysis is adopted in both parts.

Following the introduction is Part 1, consisting of Chapters 2 through 5, which deals with front/back harmony. Chapter 2 is an introduction to the typology of front/back harmony, describing the generalizations related to this type of harmony, in particular in relation to which vowels participate. It also sets up the remainder of Part 1.

Chapter 3 is about Hungarian, in which there is no clear distinction between ‘neutral’ and ‘harmonic’ vowels, because front unrounded vowels can display behaviour that is intermediate between these categories, as both triggers and targets. This chapter also serves as an introduction to and argument for the type of formal analysis adopted throughout the dissertation, namely one in which the penalty for disharmony depends on which type of vowel is the potential target. I formalize this using a target scaling factor that reflects asymmetries in which vowels participate in harmony. Chapter 3 has previously been published in *Glossa* (Ozburn 2019a).

In Chapter 4, I show how the same tools used in the analysis of Hungarian can also help to analyze some facts from Votic that have been described as a ‘paradox’. In Votic, [i] behaves as [-back] for the purposes of conditioning lateral [l]~[ɫ] alternations, but is neutral to harmony. Adopting the same type of target scaling as for Hungarian provides a straightforward account, as neutrality does not have a direct relationship to contrast or feature specification in this approach. In this chapter, I also examine the participation of consonants in vowel harmony systems.

Chapter 5 differs from the other chapters, in that it is not about an individual language. Instead, it is about segment-specific participation, exploring the hypothesis that we might expect a vowel like [i], which is often neutral, to participate less in front/back harmony systems even in cases in which it is not categorically neutral. In this chapter, I develop a method of quantifying segment-specific participation in vowel harmony and apply it to lexical corpora of three languages with front/back harmony. It is primarily a proof-of-concept. This chapter has been previously published in the proceedings of AMP 2018 (Ozburn 2019b).

Following Chapter 5, I move into Part 2, consisting of Chapters 6 through 9, which deal with ATR-dominant vowel harmony. Chapter 6 is an introduction about ATR harmony, in the same vein as Chapter 2. This chapter describes generalizations related to ATR harmony, focusing on ATR-dominant harmony and on which vowels tend to participate in such systems. It also sets up the remainder of Part 2.

Chapter 7 analyzes harmony in Alur, in which there is a mostly 'normal' relationship between lack of contrast and neutrality, except that in certain contexts, unpaired [a] participates by re-pairing to a mid vowel. The factors that affect Alur harmony have been described in a complex way, and I reanalyze them as being purely phonological, relating to domains and distance. The facts of Alur are consistent with the type of analysis adopted throughout this thesis, although such formal devices are not actually necessary. It is nonetheless an important case to examine, as it illustrates one of the ways in which the relationship between contrast and participation can manifest, and it illustrates that even the 'normal' cases are consistent with a target-based approach. A version of this chapter will be published in the proceedings of CLA 2019 (Ozburn to appear).

Chapter 8 provides an account of Lango, in which harmony is determined by a complex set of interactions between triggers, targets, and distance, even though all vowels are contrastive for ATR. I adapt the analysis from Hungarian and Votic to ATR harmony, showing that the same type of height effect in targets, plus factors for triggers and distance, also provide an account of Lango. I also argue against existing accounts of Lango, particularly against a notion of 'heads' of ATR domains that requires reference to both input and output ATR values.

In Chapter 9, I provide an analysis of Mayak, in which low vowels are neutral despite being contrastive for ATR, while the non-contrastive mid vowels participate allophonically.

Such behaviour is clearly inconsistent with a view equating participation and neutrality, and I argue that most existing accounts could not deal with this pattern. However, it can be analyzed by extending the account of Lango developed in Chapter 8; Mayak simply draws the boundary among potential targets in a different place, distinguishing non-low versus low targets, instead of the high versus non-high distinction used in Lango.

Finally, Chapter 10 concludes, summarizing the entire dissertation and discussing some of the opportunities for future research that arise from this way of considering participation in harmony systems.

Chapter 2: Participation in front/back harmony

1 Introduction

In this chapter, I discuss the typology of participation in front/back harmony, discussing the types of asymmetries for which special theoretical tools are required. This chapter does not outline the formal analysis; the formalism that I will adopt in this dissertation is instead discussed in detail in the chapter on Hungarian, and refined through the chapters on other languages.

Front/back harmony is a phonological process in which co-occurring vowels (within some morphosyntactically or prosodically determined domain) must agree in the feature [back]. It is typically apparent both root-internally, as a static distributional pattern that governs attested versus unattested root shapes in the lexicon, as well as in alternations, in which affix vowels alternate in their front/back value depending on the value of the root. Some examples from Hungarian are given in (1) below; note how the suffix vowel is the back [ɔ] when the root contains the back vowels [a:,o,u:] in (a-c), while the same suffix contains the front vowel [ɛ] when the root contains the front vowels [y:,y,ø] in (d-f). (Length does not affect harmony in Hungarian, except where it changes vowel quality. Note that this suffix changes in rounding between the two variants; this is not a type of harmony but rather an effect of the idiosyncratic Hungarian vowel pairing of [ɔ] with [ɛ] as the only two short low vowels.)

(1) Front/back harmony (Siptár & Törkenczy 2000); back vowel roots in (a-c), front vowel roots in (d-f)

a.	ha:z-nɔk	house-DAT
b.	va:roʃ-nɔk	town-DAT

c.	kosoru:-nɔk	wreath-DAT
d.	ty:z-nɛk	fire-DAT
e.	tykør-nɛk	mirror-DAT
f.	ørøm-nɛk	joy-DAT

Front/back harmony is relatively well-studied, though much of the research focuses on a small number of languages. The systems of Finnish, Hungarian, and Turkish are well-known and quite thoroughly examined in both descriptive and theoretical literature (e.g. Ringen and Heinämäki 1999; Siptár & Törkenczy 2000 on Hungarian; Clements & Sezer 1982 on Turkish; among many others). Several other systems have been examined in the theoretical literature but are less thoroughly studied, such as Votic (e.g. Blumenfeld & Toivonen 2016; Hall 2018) and Uyghur (e.g. Halle, Vaux & Wolfe 2000; Vaux 2000; McCollum 2019); McCollum (2019) also examines harmony in Kazakh, Kyrgyz, and Uzbek. In general, front/back harmony tends to occur primarily in Uralic and Turkic languages, though there exist some cases analyzed as front/back harmony that have been described in languages outside of these families, such as Kaska (e.g. Hansson & Moore 2011). Notably, however, several cases of harmony originally analyzed as being front/back, in Mongolic and Tungusic languages, have been reanalyzed as being tongue root harmony (e.g. Rialland & Djamouri 1984; Svantesson 1985; Li 1996; see also Vaux 2009). Section 3 discusses the hypothesis from Vaux (2009) that Turkic shifted an original tongue root harmony system from Proto-Altaic, preserved in Mongolic and Tungusic, into a front/back one.

2 General patterns

2.1 Participation and neutrality

It is well-known that (non-low) front unrounded vowels like [i] or [e] are often neutral to front/back harmony systems; this pattern is true in two of the most widely studied languages: Finnish and Hungarian (e.g. Ringen and Heinämäki 1999 on Finnish; Siptár & Törkenczy 2000 on Hungarian). In these and other languages with this pattern, the inventories lack (non-low) back unrounded vowels such as [ɯ], [ɨ] or [ɤ]; traditional analyses connect the neutrality of front unrounded vowels with their lack of back harmonic counterparts in the phoneme inventory, in claiming that neutral vowels are those without counterparts and vowels without counterparts are neutral (e.g. Vago 1980a; Ringen & Vago 1998; etc.). The Hungarian vowel inventory, as an example of the unpairedness of front unrounded vowels, is given in Table 2.1, while some examples of the neutrality of Hungarian front unrounded vowels are given in (2). Note how in the examples, the final root vowel is front, which is disharmonic with both the rest of the root and the suffix.

	Front				Back			
	Unrounded		Rounded		Unrounded		Rounded	
	Short	Long	Short	Long	Short	Long	Short	Long
High	i	iː	y	yː			u	uː
Mid		eː	ø	øː			o	oː
Low	ɛ					aː	ɔ	

Table 2.1 Hungarian vowel inventory (Siptár & Törkenczy 2000)

(2) Transparency of front unrounded vowels (Siptár & Törkenczy 2000; Rebrus et al. 2012)

a.	pɔpi:r-nɔk	paper-DAT
b.	rɔdi:r-nɔk	eraser-DAT
c.	kuvik-nɔk	little owl-DAT
d.	ka:ve:-nɔk	coffee-DAT
e.	korde:-nɔk	cart-DAT
f.	ta:ɲe:r-nɔk	plate-DAT

This neutrality is particularly common in Uralic languages, which often have the relevant inventory asymmetries (e.g. Siptár & Törkenczy 2000 on Hungarian; Ariste 1968 on Votic; Ringen and Heinämäki 1999 on Finnish; Kangasmaa-Minn 1998 on Mari; etc.). In contrast, most Turkic languages are not particularly interesting from the standpoint of participation in front/back harmony, in that these languages typically have symmetric inventories for front/back, and all vowels participate (Johanson 2015). One notable and relatively well-studied exception is Uyghur, which has neutral front unrounded vowels and a lack of back unrounded vowels in the inventory (Halle, Vaux & Wolfe 2000; Vaux 2004; Johanson 2015). Uyghur thus patterns like Finnish and Hungarian in this respect, although there are additional complications to the Uyghur pattern due to raising of harmonic low vowels to neutral [i] in certain contexts (Vaux 2000). However, McCollum (2019) argues that Uyghur does not involve transparency.

In languages in which they are neutral, the front unrounded vowels are transparent, in that they are skipped over in determining harmony of further vowels (Kiparsky & Pajusalu 2003). Examples were shown in (2); the suffix takes a [+back] value following the back vowel that precedes the neutral one. It is generally accepted that these vowels are genuinely transparent; a number of phonetic studies have been conducted on Hungarian, which have found not found perceptible acoustic differences in neutral vowels depending on harmonic context (e.g. Blaho & Szeredi 2013). Nonetheless, Benus & Gafos (2007) did find some small articulatory differences, although given the acoustic and perceptual results that have not found such differences, it is reasonable to express skepticism about the phonological relevance of those results. I follow most literature on Hungarian in assuming that the vowels are truly transparent.

There is a known height effect in neutrality in front/back harmony, in that higher vowels are more likely to be neutral (e.g. Hayes & Londe 2006). This effect is present cross-linguistically; admittedly the research focuses on only Uralic and Turkic families, but there is an implicational generalization that if a lower front unrounded vowel like [e] is neutral to front/back harmony in a given language, then so is a higher one like [i] (e.g. Anderson 1980, Benus 2005, Finley 2015, etc.). The effect also exists within individual languages; for example, there is an effect in Hungarian in which the front unrounded vowels are neutral to different extents, with higher vowels more likely to be neutral in terms of factors like participating in harmonic suffix alternations (e.g. Hayes & Londe 2006). This will be discussed in more detail in Chapter 3, but some examples are given in (3), in comparison to the examples in (2) above. Specifically, while [i] is consistently transparent (2a-c), [e:] is usually transparent (2d-f) but may be variably opaque (3a), while [ɛ] is usually opaque (3b) but may variably be transparent (3c-d).

(3) Variability in neutrality (Siptár & Törkenczy 2000; Rebrus et al. 2012)

a.	ɔste:k-nɔk/nɛk	Aztec-DAT
b.	hɔvɛr-nɔk	pal-DAT
c.	ma:gnɛʃ-nɔk/nɛk	magnet-DAT

d. ko:dɛks-nɛk codex-DAT

Notably, this Hungarian pattern suggests that there is not a categorical distinction between which vowels are neutral and which are harmonic: the degree to which a given vowel participates in harmony can fall somewhere between these two extremes. This variability occurs in two senses: particular lexical items with the same front unrounded vowel can differ in whether they show transparency or not (e.g. (3b) versus (3d)), and the same root can show transparency sometimes but not others (e.g. (3c)).

There is additional evidence from Finnish of this height effect, as well as for the idea that the divide between 'neutral' and 'harmonic' is not categorical: high [y], which is harmonically paired with [u] and generally participates in harmony, is more likely to be (variably) neutral than mid [ø], which is paired with [o] (Ringen and Heinämäki 1999). Specifically, disharmony with [y] is more likely to be permitted in loan words. This ties in with the general height effect in neutrality in front/back harmony, and shows that even among paired vowels, there is not a categorical divide in which vowels participate and which are neutral. Instead, [y] and [ø] are both paired in Finnish, yet differ in their propensity to being neutral, with the higher one more neutral than the lower one. This could conceivably be explained in the same way as the height effect among front unrounded vowels: disharmony with lower vowels is tolerated less than disharmony with higher ones.

There is also a height effect in one non-Uralic/non-Turkic language with front/back harmony, namely the Athabaskan language Kaska (Hansson & Moore 2011). Kaska shows harmony through an [æ]~[a] alternation (or maybe [ɛ]~[a]). Particularly interesting is that mid [e] is neutral and transparent to this harmony; however, /e/ can be allophonically lowered to [æ] by a coda [h], and in such cases, it becomes a target of harmony. Thus, in Kaska, like in Uralic languages, lower vowels are more likely to be harmony targets, while higher ones are more likely to be neutral and transparent.

Certain types of consonants may participate in front/back harmony systems, notably laterals (as in Votic; Ariste 1968; Blumenfled & Toivonen 2016; see Chapter 4) and dorsal segments (as in Uyghur, where velars precede front vowels and uvulars precede back vowels; Hahn 1991). In many such cases, there is no contrast in the language between the relevant consonants, and their [back] value is determined entirely by the harmony system; this is the case for the Votic plain versus velarized lateral, as well as for velars versus uvulars in most Turkic languages, although they are contrastive before [i] in Uyghur (e.g. [kij] 'wear' vs. [qij] 'cut'; Lindblad 1990). In such cases, it can be said that the consonants are targets of harmony, but not triggers themselves; their underlying value is never apparent, as there is consistent agreement with the vowels (which do clearly contrast). However, in Turkish, there are some special cases in which palatalization of consonants is contrastive, and of course therefore not determined by harmony; these palatalized consonants can be disharmonic and can disrupt the propagation of harmony (Clements & Sezer 1982; Ní Chiosáin & Padgett 2001). A full discussion and treatment of Turkish is beyond the scope of this dissertation, but at least in the described cases of front/back harmony, it appears to be only in the exceptional cases (e.g. loans) that consonants have a role other than target in front/back harmony systems. A full typology of what types of consonants might interfere with front/back harmony (e.g. whether secondary palatalization is more likely to interfere than other types of consonants[3]) is beyond the scope of the present dissertation, but an important question for future research.

2.2 Other generalizations in front/back harmony

This dissertation focuses on patterns in participation and neutrality, but it is worth briefly considering the other types of generalizations that can be made about front/back harmony systems.

In both the Uralic and Turkic families, languages are suffixing, and front/back harmony propagates from root to suffix (Abondolo 1998 on Uralic; Johanson 2015 on Turkic). On the surface, then, front/back harmony is typically progressive, but the directionality can possibly be explained by morphological structure: these systems are root-controlled. Both epenthetic vowels and phonetic coarticulation studies have been used to address the directionality question. There is some discussion and debate about whether epenthetic vowels in loanwords in Turkish show regressive harmony, which would be suggestive of harmony not being solely progressive; however, Kabak (2011) expresses skepticism of the view that the epenthetic vowel is determined by harmony, since back [ɯ] breaks up certain types of initial CC clusters regardless of whether the following vowels are back or front. Moreover, in the Turkic language Tuvan, backness harmony affecting word-medial epenthetic vowels is strictly progressive, with vowel quality coming from the preceding vowel rather than the following one (Harrison 2000), suggesting that the progressive directionality is not just an artifact of suffixing morphology. In terms of phonetic studies, work on Finnish has suggested that coarticulation to neutral vowels differs progressively versus regressively, which has been argued to reflect a progressive directionality in the

---

[3] The reason we might suspect this is that there is a connection between palatalization and front/back harmony; for instance, Karaim palatalization consonant harmony has been hypothesized to have developed from Turkic vowel harmony (see e.g. Nevins & Vaux 2004).

phonological harmony system (Gordon 1999). On the other hand, Beddor & Yavuz (1995) found that in Turkish, regressive coarticulation is actually stronger than progressive coarticulation, exactly opposite to the phonological process of vowel harmony; it is unclear if this is relevant to the harmony system, as they suggest that it may be related to stress being (usually) word-final[4], so that the regressive coarticulation mostly involves stressed-to-unstressed. In other words, front/back harmony tends to be progressive, and it is not entirely clear whether this is simply an artifact of the morphology of Uralic and Turkic languages, or something deeper about the harmony.

While harmony is typically discussed in terms of suffix alternations for front/back harmony systems, the languages also have root-internal harmony, in which (non-neutral) vowels typically must agree in [back] (Siptár & Törkenczy 2000 on Hungarian; Ringen and Heinämäki 1999 on Finnish; Clements & Sezer 1982 on Turkish, though they also note there are numerous exceptions to root-internal harmony in Turkish). Compounds generally do not harmonize in Uralic and Turkic, although there is evidence in Turkish that compounds tend to be harmonic at a level greater than expected by chance; that is, the words chosen to make a compound tend to belong to the same harmonic class (Martin 2007). Recent loans are often not harmonized (e.g. Siptár & Törkenczy 2000 on Hungarian; Ringen and Heinämäki 1999 on Finnish; Clements & Sezer 1982 on Turkish), suggesting that root-internal harmony may no longer be an active part of the phonology of these languages, or that loans have their own stratum in the phonology; as such, some analyses adopt high-ranking root faithfulness and do not analyze root internal harmony (e.g. Ringen & Vago 1998). However, at least in Hungarian, disharmonic loans involving vowels that are usually harmonic are considered to be 'foreign' by native speakers, while disharmonic

---

[4] Though see more recent work such as Kabak & Vogel (2001) and Inkelas & Orgun (2002) on the complications of Turkish stress.

loans involving neutral vowels are not (Törkenczy 2010: 8). This suggests that native speakers do in fact note root-internal harmony as part of their phonology.

In the case of disharmonic loans in front/back harmony systems, the suffix [back] value is typically determined by the final non-neutral vowel of the base, as in Hungarian [ʃofør-nɛk] 'chauffeur-DAT'(Siptár & Törkenczy 2000 on Hungarian; Ringen and Heinämäki 1999 on Finnish; Clements & Sezer 1982 on Turkish). Nonetheless, in harmonic words, it could be claimed that harmony propagates from the initial vowel, given that harmony is progressive. Moreover, Kiparsky & Pajusalu (2003) note in their survey that a number of languages with front/back harmony, such as Votic and Seto, have a larger vowel inventory in initial syllables than non-initial ones; for instance, Kiparsky & Pajusalu (2003) note that Votic dialects and Southern Seto have [ø] and [ɨ] in initial syllables but not in non-initial ones. Such facts independently suggest that these languages have a preference for faithfulness in the first vowel (also the stressed syllable). In certain front/back harmony systems, such as Khanty, the stressed vowel may also be the relevant trigger from which harmony propagates, as in [saˈreenä] 'salmon-LOC' versus [ˈaaŋkeeja] 'mother-LAT' (Vaysman 2009). It is difficult to determine any broader role of stress in these harmony systems, given that a number of Uralic and Turkic languages have fixed stress, either initial (as in Hungarian and Finnish; e.g. Siptár & Törkenczy 2000 on Hungarian; Vroomen et al. 1998; Antilla & Boersma 2001 on Finnish) or final (as in the simplest case in Turkish; note that the phonology of stress in Turkish is complicated by morphological and lexical exceptions, but this has no consequences for the application of vowel harmony; e.g. Kabak & Vogel 2001, Inkelas & Orgun 2003).

To my knowledge, beyond arguments based on neutral front vowels, there is no basis for assuming that either value of [back] is the only active one in front/back harmony systems. This is in contrast to some tongue root systems in which ATR is clearly the active value, given that ATR anywhere in the harmony domain, including in an affix, will cause all other vowels in the domain to become ATR (e.g. Casali 2003).[5] In front/back harmony, as mentioned, the root controls harmony, and we consistently see agreement between the suffix and the non-neutral vowels in the root. There is no reason to suppose that only [+back] or only [-back] is the active feature in the system, and dominant/recessive harmony systems do not appear to exist for front/back harmony. However, there are some researchers who suggest that [Front] is the active feature in front/back harmony at least in Finnish and Hungarian (e.g. Goldsmith 1985); Hyman (2002: 6) notes, without explaining further, that it is not clear whether both values need to be active, marking Front harmony as existing but Back harmony as questionable.

Many languages with front/back harmony also have rounding harmony, particularly in Turkic languages, though this is also the case in some Uralic languages such as Hungarian and Eastern Mari (Abondolo 1998). Rounding harmony may be dependent on front/back harmony, in cases like Hungarian, where it occurs only among front vowels (Törkenczy 2010). However, front/back harmony is not typically dependent on any other type of harmony.

3 Harmony in the proto-language

Important in the understanding of front/back harmony is the extent to which the system can be reconstructed as a feature of the proto-language in Uralic and Turkic, and what form the harmony took. Specifically, information about how the harmony systems evolved to their present state can provide evidence on what the best participators in harmony are, which is relevant to the type of analysis to be developed here.

---

[5] It is possible that this is just a particular property of ATR systems, as it is unclear whether harmony in any other features behaves like this. In particular, rounding, front/back, and height harmony generally do not, with the exception of Moro, which might be derived from an ATR system (Ritchart & Rose 2017).

Vaux (2009) has suggested that Turkic front/back harmony developed out of tongue root harmony in Proto-Altaic. This argument assumes the Altaic hypothesis (Poppe 1960) in which Turkic, Mongolian, and Tungusic shared a common ancestor, which is not without controversy. Mongolian and Tungusic were originally thought to have front/back harmony, but based on more careful phonetic study of some of the relevant languages, more recent work (e.g. Svantesson et al. 2005) argues that in fact, it is tongue root harmony. Vaux (2009) suggests that the more articulatorily complex ATR/RTR distinctions, which he notes are not present in all languages (unlike [back] distinctions) in the proto-language evolved in Turkic to become a front/back distinction and harmony system instead; he argues that this makes sense as a way in which tongue root distinctions, which he claims are more marked than [back] distinctions[6], can be lost in a language. Particularly interesting here is that [i] was previously claimed to be neutral in the tongue root harmony in other Altaic branches; this is well known for Mongolian in particular. However, it has been shown at least for Halh/Khlakha Mongolian that [i] alternates allophonically with [ɪ] (Rialland & Djamouri 1984; Svantesson et al. 2005), although this may be gradient/phonetic coarticulation rather than categorical/phonological allophony, as Svantesson et al. (2015) characterize [i] as phonologically transparent to harmony. Exploring the connection between the behaviour of [i] in these Altaic tongue root harmony systems and in front/back harmony is an interesting avenue for future research.

In terms of the history of harmony in Proto-Uralic, the Proto-Uralic vowel system has been reconstructed as having a four-way contrast in high vowels, defined by backness and rounding, but only a two-way or three-way contrast in mid vowels, between [e] and [o] and

---

[6] This claim, however, is not entirely clear, as he does not provide evidence that making tongue root distinctions is in fact more common than making [back] distinctions involving vowels of the same rounding. In particular, one question is whether making rounding distinctions in front vowels is in fact more common than making tongue root distinctions. It is also not entirely clear what concept of the word 'marked' he is adopting, though he seems to reference articulatory complexity.

possibly a back unrounded mid vowel (Ultan 1973; Abondolo 1998). Many languages, such as Finnish and Hungarian, lost the back, high, unrounded vowel, as well as the back, mid, unrounded vowel if one was present (Ultan 1973), though others like Nganasan maintain a four-way alternation (Abondolo 1998). It has been proposed that harmony in Proto-Uralic began as a process that affected only low vowels; non-low vowels, even rounded ones like /u/ and /o/, were exempt from harmony, and opaque (e.g. Ultan 1973). Harmony was only later extended from low vowels to non-low ones. In some languages, like Finnish, this resulted in the creation of vowel qualities in non-initial syllables that had previously existed only in initial syllables. For instance, [ø] is reconstructed as existing only in initial syllables in Proto-Fennic (e.g. Viitso 1998), and still exists only initially in some daughter languages like Votic (e.g. Ariste 1968), but non-initial [ø] does exist in Finnish through an extension of the harmony system to target [o]. The fact that harmony began as a process only targeting low vowels is particularly interesting given that lower vowels appear to be better targets of front/back harmony, in the context of the height effect discussed above.

4 Why we might want to rethink how participation gets defined

On the surface, many front/back harmony systems look exactly like we might expect from a view in which participation correlates with having a counterpart in the harmonic feature. However, when we consider many of these systems in more detail, we find instead that participation is more complicated. In Hungarian, for instance, there is not a categorical divide between which unpaired vowels are neutral and which participate in harmony; the vowel [eː], for example, is intermediately neutral as a target, in that it sometimes alternates and sometimes does not in suffixes (e.g. Törkenczy 2010). This pattern and additional complications in Hungarian neutrality are discussed and analyzed in Chapter 3. In Votic, which is addressed in Chapter 4, [i] does not participate in front/back harmony, but does participate in other front/back processes in the language, suggesting that a lack of [back] specification, due to lack of contrast, cannot be the reason for its neutrality to harmony.

5 Hypothesis based on the typology

Given the consistency with which certain vowels tend to be neutral to front/back harmony, and the questions specific languages raise about whether this can truly be attributed to patterns of contrast, I will propose that individual vowel qualities (in a language-independent sense) vary in their propensity to be targets of harmony, independent of contrast. In other words, participation is not simply about whether a vowel has a harmonic counterpart within the phoneme inventory[7], but about how problematic disharmony is involving that vowel; whether a vowel contrasts for the harmonic feature is not necessarily relevant. In front/back harmony specifically, disharmony involving rounded vowels is more problematic than that involving unrounded ones, and disharmony involving lower vowels is more problematic than that involving higher ones. These generalizations are independently motivated by phonetic findings about articulation and perception of the front/back contrast in vowels of different rounding and height.

This predicts that we might see the same kind of variability in whether a vowel is a target of harmony in non-categorical ways too. As mentioned, while many front/back harmony systems do have neutrality, there are others that have been described as having all vowels participate. However, harmony systems tend to have at least some exceptional cases of disharmony within their lexicon, from sources like unassimilated loans, compounding, and non-alternating affixes. We might expect that vowels that are more likely to be categorically neutral (within and across

---

[7] The question of how harmonic counterpart gets determined is itself an often ignored question too.

languages) are also more likely to be gradiently neutral within a language, in that they participate less consistently in the harmony system and are more likely to be part of exceptional disharmonic co-occurrences. In the final chapter of this section on front/back harmony, I will develop a way of testing this hypothesis quantitatively and apply it to lexical corpora of three front/back harmony languages. In this dissertation, the purpose is a proof-of-concept and illustration of this tool, but it could be profitably applied in future research, for other cases of harmony, other types of data or corpora, or other specific questions.

6 Structure of the rest of the section on front/back harmony

The rest of this section is structured as follows. First is a chapter on Hungarian vowel harmony (previously published as Ozburn 2019a), which explains why the nuances of the relationship between neutrality and lack of contrast in Hungarian force a new view of participation in harmony. This chapter describes the theoretical framework to be used through the rest of the dissertation. Following the Hungarian, there is a chapter on Votic, which shows how the same kind of approach I argued to be necessary for Hungarian can also solve the 'paradox' in Votic harmony and lateral alternations (Blumenfeld & Toivonen 2016) in a very straightforward way. The final chapter in this section (previously published as Ozburn 2019b) deals with the hypothesis that the patterns we tend to see categorically in participation versus neutrality in front/back harmony should also be apparent gradiently, including in languages that on the surface seem to have a direct match between which vowels participate in front/back harmony and which are contrastive for front/back.

Chapter 3: A target-oriented approach to neutrality in vowel harmony: Evidence from Hungarian

1 Introduction

Widely discussed in the study of vowel harmony (e.g. Vago 1973; Goldsmith 1985; van der Hulst 1985; Ringen 1988; Archangeli & Pulleyblank 1994; Pulleyblank 1996; Ringen & Vago 1998; Kiparsky & Pajusalu 2003; Krämer 2003; Pulleyblank 2004; Archangeli & Pulleyblank 2007; Nevins 2010; Gafos and Dye 2011; Kimper 2011; Törkenczy et al. 2013; etc.), *neutral* vowels are those that are exempt from undergoing harmony. Neutral vowels are typically classified into two types: *opaque* and *transparent*, where the former type interrupts harmony while the latter allows harmony to pass through it (Archangeli & Pulleyblank 2007; Gafos and Dye 2011). While considerable attention has been paid to distinguishing between these types (e.g. Kiparsky & Pajusalu 2003; Kimper 2011; Walker 2012; van der Hulst 2016), less research has focused on why neutrality occurs at all. In much of the literature, either implicitly or explicitly, the standard assumption is that a vowel will participate in harmony except when doing so is impossible due to a lack of harmonic counterpart in the inventory (e.g. Ringen and Vago 1998; Kimper 2011; Walker 2012; van der Hulst 2016). For example, van der Hulst (2016) draws a direct equivalence between neutrality and unpairedness in stating that “The non-advanced low vowel /a/ in Tangale misses a harmonic counterpart. This is what we call a neutral vowel.”

However, the connection between neutrality and incompatibility with the harmonic feature is false in both directions. Vowels that are perfectly compatible with the harmonic feature may nonetheless be neutral, as in Mayak (Chapter 9), where low vowels contrast in [ATR] (as shown by Andersen (1999), who provides evidence that it is an ATR alternation by showing a low vowel suffix that does alternate) but do not undergo tongue root harmony in situations in which other vowels do (Andersen 1999). Similarly, vowels that are incompatible with the harmonic feature may nonetheless undergo harmony through a change in an additional feature, in a process known as *re-pairing* (Baković 2000). For example, in Maasai, [a] has no [ATR] counterpart, yet undergoes progressive ATR harmony by re-pairing with [o] (Archangeli & Pulleyblank 1994; Baković 2000). Alternatively, unpaired vowels may alternate allophonically, as in Kinande, without necessitating an additional feature change (e.g. Archangeli & Pulleyblank 2002).

Even in Hungarian, one of the most studied languages for neutrality and its connection to unpairedness, it has been noted that neutrality is not as simple as the traditional view might suggest, and that this view is fundamentally flawed. For example, Törkenczy (2013) notes problems with many common claims about neutrality, including that some Hungarian neutral vowels are not harmonically unpaired, and that whether a given neutral vowel is opaque or transparent can depend on whether it occurs in an alternating or invariant suffix. Similarly, Siptár (2015: 3), in discussing Hungarian, notes that the concept of neutral vowel is "variable, gradual, and context-dependent". Benkő et al. (2017) further note that it is a problem for Hungarian that traditional approaches consider neutrality to be categorical and that gradience (in the sense of vowels differing in their phonological degree of neutrality, with morpheme-specific neutrality versus participation) is often unintegrated into the phonological analysis. In front/back harmony more broadly, Kiparsky & Pajusalu (2003) mention in their typology that there are multiple possible reasons for a vowel to be unable to alternate harmonically.

This chapter proposes a revised analytical and intuitive framework for viewing participation and neutrality in harmony, by considering a variety of data from Hungarian that are problematic for a traditional view of neutrality and most theoretical analyses. For example, the vowel [e:], while generally transparent in roots, alternates harmonically in some suffixes by re-pairing with [a:]. I argue that this and other patterns in Hungarian necessitate a new view of participation in harmony, in which the main cause of neutrality is not necessarily incompatibility with the harmony feature, but rather a weak vowel-specific drive to undergo harmony. Articulatory and perceptual facts, in terms of the space along the front/back dimension that vowels can move around without changing category, suggest that vowels that are rounded and/or low are better participators in front/back harmony than vowels that are non-low and unrounded (e.g. Benus 2005; Gafos and Dye 2011). These properties give independent motivation for the height-based vowel-specific asymmetries in the requirement to undergo harmony. I formalize this view by extending Kimper's (2011) trigger-based analysis of harmony to targets, showing that such an account can capture the basic facts of Hungarian, in addition to predicting a variety of other patterns that are problematic in other analyses. Moreover, the inherent gradience of this approach allows for integration into the account of the non-categorical divide between harmonic and neutral vowels, and allows for the incorporation of a morpheme-specific component for lexically determined behaviour.

The chapter is organized as follows. Section 2 describes the facts about Hungarian harmony, and Section 3 details the phonetic motivations for the relevant patterns. Section 4 discusses the implications that the Hungarian data has for the view of neutrality and why it is problematic to previous accounts. Section 5 presents the target-oriented approach, while Section 6 applies it to the Hungarian data. Section 7 describes some further implications of target-focused harmony for other languages, and Section 8 concludes.

2 Hungarian harmony

This section overviews the facts about Hungarian vowel harmony. Section 2.1 discusses the vowel inventory and basic harmony pattern. Section 2.2 describes the patterns of neutrality within roots, while Section 2.3 deals with neutrality in suffixes. Section 2.4 summarizes the patterns in terms of gradience in triggers and targets of harmony.

2.1 Vowel inventory and basic harmony pattern

The Hungarian vowel inventory, with feature representations, is given in Table 3.1.

	Front [-back]				Back [+back]			
	Unrounded [-round]		Rounded [+round]		Unrounded [-round]		Rounded [+round]	
	Short	Long	Short	Long	Short	Long	Short	Long
High [+high,-low]	i	i:	y	y:			u	u:
Mid [-high,-low]		e:	ø	ø:			o	o:
Low [-high,+low]	ɛ					a:	ɔ	

Table 3.1: Hungarian vowel inventory (Siptár & Törkenczy 2000)

Phonetic descriptions have noted that [a:] is produced as central or even front, but it patterns phonologically as a back vowel and has been argued to be phonologically classified as such (e.g. Gósy & Siptár 2015). Throughout this chapter, I will consider [a:] a back vowel, despite its phonetic implementation. It is worth noting that for all other vowels, the distinction between front and back is synchronically based in the phonetics; the phonetics does have bearing on the phonological patterning, despite the phonetics/phonology mismatch in the behaviour of [a:].

By some researchers, [e:] has been analyzed as being at least sometimes a phonologically low vowel, due to its alternations with [ɛ], [a:], and [ɔ]. For example, Siptár & Törkenczy (2000: 158) class non-alternating [e:] as mid in their feature system, but alternating [e:] emerges as a low vowel phonologically in their analysis; in standard Hungarian, however, non-alternating [e:] and derived [e:] are the same vowel. However, I follow the phonetic inventory in Siptár & Törkenczy (2000: 52), Hayes & Londe (2006), and Törkenczy (2010), among many others, in categorizing it as mid. Given the height effect among front unrounded vowels in Hungarian harmony, to be discussed in Section 2.2, the distinction between [e:] and [ɛ] in height is relevant to the phonological system. In general, as Törkenczy (2010) notes, the system requires two degrees of backness and three of height. In the present analysis, I assume that the crucial factor of height is determined based on the phonology, though of course that is founded in the phonetics.

Notable in this inventory are the gaps where vowels do not have a counterpart that differs only along the front/back dimension. Only non-low rounded vowels have direct harmonic counterparts; the front unrounded vowels and the low vowels ([i], [i:], [e:], [ɛ], [a:], and [ɔ]) do not have counterparts that differ only in the feature [back].[8] Given that short, rounded, back vowels have a three-way height distinction, both [high] and [low] are active features within the Hungarian phonological system, assuming these are the relevant features for characterizing

---

[8] However, there is ambiguity in the identity of the vowel [ɔ], with some researchers using the symbols [ɒ], [ɑ], or [a], suggesting that this vowel may not in fact be rounded; it is unclear whether these researchers mean that the vowel is rounded or not. If it is not, then [ɛ] and [ɔ] are a direct harmonic pair. Whether they are or not has no bearing on the analysis here.

vowels.[9]

Hungarian has an extensively studied system of front/back harmony, as well as a more limited system of rounding harmony that will not be discussed here (e.g. Vago 1973; Goldsmith 1985; van der Hulst 1985; Ringen 1988; Ringen & Vago 1998; Siptár & Törkenczy 2000; Hayes & Londe 2006; Nevins 2010; Törkenczy 2010; Gafos and Dye 2011; Kimper 2011; Törkenczy et al. 2013; etc.). In general, disregarding the front unrounded vowels, vowels that disagree in the feature [back] are not permitted to co-occur within (native) roots, and suffixes alternate according to the [back] value of the root vowels. Examples are given in (1), with the dative suffix [nɔk]~[nɛk].[10]

(1) Front/back harmony (Siptár & Törkenczy 2000); back vowel roots in (a-c), front vowel roots in (d-f)

a.	ha:z-nɔk	house-DAT
b.	va:roʃ-nɔk	town-DAT
c.	kosoru:-nɔk	wreath-DAT
d.	ty:z-nɛk	fire-DAT
e.	tykør-nɛk	mirror-DAT
f.	ørøm-nɛk	joy-DAT

---

[9] Following a Contrastive Hierarchy style approach (Dresher 2009), we could claim that [low] is not contrastive or specified among unrounded vowels, which would allow [e:] and [a:] to be direct harmonic counterparts. Nonetheless, doing so would not resolve the issue of their asymmetrical behaviour (to be discussed in Section 2.3), and would require [ɛ] and [e:] to be of the same height, which is inconsistent with their distinct phonological patterning (see e.g. Ringen & Vago 1998 for discussion and references about [ɛ], and Benus 2005 for arguments that the distinction results from the difference in height).

[10] Note that this alternation involves re-pairing, since neither vowel has a direct harmonic counterpart.

2.2 Neutrality in roots

The front unrounded vowels are generally considered transparent to this harmony; they can co-occur in roots with any other vowel, and are skipped over in determining the harmonic value of the suffix (Siptár & Törkenczy 2000; Törkenczy 2010). Thus, suffixes generally surface as back when they attach to stems that contain a back vowel followed by a (non-low) front unrounded vowel. Examples are shown in (2a-f). While the high front unrounded vowels [i] and [iː] are consistently transparent[11] (as long as there is only a single one; see below), the mid front unrounded vowel [eː] shows variable behaviour; some back-[eː] stems show consistent transparency (2d-f), while others are able to take either front or back stems ("vacillating"), as in (2g) (Törkenczy et al. 2013). The low front unrounded vowel [ɛ] shows an even greater degree of variable behaviour and is sometimes considered front harmonic (e.g. Ringen 1975; 1978; 1980); while there exist some back-[ɛ] stems that take only back suffixes (2h), most such stems either vacillate (2i) or take only front suffixes (2j) (Törkenczy et al. 2013). This is a morpheme-specific property: different stems have different preferences for front or back suffixes, and the choice is lexically arbitrary (Hayes et al. 2009; Rebrus et al. 2017). Moreover, although a single [i] or [iː] is always transparent, multiple high front unrounded vowels are variably opaque, as shown in (2k).

---

[11] A reviewer points out that in very rare cases, a single [i] may be opaque, as in "abszint-nak/nek" and "Auschwitz-ból/ből", but also notes that these words may be analyzed as complex, as [sint] and [vitsː] are both independent words. I will assume that the variability possible in these words comes from ambiguity in their morphological structure, rather than from any phonological factors. In other words, speakers may or may not treat these words as compounds. When they are treated as compounds, the back vowel is outside of the harmony domain, allowing the possibility of front suffixes. In contrast, when they are not treated as compounds, they behave like other back-[i] words, in taking back suffixes.

(2) Transparency of front unrounded vowels (Siptár & Törkenczy 2000; Rebrus et al. 2012)

a.	pɔpi:r-nɔk	paper-DAT
b.	rɔdi:r-nɔk	eraser-DAT
c.	kuvik-nɔk	little owl-DAT
d.	ka:ve:-nɔk	coffee-DAT
e.	korde:-nɔk	cart-DAT
f.	ta:ɲe:r-nɔk	plate-DAT
g.	ɔste:k-nɔk/nɛk	Aztec-DAT
h.	hɔvɛr-nɔk	pal-DAT
i.	ma:gnɛʃ-nɔk/nɛk	magnet-DAT
j.	ko:dɛks-nɛk	codex-DAT
k.	ɔnɔli:ziʃ-nɔk/nɛk	analysis-DAT

The distinction between the degree of transparency for high, mid, and low vowels, in terms of how consistently they are transparent, is known as the height effect; it is strongly supported by a large body of data and wug test judgements from native speakers (Hayes & Londe 2006). Essentially, [i(:)] is more transparent than [e:], which is more transparent than [ɛ]; this "transparency scale" is supported by ratios of front suffixation to all suffixation for back-neutral stems calculated from corpus counts (Rebrus & Törkenczy 2016b), with total transparency for [i(:)], a low ratio for [e:], and a high ratio for [ɛ]. The example in (2k), showing that multiple neutral vowels are less transparent, illustrates what is known as the count effect, and also applies to the other front unrounded vowels; again, this effect has been shown to be productive and robust (Hayes & Londe 2006). The height and count effects also combine, such that if a back vowel is followed by at least two neutral vowels, there is vacillation (either front or back suffixation) with [i(:)], vacillation or front suffixation with [e:], and only front suffixation with [ɛ] (Siptár & Törkenczy 2000).

While (2) shows that (non-low) front unrounded vowels do not generally trigger harmony when other potential triggers are present, they are "last-resort" triggers, in that roots with only front unrounded vowels typically take front suffixes (e.g. Siptár & Törkenczy 2000).[12] Examples are shown in (3), with the dative suffix seen in previous examples and the ablative suffix [tø:l]~[to:l]. It is worth noting that within roots, neutral vowels can precede back vowels, as in [bikɔ] 'bull'; the pattern in (3) occurs solely across morpheme boundaries, in that root-internal harmony for front unrounded vowels is not required even in native vocabulary.

(3) Last-resort triggering by front unrounded vowels (Siptár & Törkenczy 2000)

a.	vi:z-nɛk	water-DAT
b.	vi:s-tø:l[13]	water-ABL
c.	sɛge:ɲ-nɛk	poor-DAT
d.	sɛge:ɲ-tø:l	poor-ABL

I will consider these examples in (3) to be last-resort triggering, as opposed to a default value, because it is clear from vacillating stems that front unrounded vowels are able to trigger front

---

[12] While Siptár & Törkenczy (2000) note that most neutral-only roots behave like [vi:z] in taking front suffixes, there is a small class of about 50-60 exceptions to this generalization ("anti-harmonic" stems), exemplified by *híd+nak* (see e.g. van der Hulst 1985; Ringen and Vago 1998). A full treatment of these roots is beyond the scope of this dissertation, though I offer some possibilities for analysis in Section 6.5. Most of these roots have [i] or [i:]; a few have [e:], in keeping with the general height effect.

[13] Note that Hungarian has obstruent cluster voicing assimilation; the root-final consonant in this form assimilates in voicing to the suffix-initial [t].

harmony in some contexts; the same can be thought to apply here.[14] In new loans and nonce words, front unrounded vowels also always take front suffixes, providing an additional argument for them to behave as triggers.[15] As I adopt binary features (arbitrarily), I also assume for simplicity that both [+back] and [-back] are active in triggering harmony in Hungarian, as has been widely assumed (e.g. Vago 1973; 1976; Farkas & Beddor 1987).

Whether Hungarian has root-internal harmony is a matter of some debate. Many analyses assume that it does not, and that harmony only affects suffixes; in analyses of this type formulated in OT, a high-ranked root faithfulness constraint is assumed (e.g. Ringen & Vago 1998). However, aside from recent loans, Hungarian roots obey similar harmonic restrictions to root-suffix combinations: Törkenczy et al. (2013) note that bisyllabic stems containing both front and back vowels are rare with front rounded vowels ("disharmonic roots"), with all cases being recent loans, but frequent with front unrounded vowels ("mixed roots"), though with [ɛ] most are loanwords. This difference corresponds to the strong versus weak disharmony distinction drawn by Rebrus & Törkenczy (2015). Törkenczy (2010: 8) provides corpus counts supporting this distinction and notes that intuitively for native speakers, disharmonic roots "feel foreign", while mixed roots do not; experimental data on this distinction does not yet exist, to my knowledge. In other words, front rounded vowels cannot co-occur with back vowels within roots, while mixed roots containing front unrounded vowels with back vowels are permitted. This fact

---

[14] An additional argument, well-known but most often rejected in the recent literature, comes from the behaviour of personal pronouns that are formed by adding person suffixes to case suffixes. The front/back value of the vowels in such pronouns is consistent for a given case, but varies across cases. For example, the ablative case [toːl]~[tøːl] is front in personal pronouns, as in [tøːl-ɛm] 'ABL-1SG'; in contrast, the delative case [roːl]~[røːl], in which the same vowels alternate, is back in personal pronouns, as in [roːl-ɔm] 'DEL-1SG'. Vago (1973; 1980a) argues that such forms, in absence of a stem, allow the underlying suffix vowel to surface and therefore provide evidence for its feature value. However, this argument is generally rejected; most researchers assume that personal pronouns are formed from independent stems that look like alternating suffixes, but whose vowels have no bearing on the question of underlying suffix vowels (e.g. Ringen 1978). Moreover, a reviewer points out that most suffixes do not have a stem form, so Vago's argument is at most applicable to a small number of suffixes.

[15] Thank you to a reviewer for this information.

suggests that, disregarding recent loans, Hungarian roots do show a harmony pattern that should be analyzed, because it works in essentially the same way as the root-to-suffix harmony.

To summarize, Hungarian has a pattern of front/back harmony, in which front unrounded vowels tend to be transparent, but can trigger harmony in certain stems or when there are multiple of them, and behave as last-resort triggers when they are the only vowel in the root (although they can fail to trigger for some roots, as discussed previously). The different front unrounded vowels behave differently with respect to various aspects and degrees of neutrality, with height affecting whether a single neutral vowel is consistently transparent ([i(:)]), usually transparent ([e:]), or usually variable or opaque ([ɛ]). The back vowels and front rounded vowels are generally harmonic, and typically occur only with neutral vowels or vowels in the same harmonic class.

2.3 Patterns of neutrality in suffixes

I turn now to additional complications from suffix alternations and their interaction with neutrality. Törkenczy et al. (2013) discuss the fact that there are both invariant and alternating suffixes containing neutral vowels, but their analysis focuses on paradigm uniformity, not the issue of neutrality. Other literature also notes the presence of alternating suffixes as a factor in the scale of neutrality (e.g. Rebrus & Törkenczy 2016a), but again does not provide any detailed theoretical implementation.

The Hungarian inventory in Table 3.1 showed that only non-low rounded vowels are paired for [back], but other vowel pairs also alternate harmonically in suffixes. An example is the dative suffix in (1-3), since neither [ɔ] nor [ɛ] is paired. Siptár & Törkenczy (2000), among others, provide a complete list of alternating vowel pairs, shown with examples in Table 3.2.

Back vowel	Front vowel	Example
u:	y:	la:b-u: ‘legged’ fɛj-y: ‘headed’
u	y	ha:z-unk ‘our house’ kɛrt-ynk ‘our garden’
o:	ø:	va:r-o: ‘waiting (adj.)’ ke:r-ø: ‘asking (adj.)’
o	ø, ɛ (varies based on rounding harmony; recall that there is no short [e], so that a rounding alternation here also results in a change in [low])	ha:s-hoz ‘to (the) house’[16] føld-høz ‘to (the) land’ kɛrt-hɛz ‘to (the) garden’
a:	e:	va:r-na: ‘he would wait for it’ ke:r-ne: ‘he would ask for it’
ɔ	ɛ	ha:z-bɔn ‘in (the) house’ kɛrd-bɛn ‘in (the) garden’[9]

Table 3.2: Suffix alternations in Hungarian (Siptár & Törkenczy 2000)

Despite the large literature on Hungarian vowel harmony, the suffix alternations between [a:] and

---

[16] As in (3b), voicing assimilation from the suffix consonant applies here to the root-final obstruent.

[e:], in the penultimate row of Table 3.2, are rarely explicitly analyzed theoretically. However, this pair is problematic for the widely stated view that neutral vowels are harmonically unpaired, which would suggest that they should be non-alternating (e.g. van der Hulst 2016; see also Törkenczy 2013 on this issue). Under categorical definitions of neutrality, [e:] (unlike [ɛ]) is widely considered a neutral vowel in Hungarian (e.g. Vago 1978, Ringen 1988; Siptár & Törkenczy 2000), but Table 3.2 shows that it does participate in suffix alternations. Indeed, [a:]~[e:] alternations occur regularly and consistently through Hungarian morphology, in a variety of suffixes. Like for other pairs in alternating suffixes, the front vowel [e:] occurs after stems with front vowels, while the back vowel [a:] occurs following stems with back vowels (Siptár & Törkenczy 2000). Moreover, [e:] is the only possible counterpart for the harmonic vowel [a:] in alternating suffixes (cf. Table 3.2). Additional examples of the [a:]~[e:] alternation are shown in (4) with four different suffixes, namely the adessive, translative, 3sg definite conditional, and 2sg indefinite conjunctive/imperative; Törkenczy (2010) and Rebrus et al. (2012) note 10 productive suffixes containing this alternation.[17] Note that these are native Hungarian suffixes.

---

[17] Other sources differ in the number; for instance Törkenczy et al. (2013) note only 8 suffixes with this alternation.

(4) Suffix alternations with [a:] and [e:] (Törkenczy 1997)

	Word	Gloss
a.	dob-na:l	drum-ADE
b.	tøk-ne:l	pumpkin-ADE
c.	dob-ba:	drum-TRANSL[18]
d.	tøk-ke:	pumpkin-TRANSL
e.	ɔd-na:	give-3.SG.DEF.COND
f.	øl-ne:	kill-3.SG.DEF.COND
g.	ɔɟ:a:l (ɔd-ja:l)[19]	give-2.SG.INDF.CONJ/IMP
h.	øj:e:l (øl-je:l)[12]	kill-2.SG.INDF.CONJ/IMP

Despite such alternations, there is a distinction between [e:] and [a:] in terms of the number and nature of non-alternating suffixes with these vowels. Siptár & Törkenczy (2000) list twelve non-alternating suffixes with [e:], comparable to the number with [i]. These are common suffixes that are fundamental to the grammatical system of Hungarian, such as the possessive and the causal; implicit in Siptár & Törkenczy's (2000) discussion of these affixes is that they are native to Hungarian. In contrast, there are only three non-alternating suffixes with [a:], comparable to the four containing [o]; these suffixes are borrowed, mostly Latinate, and their exact semantic content is not always clear from the glosses. Given this and the fact that all three end in [a:l], they may in fact all be the same suffix. [20] Examples of non-alternating [e:] suffixes are given in (5),

---

[18] The translative suffix is [va:]~[ve:] with vowel-final stems; the /v/ assimilates to the root-final consonant in consonant-final stems like the examples given here (Siptár and Törkenczy 2000).

[19] The combination of [d]+[j] and [l]+[j] in these examples cause consonant changes; the examples in parentheses indicate the morpheme-by-morpheme transcription, while outside the parentheses are the combined surface forms.

[20] A reviewer points out that this is indeed considered one suffix, and it is lexicalized and not productive, because it truncates stems and/or shortens the stem vowels; these examples in (6) might be taken to be monomorphemic.

and a complete list of the potential non-alternating [a:] suffixes is given in (6).

(5) Non-alternating suffixes with [e:] (Siptár & Törkenczy 2000)

	Word	Gloss
a.	la:ɲ-e:	girl-POSS
b.	hɔza:-e:rt	(one's) country-CAUS
c.	kova:tʃ-ne:	Kovács (surname)-Mrs.
d.	ɔd-ne:k[21]	give-1.SG.INDF.COND
e.	kultʃ-ke:nt	key-ESS

(6) Non-alternating suffix(es) with [a:] (Siptár & Törkenczy 2000)

	Word	Gloss
a.	rɛtsɛnz-a:l	review (verb)
b.	ɛlɛktr-ifika:l	electrify
c.	fɛtiʃ-iza:l	make a fetish of

As mentioned, the status of [a:] as a harmonic vowel and of [e:] as a transparent vowel in Hungarian is widely assumed; this is supported by the data on non-alternating suffixes in (5-6), because [e:] is neutral in a variety of suffixes, while [a:] is not (see Ringen 1978 for argumentation connecting neutrality to non-alternating suffixes in Hungarian). However, the data in (4) clearly illustrate that [a:] and [e:] are a harmonic pair in suffix alternations.

Törkenczy (2010), in listing vowels occurring in invariant suffixes in Hungarian, does not list any with [a:], suggesting that he may consider them monomorphemic.

[21] However, a reviewer points out that for many speakers, this suffix is alternating, with the example ɔd-naːk 'give-1.SG.INDEF.COND', compared to fɛd-neːk 'cover-1.SG.INDEF.COND'. For these speakers, this suffix simply patterns like those in (4).

Older work has at least implicitly assumed that the underlying vowel in [a:]~[e:] alternations is /a:/, with /e:/ in suffixes surfacing faithfully as a non-alternating neutral vowel. In contrast, in more recent literature on Hungarian, the presence of a neutral vowel in alternating suffixes is considered an aspect of gradient neutrality (e.g. Rebrus & Törkenczy 2016a). As illustrated in Table 3.2, there are no alternating suffixes with high front unrounded vowels; these vowels occur only in invariant suffixes (Törkenczy 2010; Törkenczy et al. 2013).[22] The data in (4-5) has shown that [e:] is present in both invariant and alternating suffixes. Table 3.2 also shows that [ɛ] occurs in alternating suffixes; however, it never occurs in invariant suffixes (Törkenczy 2010; Törkenczy et al. 2013). Thus, in this respect too, there is a height effect: high [i] and [i:] are the most neutral, in that they occur only in invariant suffixes; low [ɛ] is the least neutral, in that it occurs only in alternating suffixes; and mid [e:] is intermediate. Under a gradient view of neutrality, this pattern is sensible and consistent with the patterns in Section 2.2, but it is problematic under a categorical view of neutrality, and it needs to be analyzed.

2.4 Summary of patterns

In Hungarian, the distinction between neutral and harmonic vowels is not always clear, and there are multiple distinct aspects and degrees of neutrality; the front unrounded vowels show a variety of behaviours, as shown in Sections 2.2 and 2.3. This has been noted extensively in recent literature on Hungarian. For example, Rebrus & Törkenczy (2016a) note three different notions that are often included as part of neutrality: (i) neutrality as a target (invariant suffixes; see (5)), (ii) neutrality as a trigger with respect to other targets (transparency; see (2)), and (iii) neutrality as the source trigger (anti-harmony; see footnote 5). In Hungarian, these three properties create a scale of neutrality among the front unrounded vowels, which show [i(:)] to be the most neutral,

---

[22] [i] alternates in only a single suppletive suffix [-i]/[-jɔ] '3.SG.DEF' (Rebrus et al. 2012)

followed by [e:], and finally [ɛ].

In slightly different terms, we can summarize the behaviour of Hungarian vowels based on whether and when they are triggers or targets of harmony. As is evident here, there are two primary factors that can affect whether a given vowel is a trigger or target of harmony: the specific vowel quality and the specific morpheme in which the vowel appears (see Rebrus et al. 2012). Additionally, harmony is affected by whether a vowel is in a root or suffix and how many neutral vowels there are.

Vowel	Trigger (root-to-suffix)?	Target (in suffix)?
i/i:	Only if there is no alternative or (variably) when multiple i/i: are present	Never
e:	Variably in a few morphemes, or if no alternative is present	Sometimes
ɛ	In most morphemes (often variably)	Always
All others	Always	Always

Table 3.3: Trigger and target status of Hungarian vowels

Despite the discussion of these issues of variability in the literature, theoretical treatments of these issues are lacking. In the present chapter, I focus primarily on one aspect of differential neutrality, namely the non-neutrality of [e:] in suffixes. As will be discussed in more detail in the analysis section, the way of understanding harmony targets presented here ties in closely with the view of neutrality as gradient and variable that is dominant in modern literature on Hungarian (e.g. Rebrus et al. 2012; Törkenczy et al. 2013; Rebrus & Törkenczy 2016a).

3 Motivations for the asymmetries

Section 2 has shown that Hungarian neutrality has a height effect, with variability in the degree of neutrality depending on the height of the vowel. The lowest front unrounded vowel [ɛ] is sometimes considered harmonic (e.g. Ringen 1975; 1978; 1980; Ringen and Vago 1998); it is less transparent than higher [e:], which is in turn less transparent than high [i, i:]. This generalization holds across multiple facets of neutrality, and is clearly an active part of Hungarian phonology.

This height effect is a manifestation of a broader, cross-linguistic pattern in front/back harmony, namely that lower vowels are more likely to undergo this type of harmony (e.g. Anderson 1980, Benus 2005, Finley 2015, etc.). Within front unrounded vowels, which are commonly transparent to front/back harmony, the most common transparent vowel is [i], followed by [e]; moreover, there is an implicational generalization that if a lower vowel like [e] is transparent, then so is [i] (e.g. Anderson 1980, Benus 2005, Finley 2015, etc.). This height asymmetry also extends beyond front unrounded vowels: in loanwords in Finnish, [y] is often transparent to front/back harmony, while the lower front rounded vowel [ø] is more consistently harmonic (Ringen and Heinämäki 1999).

These cross-linguistic and Hungarian-internal patterns are not surprising, given that there is phonetic motivation for this height asymmetry. As argued for instance by Benus (2005), Benus & Gafos (2007), and Gafos & Dye (2011), higher vowels have a greater degree of perceptual stability along the front/back dimension; higher front vowels can be articulatorily retracted to a greater degree without changing categorical front/back perception. Lower vowels, along with rounded vowels, have less perceptual stability with respect to articulatory perturbations in the front/back dimension, and are therefore less phonetically likely to be neutral, because changing the articulation along the front/back dimension changes the phonological category more easily. These authors specifically refer to vowels being less or more transparent, but that is equivalent to being a better or worse trigger and target of harmony, because they are discussing the possibility of vowels being skipped over (i.e. both not targeted by harmony and not triggering it themselves).

Benus & Gafos (2007) argue that these properties relate directly to articulatory facts. To explain the effect of rounding, they note that rounding lengthens the vocal tract, but the lingual constriction remains in the same horizontal location; this advances the area of perceptual stability, so that a front rounded vowel can be retracted less than a front unrounded one before changing perceptual categories. Similarly, for height, they note that the constriction required for lower vowels has less flexibility along the horizontal dimension, so that the potential for perceptually stable retraction of a lower front vowel is less than that of a higher vowel. Overall, they suggest that the effect of height and rounding on neutrality has a phonetic basis in the relationship between articulation, acoustics, and perception for vowels of different height and rounding.

Moreover, Beddor et al. (2001) suggest that the tendency of non-low front unrounded vowels, specifically [i] and [e], to be neutral in front/back harmony comes from smaller effects of vowel-to-vowel coarticulation on these sounds along the front/back dimension. These authors and others (e.g., Ohala 1994; Beddor & Yavuz 1995; Majors 1998) have suggested that vowel-to-vowel coarticulation can evolve into phonological harmony. As such, if coarticulation has a smaller effect on higher vowels, then these vowels are less likely to harmonize in a phonological system (i.e. are more likely to be neutral). Thus, the height effect could be attributed to phonetic properties of coarticulation, independent of phonological harmony.

In summary, cross-linguistic and phonetic evidence all suggest that lower vowels (and rounded vowels) are better potential targets of front/back harmony, or perhaps more that high vowels are worse targets, since most of the discussion focuses on high vowels.[23] A proper account of the behaviour of Hungarian front unrounded vowels must consider their quality as potential targets in order to capture this critical broader generalization. Specifically, the theory should be able to account for the Hungarian pattern, but not a case where lower vowels are neutral and higher vowels are not, which is unattested. Additionally, beyond front unrounded vowels, the [a:]~[e:] alternation follows this height generalization, in that low back vowels like [a:] are consistently harmonic despite being unpaired. A "reverse Hungarian" pattern in which [e:] consistently harmonizes to [a:], while [a:] generally behaves neutrally, should be impossible, because it contradicts the cross-linguistic, phonetically motivated generalization that low vowels are better undergoers of front/back harmony.

4 Implications for views of neutrality

Although it is typically ignored in analyses, the non-neutrality of Hungarian [e:] in suffixes (i.e. the [a:]~[e:] alternation described in Section 2) has significant implications for the understanding of neutrality in vowel harmony. This section describes the problems that this alternation poses to standard categorical analyses of Hungarian and the new view of neutrality that it suggests.

4.1 The problem

In standard analyses of Hungarian vowel harmony (e.g. Vago 1980a; Ringen & Vago 1998; etc.), [e:] does not undergo vowel harmony because there is no mid, back, unrounded, long vowel to

---

[23] Worth noting is that the target-based approach here does not say anything about the outcome of the harmony process, which is regulated by independent factors like faithfulness. Instead, it is simply a way of establishing whether harmony occurs at all between a given pair of vowels.

which it could harmonize. For example, Ringen & Vago (1998) propose a constraint *ɨʌ, which requires [+back, -low] vowels to also be round. Ranked above the harmony-enforcing constraint, this general markedness constraint prevents [e:] from undergoing harmony and allows it to be neutral. To my knowledge, there is no discussion in the theoretical literature about why [a:] does undergo harmony, despite also being unpaired, where ‘unpaired’ is defined as differing only on the harmonic feature (regardless of the presence of alternations), and therefore includes all of the unrounded vowels and low vowels in Hungarian.

The fact that [e:] can act as the harmonic counterpart for [a:] raises critical questions about this view of neutrality. Even if the [a:]~[e:] alternation is always underlying /a:/, precisely the same mechanisms that allow [a:] to undergo re-pairing to [e:] should force the same possibility for [e:], under the assumption of symmetry in harmonic alternations (though perhaps raising is an acceptable change while lowering is not). Indeed, if changing the feature [low] is allowed in order for /a:/ to harmonize to [e:], then it is unclear what blocks changes in [low] in so many cases of /e:/ in suffixes surfacing faithfully. The tableaux in Tables 3.4 and 3.5 demonstrate this problem, adapting the analysis of Maasai re-pairing from Pulleyblank et al. (1995) and using BACKNESS-HARMONY to stand in for any constraint that enforces harmony. Table 3.4 derives the correct result, with /a:/ re-pairing to [e:] when it follows a front vowel, but Table 3.5 incorrectly predicts that /e:/ should consistently surface as [a:] in a back vowel environment. The reverse ranking of BACKNESS-HARMONY and IDENT-IO[low] would predict that /e:/ can surface faithfully in Table 3.5, but gives an incorrect result in Table 3.4. Either way, the result is incorrect according to a categorical view, and it is also impossible given these assumptions to derive a gradient view in which /e:/ is variably neutral but /a:/ is not. (See Section 4.2 on possible alternatives with asymmetrical faithfulness to the [low] feature.)

/y…a:/	*ɤ:	*æ:	BACKNESS-HARMONY	IDENT-IO[low]
y…a:			*!	
☞ y…e:				*
y…æ:		*!		

Table 3.4: /a:/ undergoes re-pairing to [e:]

/u…e:/	*ɤ:	*æ:	BACKNESS-HARMONY	IDENT-IO[low]
☹ u…e:			*!	
☞ u…a:				*
u… ɤ:	*!			

Table 3.5: Incorrect prediction that /e:/ undergoes re-pairing to [a:]

Thus, the disparate behaviour of the two unpaired vowels [a:] and [e:] in Hungarian means that we cannot assume that general markedness, alone or in combination with faithfulness to other features, determines whether a vowel is a harmony target. Instead, neutrality must be derived in some other, more vowel-specific way.

4.2 Problems for possible solutions within traditional frameworks

Before considering the target-based view adopted here, I outline two potential solutions that are formulated within more traditional views of harmony, as well as the reasons for rejecting them. First, I consider separating IDENT-IO[-low] from IDENT-IO[+low], and then the possibility that harmony is triggered only by [-back]. Note that neither of these solutions has actually been proposed in the literature to deal with the Hungarian [e:]~[a:] alternation, unsurprisingly given the lack of discussion about this pattern within these kinds of analyses.

4.2.1 Ident-IO[-low] vs. Ident-IO[+low]

The problem illustrated in Tables 3.4 and 3.5 is that BACKNESS-HARMONY needs to outrank IDENT-IO[low] for /a:/, but the reverse ranking is required for /e:/. A possible way to deal with such an asymmetry is to incorporate it into the faithfulness constraint, with separate constraints IDENT-IO[-low] (violated for a change from [-low] to [+low] between an input segment and its output correspondent) and IDENT-IO[+low] (violated for a change fro [+low] to [-low]); this type of separation of faithfulness constraints by feature value has been used by McCarthy and Prince (1995; 1999). Respectively, these constraints prevent changing a [-low] feature to [+low], and a [+low] feature to [-low]. With a ranking of IDENT-IO[-low] >> BACKNESS-HARMONY >> IDENT-IO[+low], changing [+low] to [-low] is permitted to satisfy harmony, but the reverse feature change is not. Tables 3.6 and 3.7 demonstrate that this approach would correctly account for the Hungarian pattern, if we consider all such alternating suffixes to be underlyingly /a:/.

/y…a:/	*ɤ:	*æ:	IDENT-IO[-low]	BACKNESS-HARMONY	IDENT-IO[+low]
y…a:				*!	
☞ y…e:					*
y…æ:		*!			

Table 3.6: /a:/ undergoes re-pairing to [e:]

/u…e:/	*ɤ:	*æ:	IDENT-IO[-low]	BACKNESS-HARMONY	IDENT-IO[+low]
☞ u…e:				*	
u…a:			*!		
u… ɤ:	*!				

Table 3.7: /e:/ does not re-pair to [a:]

However, this option is highly stipulative, in that it just as easily predicts the reverse of the Hungarian pattern, with neutral [a:] and harmonic [e:], simply by switching the positions within the ranking of IDENT-IO[-low] and IDENT-IO[+low]. As discussed in Section 3, this result is undesirable, because it misses a crucial cross-linguistic, phonetically motivated generalization: low vowels are better undergoers of front/back harmony than non-low vowels, and so “reverse Hungarian” is an improbable language. Moreover, it would be typologically odd to claim that the ranking IDENT-IO[-low] >> IDENT-IO[+low] is universal, in contexts beyond front/back harmony, as it would predict that it would always be more important to preserve [-low] than [+low]. Thus, the height generalization cannot be captured in an analysis where the asymmetry falls to the separation of faithfulness constraints. In this type of analysis, it is also impossible to connect this alternation to the idea that [e:] is intermediately neutral.

4.2.2 Triggering by only [-back]

A second possibility framed within more traditional work is that only [-back] vowels trigger harmony, and (non-low) front unrounded vowels are unspecified for [back]. In that case, only back vowels can undergo harmony. Assuming again for the purposes of this argument that alternating [a:]~[e:] suffixes are all underlyingly /a:/, this would account for why /e:/ is exempt.

Such an approach fails to account for why front rounded vowels are not neutral in suffixes. Given Richness of the Base and the wide inventory of suffixes in Hungarian, we expect the possibility of underlying front rounded vowels in Hungarian suffixes. Under an account where only [-back] triggers harmony, any suffixes with underlying front rounded vowels would have no reason to harmonize, because a root containing [+back] vowels cannot trigger harmony. For example, for a UR like /ha:z+tø:l/, /a:/ would be unable to trigger harmony; since Hungarian harmony is strictly root-to-suffix, this UR should surface as *[ha:ztø:l]. In other words, we would predict a divide in which all front vowels can occur in native non-alternating suffixes, while back vowels cannot. Instead, in Hungarian, only /i, i:, e:/ as underlying suffix vowels are exempt from being targets of harmony; these vowels regularly behave as non-alternating in suffixes, while rounded front vowels are consistently alternating. This is unexplained if only [-back] can trigger harmony. Thus, this solution is undesirable as a way to explain the asymmetry between [a:] and [e:].

Moreover, if only [-back] can trigger harmony and front unrounded vowels are unspecified, there is no way to derive the variability described in Section 2.2, particularly that which occurs with the vowel [ɛ]. If [ɛ] is unspecified for [back], then it should be unable to trigger harmony ever, which leaves the possibility of front suffixes in (2i-j) unexplained. Similarly, if [ɛ] is specified as [-back], then it should always trigger front harmony, which leaves the possibility of back suffixes in (2h-i) unexplained. This argument also extends to the behaviour of disharmonic roots (i.e. those with a combination of back vowels and front rounded vowels), which consistently take suffix values based on the last stem vowel. If only [-back] can trigger harmony, then back vowels cannot trigger backness in the suffix, and so it is incorrectly predicted that front-back disharmonic stems should take front suffixes. We might try to claim that stem-internal disharmony begins a new harmony domain, and that back suffixes are a default value; however, doing so would leave the last-resort triggering in (3) unexplained. Indeed, if front unrounded vowels are unspecified for [back], suffixes for neutral-only stems should also contain a default value, yet the value in this case is front. Thus, it is inconsistent with the Hungarian data to claim that only [-back] can trigger harmony.

4.3 New view of neutrality

Due to the aforementioned problems with solutions in which neutrality is determined by lack of a harmonic counterpart, I argue that the Hungarian patterns require a fundamental change in how we view targets of vowel harmony: it suggests the need to consider the nature of each vowel as a possible harmony target. This view works not only for the [e:]~[a:] asymmetry, but more broadly for the height effect in Hungarian. I propose that the drive to undergo vowel harmony differs by vowel, in a way motivated by cross-linguistic and phonetic properties, and that cases of neutrality do not occur simply because harmony is impossible, but because the vowel-specific drive to undergo harmony is too weak to force unfaithfulness. This notion is gradient, and so will be able to avoid the issues of the categorical views.

In the case of front/back harmony, phonetic and cross-linguistic facts outlined in Section 3 show that vowels that are higher and unrounded are dispreferred targets. As such, the drive for lower vowels, such as [a:], to undergo harmony should be stronger than the drive for the non-low, unrounded vowel [e:]. The result, in Hungarian, is that [a:] is consistently a harmonic vowel, while [e:] can be neutral. The additional fact that [a:] is unpaired in the inventory, yet required to be a target, is what forces it to re-pair to [e:]. Moreover, the basic height effect in Hungarian is in direct accordance with the phonetic facts: the higher the vowel, the worse it is as a target of harmony, and therefore the more likely it will be to be neutral.

5 New theoretical approach: Target quality

To formally implement the new view of neutrality described in Section 4.3, I adopt the Harmonic Grammar (HG) framework (Legendre et al. 1990a; b; Smolensky & Legendre 2006; Pater 2009) and propose to extend Kimper's (2011) trigger strength analysis to targets. I also posit a target-oriented harmony constraint, which is a modification to Pulleyblank's (2004) no-disagreement harmony-driving constraints. This section provides the theoretical background, assumptions and motivations of my analysis; the application to Hungarian is left for Section 6.

5.1 Harmonic Grammar (HG)

Harmonic Grammar (HG; Legendre et al. 1990a; b; Smolensky & Legendre 2006; Pater 2009) is a model of phonology that is similar to standard Optimality Theory (OT; Prince & Smolensky 1993) except that constraints are weighted (Potts et al. 2010). In the specific implementation that I adopt here, I follow Potts et al. (2010) in assuming that constraints are negatively formulated, as in standard OT, and weights are positive integers. There are often many options that will work for the weight of each constraint; the specific choice is arbitrary (e.g. Bowman 2013).

Within HG, it has been proposed that violations of a constraint can be multiplied by a *scaling factor*, which allows properties of a specific violation to influence the degree to which the violation is weighted (e.g. Kimper 2011; see also Coetzee and Kawahara 2013 for scaling factors of a different type). A scaling factor x is a positive constant by which the constraint weight is multiplied; this has the effect of increasing the penalty incurred for a violation if $x > 1$ and decreasing it if $0 < x < 1$. As an example, in this dissertation, a scaling factor based on target quality will be applied to the harmony constraint. For example, the disharmonic combinations [u…y] and [u…i] both violate a progressive front/back harmony constraint once. However, if the vowels [y] and [i] have target scaling factors such as 4 and 2 respectively, then [u…y] will incur twice the penalty for this disharmony than will [u…i]. Further details are reserved for Sections 5.3 and 5.4.

For a given violation, the total penalty assigned by a constraint to a candidate (for a single locus of violation) is –W*S, where W is the weight of the constraint and S is any relevant scaling factors; such penalties are summed across all violations that the candidate incurs for all constraints (adapted from Potts et al. 2010). For example, given a harmony constraint with weight 3 and the target scaling factors defined in the previous paragraph, the harmony violation incurred by [u…y] would receive a penalty of −3*4 = −12, while [u…i] would have one of −3*2 = −6. The total of all penalties for all constraints is called the harmony score, and the candidate with the highest harmony score wins. Since the expression is always negative, the winning candidate will be the one with a harmony score closest to zero (Potts et al. 2010).

5.2 The nature of the harmony-driving constraint

An important theoretical implication of the Hungarian pattern is that the target-sensitive harmony constraint must be negatively formulated, in the sense that disharmony is penalized. While negative constraints are quite common, they are not universal; Kimper (2011) uses a positively formulated harmony constraint, in which harmony is rewarded, and Archangeli & Pulleyblank (2002) propose positive target conditions. However, in Hungarian, the good potential target [a:] becomes a poor potential target[24], [e:], while the reverse does not always occur. If harmonizing

---

[24] I use the term ‘potential target’ to refer to a vowel that is in a position to have undergone harmony but has not. It can therefore refer to both surface and underlying forms.

to good targets were rewarded, then we would predict the opposite pattern; the poor target [e:] would become the good target [a:] in order to benefit from the reward; /a:/ surfacing unfaithfully as [e:] would be rewarded less and therefore be less likely to occur. Instead, the direction of the Hungarian asymmetry necessitates a penalty for good targets that have not undergone harmony. Non-harmonized [a:] will be heavily penalized, and so [a:] is required to harmonize, whereas non-harmonized [e:] will be penalized less and therefore permitted.

As long as there is a mechanism to distinguish trigger and target, the specific constraint used to enforce harmony is independent of both the principle of target scaling and the perspective that neutrality is caused by the drive to undergo harmony being too weak. I adopt a variation on Pulleyblank's (2004) infinite distance no-disagreement constraints, defined in (7) and subject to trigger and target scaling as in Sections 5.3 and 5.4.

(7) $*[-\alpha\text{Back}]\infty\underline{[\alpha\text{Back}]}$: Assign a violation for every vowel that is [αBack] and is preceded at any distance by a vowel that is [-αBack].

The underline indicates that this constraint is oriented towards the second vowel; it is the focus/target, and therefore the locus of violation, while the first vowel is the context/environment. Violations are counted for each [αBack] preceded by a [-αBack] vowel, and the infinity symbol indicates that such a configuration is a violation regardless of the distance or intervening segments. Only one violation is possible per potential [αBack] target. For example, a sequence such as [u… y_1…y_2] violates (7) twice (i.e. there are two targets/foci), once for [y_1] and once for [y_2], both of which are preceded at some distance by a back vowel [u]. The fact that [y_1] intervenes between [u] and [y_2] is irrelevant to the fact that a violation is assigned.

On the other hand, the sequence [y_1…y_2…u] violates (7) only once, for the [u] that is preceded by a front vowel; this is because the constraint counts violations based only on vowels that are preceded by another vowel that disagrees in [back].

Since Hungarian harmony is progressive, this constraint counts violations based on the number of vowels that could be targets of harmony but are not. For the purposes of trigger scaling, it is counted from the closest possible potential [-αBack] trigger. Thus, for each pair of vowels violating the constraint, the target scaling factor is based on the vowel for which the violation is counted (the focus), while the trigger is the closest preceding vowel of the opposite [back] value (the context). For example, in [y_1…y_2…u], the one violation has trigger [y_2] and target [u]. The two violations in [u… y_1…y_2] both have trigger [u], and the targets are [y_1] and [y_2] respectively.

The constraint in (7) is equivalent to the two separate constraints *[-Back]∞[+Back] and *[+Back]∞[-Back]. This approach therefore differs from that of Pulleyblank (2004) in that the constraint is consistently oriented not to a specific value of the feature [back], but rather to the second vowel in a disharmonic sequence. This choice offers a conceptual advantage for a case of harmony like Hungarian, which is arguably progressive and triggered by either feature value, because it offers an easy way of identifying the trigger and target for each violation. For later chapters that deal with ATR-dominant harmony, I adopt Pulleyblank's mechanism of identifying a single feature value (in these cases, RTR) as the focus/target. In comparison, other types of harmony constraints, such as AGREE[back] (e.g. Baković 2000), do not have a natural way of referring to the potential triggers and targets of harmony.

In addition to providing an inherent definition of trigger and target, this type of constraint has a more natural way of dealing with transparency than other options. This is somewhat reminiscent of headed spans (McCarthy 2004) and Optimal Domains Theory (ODT; Cole & Kisseberth 1994), which identify the head/trigger of the harmonic domain; however, in the current approach there is no inherent reason for the vowel scaled as trigger to also be the actual trigger; this comes from the interaction of other constraints. Spreading-type constraints like ALIGN[back] or SPREAD[back] generally require significant representational complexity, such as line crossing, gapped configurations, or output underspecification, to deal with transparency (e.g. Jurgec 2011; Ní Chiosáin & Padgett 2001; Ringen & Vago 1998; Walker 1998; 2012). Further, constraints like AGREE[back] do not have a natural way of referring to non-adjacent vowels, and they naturally prefer opacity to transparency in order to reduce the number of disharmonic transitions (Baković & Wilson 2000). In contrast, a target-oriented constraint captures the generalization that all vowels should be targets of harmony when a potential trigger is present; no additional representational complexity is required, and transparency is derived because whether other potential targets intervene is irrelevant to whether the constraint is violated.[25] See Pulleyblank (2004) for additional argumentation for this style of constraint. Some of the constraints used by Hayes & Londe (2006) in dealing with Hungarian are of a similar type, though formulated differently.

It is worth noting that the main cases of harmony that Pulleyblank (2004) discusses are dominant-recessive systems, and the constraints are oriented towards the recessive (target) feature value, such as ATR in Yoruba. Thus, although Pulleyblank (2004) does not discuss targets, the constraint here is consistent with the insight of his constraints, despite the distinct implementation. Nonetheless, in the standard OT framework adopted by Pulleyblank (2004),

---

[25] Note, however, that it is not necessarily irrelevant to the weight of that violation: the violation can be scaled by distance, as Kimper (2011) and Bowman (2013) do, so that harmony violations at a greater distance will be weighted less than those at a closer distance. This type of scaling can capture the count effect in Hungarian transparency.

without weights or scaling, the constraint in (7) could not predict the Hungarian transparency facts; both Back-Front-Back and Back-Front-Front sequences violate *[-αBack]∞[αBack] twice, and so this constraint without scaling could not distinguish these candidates. The next two subsections deal with how this constraint will be scaled in the current framework to allow for transparency.

5.3 Kimper's (2011) trigger strength

In order to explain asymmetries in triggering behaviour in vowel harmony, Kimper (2011) introduces a trigger strength scaling factor for his harmony constraint. As described in Section 5.1, this scaling is a constant for a given vowel, and the weight of a violation of the harmony constraint is multiplied by this scaling value to determine the penalty (or, in Kimper's framework, reward) of a particular trigger configuration. Thus, it has the effect of increasing the likelihood of harmony when the potential trigger is perceptually impoverished for the harmony feature, which is a characteristic of good harmony triggers (Kimper 2011). Trigger strength is therefore vowel-specific, but cross-linguistic and phonetically motivated, exactly parallel to the concept of target-specific harmony advocated for here.

Kimper (2011) employs a more complex formal model than the one necessary here, but the scaling factor that he develops for trigger strength forms the basis for the current approach to vowel-specific targeting. A simplified version of Kimper's (2011) trigger scaling factor, quoted from Bowman (2013: 4), is given in (8).

(8) "Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the reward earned by a constant x (such that x

> 1) for each degree *i* to which α is perceptually impoverished with respect to ±F (simplified from Kimper 2011)"

This constraint refers to rewards, since Kimper (2011) and Bowman (2013) use a harmony constraint that is formulated positively, meaning that harmonic configurations are rewarded. As discussed in Section 5.2, such a constraint is impossible here. However, the trigger strength scaling factor can easily be reformulated for a negative constraint, as in (9).[26]

(9) Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that x ≥ 1) for each degree *i* to which α is perceptually impoverished with respect to ±F.

This scaling factor has the effect of increasing the penalty incurred for a disharmonic trigger/target pair when the trigger is perceptually impoverished. The result is that not harmonizing is worse with a good trigger than with a bad trigger. For example, in a front/back harmony system, if [y] is more perceptually impoverished for [back] than [i] (i.e. [y] is a better trigger), then the trigger strength scaling for [y] should be greater than or equal to that of [i]. This scaling factor is then applied to the weight of the harmony constraint in the way described in Section 5.1. As such, a disharmonic sequence in which [y] is the potential trigger will receive a greater penalty than one in which [i] is the potential trigger. In order to capture the fact that Hungarian neutral vowels are transparent instead of opaque, I adopt the trigger scaling factor in (9) in addition to the target one introduced in Section 5.4. (See Kimper (2011) and Bowman

---

[26] Of course, Kimper (2011) has reasons for assuming a positive harmony constraint, and using a negative one instead will have implications to the theory. However, this issue is beyond the scope of this dissertation.

(2013) on the application of trigger strength to Hungarian.)

Kimper (2011) also proposes a distance scaling factor, which is particularly relevant for the count effect in Hungarian; he suggests that the reward of harmony decreases as the distance between the trigger and target increases. Since I formulate the harmony constraint negatively here, in this case, the penalty will decrease. A definition of the distance scaling factor, based on the simplified version in Bowman (2013), is given in (10).

(10) Scaling factor: distance

For a trigger α and a target β, multiply the penalty earned by a constant x (such that $0<x<1$) for each unit (e.g. syllable) of distance *d* intervening between α and β.

In Hungarian, this factor will have the effect of penalizing disharmony less at a greater distance, which will create the count effect; with multiple neutral vowels, the back vowel in the root is further from the suffix target, and so a front suffix will be penalized less than it would with only a single neutral vowel.

5.4 Target quality

As argued in Sections 2 through 4, whether a vowel undergoes harmony or not is sensitive to its quality as a potential target. This observation is similar to Kimper's (2011) discussion about triggers, and so I formalize it by adapting Kimper's (2011) trigger strength scaling factor in (8-9) for use with targets. The statement of this scaling factor is given in (11).

(11) Scaling factor: target quality

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that $x \geq 1$) for each degree *i* to which β is a quality target (defined by phonetic factors discussed in Section 3) with respect to ±F.

Parallel to the one for trigger strength, this scaling factor has the effect of increasing the penalty for disharmonic pairs when the potential target is a good harmony target. In this definition, I remain agnostic on the precise factors that determine target quality; an in-depth analysis of the underlying mechanisms behind target quality is left for future research. However, as discussed in Section 3, it is clear from cross-linguistic patterns that low and round vowels are better targets of front/back harmony than non-low, unrounded vowels, and there is phonetic motivation for this asymmetry, based on the height-based differences in distinguishing vowels along the front/back dimension, in connection with the purported articulatory and perceptual motivations for vowels undergoing harmony. As such, for front/back harmony, the target quality scaling factor of vowels that are low and/or round should be greater than or equal to that of non-low, unrounded vowels. This scaling is implemented within the model as described in Section 5.1.

5.5 Markedness conflation

Both trigger and target scaling factors are hypothesized to be universal, but only in the sense that they are motivated by cross-linguistic phonetic facts and are options across languages. A language may choose to treat all vowels equally in terms of triggering harmony and/or being targets; in such cases, the harmony system would be symmetrical. However, if a language chooses to scale for trigger strength and/or target quality, there are limited options for the relative scaling factors across vowels, and these options are built into the definition through the phonetic motivations. For example, since target quality in front/back harmony is motivated by articulatory and perceptual properties of height and rounding with respect to backness, languages can pattern like Hungarian, with only vowels that are round and/or low undergoing harmony, but reverse Hungarian is predicted to be impossible. Specifically, for front/back harmony systems the target scaling factor for non-low, unrounded vowels must be less than or equal to that of vowels that are low and/or rounded. The option for equality allows for symmetry, but if there is asymmetry, it must be in the direction of less participation of non-low, unrounded vowels.

While the formalization differs, this approach parallels previous work on similar hierarchies, such as de Lacy's (2004) concept of markedness conflation. The target concept here is similar in that it reflects universal properties of which vowels are better targets, and the "ranking" of possible targets can be conflated. For example, considering only height, since lower vowels are better targets of front/back harmony, we predict the following possibilities: no vowels participate, only low vowels participate (neutrality of mid and high vowels), low and mid vowels participate (neutrality of high vowels), or all vowels participate (no neutrality). Languages where front/back harmony does not fit within these categories, such as with neutrality of low vowels to front/back harmony but participation of high and mid vowels, are not expected. These predictions are consistent with the cross-linguistic height generalizations in front/back harmony.

Formally, this universal hierarchy is implemented here through the phonetic motivations behind the target scaling factor in the definition, and the possibility of conflation comes from the option for multiple categories to have the same scaling factor. Thus, the target scaling values by height for front/back harmony (given consistent rounding) must be such that low ≥ mid ≥ high; the different patterns are derived by the possibility of equality. Recall that the target scaling factor itself is defined in terms of phonetic correlates of target quality, so that this hierarchy is not simply stipulated.

Similar to de Lacy's (2004) markedness conflation constraints, this approach does not require a stipulated universal ranking, because the scaling factor is defined to be motivated by universal, cross-linguistic properties. It remains for future research to determine whether the scaling factors described here and the stringency hierarchy of constraints discussed by de Lacy (2004) are notational variants or whether they in fact make distinct predictions.

5.6 Morpheme-specificity

As discussed in Section 2.4, whether a particular vowel is a trigger or target in Hungarian depends not only on vowel quality, but also whether it is in a stem or suffix, as well as the specific morpheme. The former property is simple to account for in terms of root faithfulness (Beckman 1998) in combination with constraint weights, since in general in Hungarian, there are stricter harmony requirements on suffixes than on roots (e.g. [ɛ] is required to harmonize in a suffix, but not in a root).

Morpheme-specific behaviour is more complicated, in particular given its interaction with vowel quality; variation is lexically determined only for [eː] and [ɛ] (Rebrus et al. 2012). In particular, under the aforementioned current assumption in the Hungarian literature that [aː]~[eː] suffixes are not all underlyingly /aː/, we need to be able to capture the fact that [eː], unlike all other vowels, can be a target in some but not all suffixes. Further, both [eː] and [ɛ], unlike other vowels, can be triggers in some but not all roots in which they follow back vowels. For example, (2d) showed a back-[eː] root that takes only back suffixes, [kaːveː-nɔk], while (2g) showed a similar root that can take either back or front suffixes, [ɔsteːk-nɔk/nɛk]. Similarly, (2i) showed a back-[ɛ] root that can take either back or front suffixes, [maːgnɛʃ-nɔk/nɛk], while (2j) showed one that takes only front suffixes, [koːdɛks-nɛk]. In other words, whether a given vowel is a trigger or target depends not only on its quality, but also on the morpheme it occurs in. The choice of a given morpheme is arbitrary, except that it is limited to the "zones of variation" (Hayes et al. 2009: 835) that are possible for morphemes of its type.

It would be conceivable to think of Hungarian harmony as scaled entirely by morpheme rather than by vowel quality, with each root morpheme specified to take certain types of suffixes. A full analysis does require a morpheme-specific component, in order to account for the fact that a given morpheme patterns consistently in the system, in a way that can be distinct from other morphemes of the same phonological shape (e.g. the different back-[ɛ] roots in (2)). However, allocating all of the phonological work to morpheme-specific factors would be highly stipulative, given that the set of options for the behaviour of a given morpheme is predictable based on vowel quality in Hungarian. Incorporating factors of trigger and target vowel quality into the analysis reduces this stipulation, by limiting the set of possibilities for morpheme-specific behaviour to the existing phonologically-defined trends, in a way that is typologically and phonetically motivated. Thus, an analysis that incorporates both morpheme- and vowel-specific factors is preferable to an account that deals with only one of these.

The morpheme-specific behaviour seen in Hungarian is of a very predictable type: a vowel may exceptionally have the normal behaviour of a vowel that is one step better as a trigger or target, in the scale (from low to high) of [i(ː)] – [eː] – [ɛ] – harmonic vowels. Specifically, [i(ː)] and the harmonic vowels are not variable, the former because the normal behaviour of [eː], like that of [i(ː)], is neither triggering nor being a target of harmony, and the latter because there is no category above. In contrast, [eː] can, in some morphemes, show the normal triggering and targeting behaviour of [ɛ], in that it variably triggers harmony (vacillation) in some roots (e.g. (2g)) and can be a target of harmony in some suffixes; both of these are the normal behaviour of [ɛ]. Similarly, in some morphemes, [ɛ] can show the normal behaviour of the fully harmonic vowels, in consistently triggering harmony (e.g. (2j)). Note that [ɛ] already shows the typical targeting behaviour of the other vowels, since it does not occur in invariant suffixes.

Due to these generalizations about the possible morpheme-specific behaviour, I will encode morpheme-specific behaviour into the trigger and target scaling factors, with vowels exceptionally able to take the trigger and target scaling of a vowel in the category above. It remains an open question why there appears to be this one-step limit in morpheme-specific behaviour (e.g. why [i] cannot behave like [ɛ]), but this issue is beyond the scope of the present chapter. I suggest the possibility that a one-step change can be consistent with the phonetic properties of a given vowel, which motivate its normal behaviour in relation to how the language sets its scaling factors, but a larger change cannot be.

6 Account of Hungarian

Using the principles outlined in Section 5, I now pursue a formal account of the Hungarian data described in Section 2. I show that with the correct constraint weights and scaling factors, this approach can capture the gradient, variable patterning of Hungarian neutral vowels, in addition to the basic facts that are accounted for in previous analyses.

6.1 Constraints, weights, and scaling

As discussed in Section 5.2, I adopt a target-oriented harmony constraint that reflects that Hungarian harmony is progressive and triggered by either feature value. This constraint is repeated in (12). One violation of this constraint is counted for each [αBack] vowel preceded by a [-αBack] vowel at any distance.[27]

(12) *[-αBack]∞[αBack]: Assign a violation for every vowel that is [αBack] and is preceded at any distance by a vowel that is [-αBack].

In addition to this harmony constraint, we require general markedness constraints to prevent vowels that are unattested in Hungarian from surfacing, as well as a faithfulness constraint against changes to [low]. These constraints are defined in (13).

(13)
*[ɤː]: Long, [-high, -low, -round] vowels must not be [+back].
*[æː]: Long, [+low, -round] vowels must not be [-back].[28]
IDENT-IO[low]: Assign a violation for every change in the feature [low] between a vowel in the input and its correspondent in the output.

While the harmony constraint is in conflict with faithfulness to the feature [back], the general constraint IDENT-IO[back] is not active within the Hungarian system, and so I omit it from the analysis.[29] Instead, I will assume that there is an active positional faithfulness constraint (e.g. Beckman 1998), specifically a highly weighted constraint to preserve the [back] specification of

---

[27] I will assume that this constraint applies only within the domain of the phonological word. In particular, I assume that the two parts of a compound are separate phonological words, and therefore this constraint does not affect compounds.

[28] Note that this constraint operates on phonological features. While [aː] is pronounced by many speakers as front, as mentioned previously, it behaves phonologically as [+back], and is therefore not subject to this constraint.

[29] Formally, this could be accomplished by giving it a weight of zero, or just a sufficiently low weight (see Potts et al. 2010 on zero weights in HG).

the word-initial vowel. This constraint will not be illustrated in the tableaux for space reasons, but will ensure that the initial vowel is the trigger of harmony.[30] Since Hungarian also has rounding harmony, I assume the same about IDENT-IO[round] as about IDENT-IO[back].

Many previous analyses of Hungarian harmony (e.g. Ringen & Vago 1998) ignore root-internal harmony, assuming a highly ranked root faithfulness constraint. In reality, a high ranking of such a constraint is only necessary for loanwords (see footnote 21 on loanword treatment). As discussed in Section 2, disharmonic roots consisting of non-neutral vowels are recent loans; within native Hungarian roots, non-neutral vowels must agree in backness, while front unrounded vowels may both precede and follow either front or back vowels. This pattern is in contrast to the behaviour of neutral-back sequences across morpheme boundaries, because neutral vowels trigger last-resort harmony (see (3)). I suggest that these facts can be accounted for using root-specific faithfulness to [back], defined as in (14).

(14) IDENT-IO$[\text{back}]_{\text{ROOT}}$: For vowels within the domain of the root, assign a violation for every change in the feature [back] between a vowel in the input and its correspondent in the output.

Weighted correctly with respect to harmony, this constraint will allow non-neutral vowels to harmonize within a root, but require root-internal disharmony involving neutral vowels to surface faithfully. Given that the constraint is root-specific, it will not affect last-resort triggering

---

[30] Note that something distinct needs to be proposed for loanwords, in which there can be root-internal disharmony and the final root vowel triggers harmony to the suffix. It is common for loanwords to have less restrictive phonotactic patterns compared to native words; this occurs in a variety of languages, for many phonotactic patterns (e.g. Itô & Mester 1995; Inkelas et al. 1997; Jurgec 2010). Such cases could be accounted for with co-phonologies (e.g. Inkelas & Zoll 2007), or by scaling faithfulness constraints by lexical stratum (foreign versus native), as done by Hsu & Jesney (2017). It is worth noting that Hsu & Jesney's (2017) approach is insufficient to explain morpheme-specific behaviour in Hungarian, even though many of the exceptional morphemes are loans. This is because such morphemes may exceptionally trigger harmony, which should be prevented if faithfulness is higher ranked in the loan stratum. Instead, an analysis of Hungarian that considers loans would require both the foreign stratum and the morpheme-specific category changing adopted in this chapter.

in which suffix vowels harmonize to neutral root vowels.

The constraint weights and scaling factors are set as in (15). There are many choices that would work, and the specific numbers are arbitrary, as long as certain conditions are satisfied; throughout the analysis, I will argue for these specific conditions on weighting that are required for Hungarian. Note that the scaling factors in (15b-d) are specific to the particular constraint *[-αBack]∞<u>[αBack]</u>, and are not applied to the grammar as a whole; this follows Kimper (2011) and Bowman (2013).

(15a) Constraint weights

*[ɤː]: weight 400

*[æː]: weight 400

IDENT-IO[back]$_{\text{ROOT}}$: weight 200

IDENT-IO[low]: weight 20

*[-αBack]∞<u>[αBack]</u>: weight 1

(15b) Trigger strength scaling

2 for [i] and [iː]

2 for [eː][31]

4 for [ɛ]

16 for all other vowels

---

[31] There is no difference between [i] and [eː] for trigger strength, because in the general case, both vowels are non-triggers. However, [i] and [eː] must nonetheless be separate categories, such that it is possible for [eː] to exceptionally be a trigger but not [i].

(15c) Target quality scaling

1 for [i] and [i:]

1 for [e:]

8 for [ɛ]

16 for all other vowels

(15d) Distance scaling

$1/2^n$, where n is the number of syllables intervening between trigger and target

As discussed in Section 5.6, note that both trigger and target scaling can vary morpheme-specifically, in that a morpheme may select for its vowel to behave like the vowel one category above it. This will be discussed further in Section 6.4.

6.2 Basic harmony facts

The present analysis retains the ability to capture the basic facts of Hungarian harmony that are the focus of other analyses. Tables 3.8-3.13 illustrate. Following Vago (1980b), among others, I assume that the dative suffix underlyingly has a front vowel; this assumption is not critical to the analysis, due to the low (zero) hypothesized weight of the general IDENT-IO[back].

Table 3.8 shows that suffix [ɛ] is required to harmonize following non-neutral roots; the same will apply for front rounded and back vowels. Disharmony violates *[-αBack]∞[αBack], while the harmonic form does not violate any of the constraints due to the absence of IDENT-IO[back] in the analysis (i.e. its very low hypothesized weight). This tableau does not demonstrate any weighting of the included constraints.

/va:roʃ-nɛk/	*[ɤ:] 400	*[æ:] 400	Ident-IO[back]$_{\text{Root}}$ 200	Ident-IO[low] 20	*[-αBack]∞[αBack] 1	H
a. va:roʃnɛk					−1*16(o)*8(ɛ)	−128
☞b. va:roʃnɔk						0

Table 3.8: Harmony with non-neutral vowels

Table 3.9 shows that front rounded and back vowels are required to harmonize within a root; disharmony violates *[-αBack]∞[αBack], while the harmonic form in this case violates Ident-IO[back]$_{\text{Root}}$. This tableau demonstrates the need for the weighting condition in (16). (See footnote 30 on the distinct analysis required for loanwords, in which disharmony is allowed.)

/va:røʃ/	*[ɤ:] 400	*[æ:] 400	Ident-IO[back]$_{\text{Root}}$ 100	Ident-IO[low] 20	*[-αBack]∞[αBack] 1	H
a. va:røʃ					−1*16(a:)*16(ø)	−256
☞b. va:roʃ			−1			−100

Table 3.9: Root-internal harmony with non-neutral vowels

(16)

weight(Ident-IO[back]$_{\text{Root}}$)

< weight(*[-αBack]∞[αBack]) * scaling(good trigger) * scaling(good target)

Table 3.10 demonstrates that front unrounded vowels are able to be transparent. In candidate (a),

the neutral vowel [e:] is opaque, with the suffix surfacing as front. This candidate has two violations of *[-αBack]∞[αBack]: one for the potential target [e:], with trigger [a:], and the other for the potential target [ɛ], again with trigger [a:]. The first violation is scaled by 16(a:)*1(e:) = 16. The second is scaled by 16(a:)*8(ɛ)*1/2(distance) = 64, since there is one syllable intervening between them. The total score for candidate (a) is therefore −80. Candidate (b), with transparency, similarly has two violations of *[-αBack]∞[αBack]: one is again for the potential target [e:], with trigger [a:], but the other is for the potential target [ɔ] with trigger [e:]. Again, the first violation is scaled by 16, but this time, the second is scaled by 2(e:)*16(ɔ) = 32. There is no distance scaling here, since the vowels are in adjacent syllables. The total for candidate (b) is thus −48. A fully harmonic candidate, changing the /e:/ to [a:], violates both root faithfulness to [back] and faithfulness to low, and therefore receives a score of −220. Thus, candidate (b) wins, since −48 is the highest score; transparency (b) is preferred over opacity (a) because [e:] is a poor trigger, and over full harmony (c) due to root faithfulness. The required weighting conditions demonstrated in this tableau are given in (17); note that here is where it is necessary to maintain Kimper’s (2011) trigger strength scaling in order to derive transparency rather than opacity.

/ka:ve:+nɛk/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
a. ka:ve:nɛk					−1*16(a:)*1(e:) −1*16(a:)*8(ɛ)*1/2	−80
☞b. ka:ve:nɔk					−1*16(a:)*1(e:) −1*2(e:)*16(ɔ)	−48
c. ka:va:nɔk			−1	−1		−220

Table 3.10: Transparency of neutral vowels

(17)

a. scaling(trigger [e:]) * scaling(good target)

< scaling(good trigger)*scaling(target [ɛ]) * scaling(1 syllable distance)

b. weight(*[-αBack]∞[αBack])*scaling(good trigger) * scaling(target [e:])

+ weight(*[-αBack]∞[αBack])*scaling(trigger [e:]) * scaling(good target)

< weight(IDENT-IO[low]) + weight(IDENT-IO[back]$_{\text{ROOT}}$)

Table 3.11 illustrates that neutral vowels are last-resort triggers of suffix harmony, even in cases where harmony involves an additional faithfulness violation. Candidate (a) is faithful, but violates harmony, with scaling of 2 for trigger [i:] and 16 for target [a:]. The total is therefore −32. In contrast, candidate (b) harmonizes, violating only IDENT-IO[low], which gives a score of −20. The harmonic candidate is the winner. The same result will also hold for suffixes that do not violate IDENT-IO[low]. The weighting condition from this tableau is given in (18).

/vi:z+na:l/	*[ɤ:] 400	*[æ:] 400	IDENT- IO[back]$_{\text{ROOT}}$ 200	IDENT- IO[low] 20	*[-αBack]∞[αBack] 1	H
a. vi:zna:l					−1*2(i:)*16(a:)	−32
☞b. vi:zne:l				−1		−20

Table 3.11: Last-resort triggering by neutral vowels

(18)

weight(IDENT-IO[low])

< weight(*[-αBack]∞[αBack]) * scaling(trigger [i:]) * scaling(good target)

Importantly, the last-resort triggering in Table 3.11 happens only across a suffix boundary; root-internal neutral-back sequences surface faithfully due to IDENT-IO[back]$_{\text{ROOT}}$. Table 3.12 illustrates; the faithful candidate wins because harmony violates IDENT-IO[back]$_{\text{ROOT}}$. (19) gives the required weighting conditions. This result will also hold for any combination of back vowels with front unrounded vowels (i.e. any of these vowels, in any order), because of the high weight of IDENT-IO[back]$_{\text{ROOT}}$ relative to the trigger and target scaling factors of front unrounded vowels.[32] Specifically, they highest trigger or target scaling of a front unrounded vowel is 8; because 8*16=128<200, even root-internal disharmony involving [ɛ] can surface faithfully. This is in contrast to Table 9, where root faithfulness is not maintained with harmonic vowels; in that case, the harmony constraint was scaled such that it outweighed root faithfulness. This result is consistent with the generalization that non-neutral vowels in native Hungarian roots are from the

---

[32] Unlike Turkish, Hungarian does not appear to have bisyllabic suffixes in which the first vowel undergoes harmony but the second does not. Hungarian does have some ‘opaque/domain-external’ suffixes (Siptár & Törkenczy 2000: 65) that are disharmonic, but neither vowel undergoes harmony and the first is neutral.

same harmonic class, whereas front unrounded vowels can co-occur with either class.

/bikɔ/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
☞ a. bikɔ					−1*2(i)*16(ɔ)	−32
b. bikɛ			−1			−200

Table 3.12: Root-internal neutral-back sequences surface faithfully

(19)

weight(*[-αBack]∞[αBack]) * scaling(trigger [i]) * scaling(good target)

< weight(IDENT-IO[back]$_{\text{ROOT}}$)

6.3 [a:]~[e:] alternation

Given the assumptions and constraints outlined in Section 6.1, this analysis can also accurately capture the asymmetric re-pairing behaviour of Hungarian [a:] and [e:], in which [a:] is required to re-pair but [e:] is not. Table 3.13 shows that /a:/ re-pairs to surface as [e:] when it follows a front rounded vowel. Direct harmony in [back], as in candidate (c), violates the highly weighted constraint *[æ:]. The decision is then between re-pairing, in (b), and faithfulness, in (a). Re-pairing violates IDENT-IO[low], with a weight of −20. The faithful candidate is disharmonic, therefore violates *[-αBack]∞[αBack]. Trigger and target scaling are each 16, giving a total of 1*(−1*16*16)=−256 as the score for candidate (a). Since −20 > −256, candidate (b) wins, and /a:/ is required to re-pair to [e:]. To ensure that harmony occurs in this case, we require the weighting conditions stated in (20). Note that /a:/ will not re-pair to [i:] or [y:] because such candidates are harmonically bounded by the one in which /a:/ re-pairs to [e:]: all of these candidates violate IDENT-IO[low], but re-pairing to a high vowel also violates IDENT-IO[high]. I therefore omit candidates in which /a:/ re-pairs to a high vowel, as well as the constraint IDENT-IO[high], from the tableaux for presentation reasons. The generalization here is that /a:/ re-pairs to the front vowel that requires the fewest feature changes, and so changing [high] in addition to [low] will be worse than changing only [low]. Note that Hungarian harmony never changes length; we can assume that this in combination with the constraint against a long, low, front unrounded vowel is the reason why repairing in this case does not go to [ɛ] or [ɛ:].

/tøk+na:l/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
a. tøkna:l					−1*16(ø)*16(a:)	−256
☞b. tøkne:l				−1		−20
c. tøknæ:l		−1				−400

Table 3.13: Re-pairing of [a:] with [e:]

(20)

a. weight(IDENT-IO[low])

< weight(*[æ:])

b. weight(IDENT-IO[low])

< weight(*[-αBack]∞[αBack]) * scaling(good trigger) * scaling(good target)

In contrast, Table 3.14 demonstrates that /e:/ is not required to re-pair when it follows a back vowel. Similarly to Table 3.13, direct harmony in [back], as in candidate (c), violates the highly weighted constraint *[ɤ:], and the decision is then between re-pairing, in (b), and faithfulness, in (a). Again, re-pairing violates IDENT-IO[low], with a weight of −20, while the faithful candidate is disharmonic and therefore violates *[-αBack]∞[αBack]. However, since the potential target is [e:], target scaling is only 1. Thus, the total score for candidate (a) is 1*(−1*16*1)=−16. Since −16 > −20, candidate (a) wins in this case, meaning that /e:/ is not required to re-pair to [a:]. This tableau demonstrates the need for the weighting conditions given in (21). Critically, comparing the second condition in (21) to the second condition in (20) shows the need for scaling(target [e:]) < scaling(target [a:]), illustrating the need for the target quality concept argued for here. The alternative harmonic candidate [la:ɲo:] is ruled out by rounding faithfulness not shown here.

/la:ɲ+e:/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
☞a. la:ɲe:					−1*16(a:)*1(e:)	−16
b. la:ɲa:				−1		−20
c. la:ɲɤ:	−1					−400

Table 3.14: [e:] does not re-pair with [a:]

(21)

weight(*[-αBack]∞[αBack]) * scaling(good trigger) * scaling(target [e:])

< weight(*[ɤ:])

weight(*[-αBack]∞[αBack]) * scaling(good trigger) * scaling(target [e:])

< weight(IDENT-IO[low])

These examples show that this analysis correctly predicts the asymmetry between [a:] and [e:] in Hungarian, by appealing to the fact that [a:] is a better target of front/back harmony than [e:] is. Disharmony involving the potential target [a:] is therefore penalized to a greater extent than disharmony involving the potential target [e:]; the result is that harmony outweighs faithfulness for /a:/, but the reverse is true for /e:/. Notably, unlike an analysis that uses separate faithfulness constraints for [+low] and [-low], the current approach does not predict the possibility of reverse Hungarian. Indeed, due to the phonetic basis and markedness conflation inherent in the definition of target scaling, the target scaling factor for [e:] must be less than or equal to that of [a:]. As such, these vowels must either pattern like they do in Hungarian or behave symmetrically.

6.4 Variability

The behaviour in Table 3.14 is not the only possibility for how [e:] will behave in suffixes. As noted in Section 2, [e:] has intermediate neutrality along this dimension; the recent literature seems to assume that the underlying form of affixes that alternate between [a:] and [e:] is not necessarily /a:/, but rather that /e:/ is not fully neutral and could be required to harmonize in suffixes. This is possible under the current analysis due to the option for morpheme-specific target scaling factors. As discussed in Section 5.6, Hungarian morphemes may exceptionally select for their vowels to take the scaling trigger and/or target factors of a vowel that is one step better as a trigger/target. It is worth noting that we could not posit that such cases are abstractly the missing [ɛ:], given that [ɛ] behaves in the same way, sometimes taking on the scaling factors of a fully harmonic vowel. If we have such a suffix with [e:], it will be required to harmonize, as shown in Table 3.15, where the subscript "ɛ" indicates an [e:] that morpheme-specifically behaves like [ɛ], in the way justified in Section 5.6. Note that if [e:] in this case were not behaving like [ɛ], then it would have a scaling factor of 1, and the output would be -16 rather than -128.

/dob+ne:$_{\varepsilon}$l/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
a. dobne:l					−1*16(o)*8(e:$_{\varepsilon}$)	−128
☞b. dobna:l				−1		−20
c. dobnɤ:l		−1				−400

Table 3.15: Re-pairing of [a:] with [e:]

In other words, while rounded vowels and back vowels are required to harmonize, [e:] may or may not: it is of intermediate neutrality with respect to being a target in a suffix. In this approach, it does not matter which of the alternating [a:]~[e:] suffixes are underlyingly /a:/ and which are underlyingly /e:/. This possibility does not exist for [i(:)], because the next category up is [e:], which has the same target scaling as [i(:)]. It also does not exist for [ɛ] or other vowels, which already have target scaling factors sufficient to ensure that harmony in suffixes must always occur. Thus, [i(:)] is consistently invariant, [e:] is variable, and [ɛ] and other vowels are consistently targets, reflecting the suffix height effect illustrated in Table 3.3 in Section 2.

I turn now to variability in triggering, which occurs in distinct ways for both [e:] and [ɛ]: the former can be transparent or vacillating (in the sense that a single morpheme can take either back or front suffixes), while the latter is typically vacillating or opaque. An example of transparent [e:] (i.e. [e:] that is a non-trigger, while non-transparent/vacillating [e:] is acceptable as either a trigger or not) was given in Table 3.10; for it to be variable, it must select for the morpheme-specific trigger scaling of [ɛ]. This is shown in Table 3.16, where the subscript "ɛ" indicates that the vowel morpheme-specifically behaves like [ɛ]. In such a case, the front and back variants of the suffix receive identical harmonic scores, meaning that either one could surface. Note that this result holds regardless of whether we assume that [e:] is also morpheme-specifically a better target in this case.

/ɔste:$_{ɛ}$k+nɛk/	*[ɤ:] 400	*[æ:] 400	IDENT-IO[back]$_{\text{ROOT}}$ 200	IDENT-IO[low] 20	*[-αBack]∞[αBack] 1	H
☞ a. ɔste:knɛk					−1*16(ɔ)*8(e:$_{ɛ}$) −1*16(ɔ)*8(ɛ)*1/2	−192
☞b. ɔste:knɔk					−1*16(ɔ)*8(e:$_{ɛ}$) −1*4(e:$_{ɛ}$)*16(ɔ)	−192
c. ɔsta:knɔk			−1	−1		−220

Table 3.16: Variable opacity of [e:]

Given that this trigger scaling factor is the usual one for [ɛ], the analysis predicts that back-[ɛ] roots will typically vacillate. Moreover, [ɛ] can morpheme-specifically have the trigger strength of a fully harmonic vowel; in such a case, the suffix will be consistently front, with [ɛ] opaque, similarly to the case of harmony with non-neutral vowels in Table 3.8.[33] In contrast, [i(:)] can never behave as vacillating or opaque, because its only morpheme-specific option is to have the trigger scaling of [e:], which generally behaves as transparent. Thus, the analysis is able to derive the height effect in transparency, in which high vowels are consistently transparent, mid vowels are sometimes variable, and low vowels are generally variable or opaque. This is due to the fact that morpheme-specificity in Hungarian is limited to a vowel behaving as "one category up", which I hypothesize could relate to the phonetic motivations for trigger and target; a vowel cannot morpheme-specifically take on the behaviour of one too distinct from its own phonetic properties.

Finally, multiple neutral vowels can also show vacillation, taking both front and back versions of the suffix. This is a simple effect of distance scaling, as shown in Table 3.17.

/ɔnɔli:ziʃ+nɛk/	*[ɤ:] 400	*[æ:] 400	Ident- IO[back]$_{\text{Root}}$ 200	Ident- IO[low] 20	*[-αBack]∞[αBack] 1	H
☞a. ɔnɔli:ziʃnɛk					−1*16(ɔ)*1(i:) −1*16(ɔ)*1(i)*1/2 −1*16(ɔ)*8(ɛ)*1/4	−56
☞b. ɔnɔli:ziʃnɔk					−1*16(ɔ)*1(i:) −1*16(ɔ)*1(i)*1/2 −1*2(i)*16(ɔ)	−56

Table 3.17: Variable opacity of multiple neutral vowels

---

[33] The number of examples like (2h), in which [ɛ] is consistently transparent, is quite small. These few roots are not captured here, but could be if we assume that morpheme-specificity could also occasionally go in the other direction, with [ɛ] behaving like [e:].

Note that this view of vacillation is a simplification; a more nuanced analysis, which is beyond the scope of the present chapter, would convert harmonic scores to percentages, because the Hungarian vacillation data is not as simple as each possible suffix form surfacing half the time. This issue will be discussed further in Section 6.5. Moreover, there is a role for paradigm uniformity in the behaviour of vacillating roots; generalizations of other aspects of paradigm uniformity in Hungarian harmony have previously been made by Törkenczy et al. (2013) and Rebrus et al. (2017). In this case, paradigm uniformity is required to ensure that vacillating roots behave the same way across multiple types of suffix vowels; specifically, we might otherwise predict suffixes with alternants containing [ɛ] or [eː] to behave differently than suffixes with fully harmonic vowels, because the target scaling factors are different. Again, a full implementation is beyond the scope of the present chapter. The critical aspect of the current analysis is that paradigm uniformity and morpheme-specificity do not need to bear the entire burden of explanation. Instead, the explanation here is shared by vowel-specific factors, which are less arbitrary because they are motivated by cross-linguistic and phonetic evidence.

To summarize, this section and those preceding it have demonstrated that the target-focused framework developed in Section 5 maintains the ability to capture the crucial facts about Hungarian harmony that have been previously analyzed, while also extending to the gradient nature of neutrality in the language. The specific numbers given here work, but what is crucial is that all of the relative weighting requirements are satisfied, specifically in terms of the relative ranges of trigger and target scaling factors for the different vowels.

6.5 Further directions for Hungarian

While it captures much of the gradient variability in Hungarian, this analysis still abstracts away from some of the more problematic aspects of Hungarian harmony. In this section, I consider how the analysis could be extended to additional data for which a full treatment is beyond the scope of this chapter.

First, as mentioned in Section 2, Hungarian has some anti-harmonic roots, which contain only front unrounded vowels, yet take back suffixes. An adequate analysis of these roots should explain why they all contain only front unrounded vowels, and why the majority of them contain [i(:)]. This patterning exactly parallels the other aspects of neutrality (invariant suffixes and transparency), and so should receive a similar account. One possibility in the framework here is that all alternating Hungarian suffixes are by default [+back], and that there is additional variability in the trigger strength of [i(:)], and to a lesser extent [e:]. If the trigger strength of these vowels can, in specific morphemes, be low enough to not trigger harmony at all, then default suffix vowels should occur; if the default is [+back], then anti-harmony is derived. The assumption of a default [+back] for suffixes changes a number of aspects of the above analysis, in particular the generalizations related to [e:] sometimes undergoing harmony in suffixes, but it is worth exploring as a future direction for explaining anti-harmony.

A further property that this analysis could be extended to deal with is stem/suffix interactions. Research has shown that of the vacillating stems that can take either front or back suffixes, stems vary in terms of their preference for front or back, and this further varies by the specific suffix attached to the stem. For example, Törkenczy (2013) notes that the root [fotɛl] ‘armchair’ more often takes the back version of the illative suffix [bɔ]/[bɛ], while the root [kontsɛrt] ‘concert’ takes only the front version of this suffix. With the plural [ok]/[ɛk], [fotɛl] instead takes the front version more often, while the root [hɔvɛr] 'friend' almost always takes the back version. In other words, the preferred choice for vacillating cases depends on both the stem and suffix.[34] Since the present analysis deals with morpheme-specific properties, a more nuanced version of it could be derived, in which there is some correlation between the harmonic scores of candidates and their percentage frequency, and this could depend on both the specific stem and suffix. This direction ties in with the general comment about the nuances of vacillating roots as discussed in Section 6.4. Moreover, it could reduce the need for an appeal to paradigm uniformity to explain the similarity in behaviour of vacillating roots across suffixes, regardless of whether the suffix alternants contain [eː], [ɛ], or only harmonic vowels (in comparison to some sort of word-listing/word-based morphology for at least these cases).

Recent literature on Hungarian has noted that invariant and alternating suffixes with [eː] behave differently with respect to vacillating roots: the invariant ones (e.g. [eː] in Table 3.14) are transparent, in that the root+suffix combination remains vacillating, while the alternating ones (e.g. [naːl]~[neːl] in Tables 3.13 and 3.15) are opaque, in that the root+suffix combination can only take front suffixes (e.g. Törkenczy et al. 2013). As Törkenczy et al. (2013) note, part of the explanation can be due to paradigm uniformity. In the present analysis, another factor to consider is the correlation between morpheme-specific trigger and target scaling. It is clear in Hungarian that a correlation between trigger and target behaviour exists in the general case: the height effect for triggering follows the exact same pattern as the height effect for being targeted. Since this analysis treats morpheme-specific variation as bumping a vowel into the behaviour of a vowel

---

[34] A reviewer notes that this can be influenced by whether the suffix is consonant-initial, giving the examples haver-ok (*haverek), férfi-as (*férfi-es), but haver-nak/nek, férfi-val/vel. This pattern is reminiscent of several other languages in which the harmony process is influenced by whether an affix is C-initial or V-initial, such as Alur (Kutsch Lojenga 1991; Chapter 7). There is no standard way of dealing with such cases, and it is beyond the scope of the present chapter to account for it in Hungarian. However, future work on Hungarian should examine the influence of the suffix-initial segment on the preferred suffix choice with vacillating stems.

one level better on the trigger or target scale, we predict the possibility that in cases where [e:] behaves like [ɛ] as a target, it could also behave like [ɛ] as a trigger. As such, alternating [e:] should behave differently than invariant [e:], as the former is morpheme-specifically a better trigger than the latter. This possible correlation between trigger and target factors, including in morpheme-specific cases, should be explored further in additional research.

A question arising from this analysis is why the trigger and target scaling can change morpheme-specifically. It is clear that they must, with limited options given the type of morpheme, because the variable patterning in Hungarian shows morpheme-specific properties in addition to vowel-specific ones. A direction to explore here is the effect of loans: vacillating back-[e:] roots tend to be loans, as do many of the roots in which [ɛ] co-occurs with a back vowel, regardless of their behaviour. A potential hypothesis is that loanwords come into the lexicon with trigger and target scaling that may not match the typical pattern of native words, but that a degree of closeness to the phonetically motivated behaviour of native words is required (i.e. the one category limit). This idea could also potentially be extended to "disharmonic" roots, in which back vowels co-occur with front rounded vowels. As noted, such roots are always recent loans, and perhaps the factor that prevents root-internal harmony from applying is a difference in the scaling factors for these vowels. Thus, examining the connection between loans and (dis)harmonic behaviour in Hungarian, in relation to the issue of "one category up" variability, is a worthwhile future direction.

7 Implications and discussion

This chapter has proposed a novel, target-oriented account of Hungarian harmony, arguing that the data requires a new way of incorporating targets into an analysis. Specifically, neutrality is not simply a lack of ability to harmonize, but rather that the dispreference for disharmony is too weak to outweigh faithfulness. This approach to Hungarian has many implications for the study of neutrality and targets more broadly in vowel harmony, some of which are considered in this section.

7.1 Target conditions in other languages

Hungarian [e:] is not alone in the fact that it is generally neutral to harmony despite appearing to have the option to participate through pairing with [a:]. As noted in Section 1, vowels that are able to undergo harmony do not always do so; moreover, there is a degree of cross-linguistic consistency in the choice between neutrality and participation, regardless of whether a vowel is paired. Mayak (Chapter 9) is one such case: the low vowels [a] and [ʌ] are paired for [ATR] and participate in tongue root harmony in limited contexts, but they are generally neutral (Andersen 1999). Moreover, low vowels in tongue root harmony systems are often unpaired and neutral, as in languages like Yoruba, which has an opaque [a] but no [+ATR] low vowel (Archangeli & Pulleyblank 1994). Archangeli & Pulleyblank (1994) note that this kind of consistency is in fact a cross-linguistic pattern: paired vowels that are idiosyncratically neutral, as is the case in Mayak, tend to be the same types of vowels that are unpaired and neutral in other languages.

Any theory in which neutrality is based on the impossibility of undergoing harmony (e.g. where neutrality is derived purely by a constraint against the harmonic counterpart) will have difficulty explaining cases like Mayak. More importantly, they miss the generalization that which vowels are neutral tends to be the same cross-linguistically regardless of whether the neutral vowels are paired or unpaired. This consistency is clear evidence that vowel-specific, harmony-specific factors beyond inventory pairing (and the associated general markedness) affect whether a vowel is a target for harmony or is neutral, contra many claims in the literature.

The analysis presented here can also be extended to cases not commonly considered as neutrality, namely harmony that targets only certain vowels. Previous analyses of vowel harmony have also noted that harmony may be subject to target conditions. An example is the constraint HI/ATR “if [+high], then [+ATR]”, which Archangeli & Pulleyblank (2002) use to account for high vowels being preferred targets in Kinande ATR harmony. This mechanism for incorporating features of targets differs from the current proposal in that it directly references the harmony-independent notion of feature compatibility, rather than the quality of a vowel as a potential harmony target. These notions are clearly related, as feature compatibility affects the phonetic and cross-linguistic factors that determine target quality. Though the details of the present approach are novel, the idea that harmony may be subject to target conditions is not, and target-specific harmony is necessary outside of Hungarian.

It is worth noting that the approach of Archangeli & Pulleyblank (2002) differs from the one here, for instance in the case of low vowels in front/back harmony; neither LO/BACK nor LO/FRONT would capture that low vowels are good targets, the former because it cannot explain why the low back vowel [aː] undergoes fronting in front harmony contexts, and the latter because it is unmotivated both cross-linguistically and in Hungarian, in which low vowels tend to be back, in terms of the number of such vowels in the phonologically defined inventory. [35] In contrast, in the present approach, the phonetic motivation for low vowels to be good targets for front/back harmony is built into the theory’s prediction of their behaviour.

Overall, target conditions are necessary beyond Hungarian, and the present approach has

---

[35] While a target condition that refers to [low] cannot capture the Hungarian alternation between [aː] and [eː], one that refers to [round] could. Specifically, NON-ROUND/FRONT would be a motivated target condition in Hungarian and cross-linguistically; this is effectively the *ɨʌ constraint mentioned earlier. However, such an approach is unable to account for why the low vowel [ɔ], which is round, is also subject to harmony.

the ability to extend to other languages and types of harmony. With the cross-linguistic concept of target quality, this analysis accounts for generalizations missed by a featural compatibility approach to neutrality, and can capture more complex cases like paired neutral vowels.

7.2 Trigger/target interactions

An additional implication concerns the interactions between triggers and targets, considered in this subsection. Given the use of scaling for both triggers and targets, the present analysis predicts that trigger and target should be able to interact in such a way that harmony depends on both of them. In other words, we should find cases in which harmony occurs only with both a strong trigger and a quality target, or in which disharmony is permitted only with a weak trigger and a poor target.

This prediction is borne out in rounding harmony in Kyrgyz and Altai. Kaun (1995; 2004) shows that in rounding harmony, good triggers are vowels that are non-high and/or front, while good targets are those that are high.[36] In Kyrgyz and Altai, rounding harmony is obligatory except when the potential trigger is high and back and the potential trigger is low; only in such cases is disharmony permitted (Korn 1969). Table 3.18 summarizes, with examples from Korn (1969).

---

[36] It is worth noting that this generalization fails in Hungarian, where rounding harmony targets only mid vowels. However, rounding harmony is quite restricted in Hungarian (e.g. Törkenczy 2010).

Trigger	*Target*	*Harmony?*	*Kyrgyz example*	*Gloss*
Non-high	Non-high	Required	ot-ko	‘to the fire’
Non-high	High	Required	kök-tun	‘of the sky’
High (and back)	Non-high	Optional in Kyrgyz; absent in Altai	uluk-tan~uluk-ton	‘from the magnate’
High	High	Required	su-nun	‘of the water’

Table 3.18: Kyrgyz and Altai harmony pattern

The only trigger/target combination where disharmony is permitted is precisely the one that is predicted: high, back vowels are weak triggers, while low vowels are poor targets. Thus, Kyrgyz and Altai manifest a trigger/target interaction predicted and easily accounted for by the present framework.

Nonetheless, the more common trigger/target relationship in rounding harmony is that trigger and target must agree in height, which is known as parasitic harmony and occurs for example in Yokuts (Kaun 1995). On the surface, this pattern would be problematic under the present analysis; as Kaun (1995) notes, good triggers of rounding harmony are non-high, while good targets are high, which means that height agreement involves either a bad trigger or a bad target. The approach here, without modifications, therefore predicts that if harmony occurs in cases of height agreement, then it should also occur when both the trigger and target are good.

However, this problem only occurs under the assumption that what it means to be a good trigger is independent of what it means to be a good target, and vice versa. Particularly in rounding harmony, this assumption is false; as Kaun (1995) argues, rounding is realized distinctly on high versus non-high vowels, meaning that height agreement allows for uniformity in the phonetic realization of [round]. Based on this argument, the relationship between trigger and target scaling factors should be dependent, which would permit an account of parasitic harmony. Formally, this could be implemented with either a single scaling factor for each trigger/target pair or with the addition of a scaling factor for the interaction effect.

To summarize, the prediction that trigger strength and target quality can interact is borne out in Kirghiz and Altai rounding harmony. Parasitic rounding harmony, which appears at first glance problematic for this approach, is in fact possible to analyze; it requires the addition of an interaction factor, which is motivated by the fact that triggers and targets are not independent.

8 Conclusions

To conclude, this chapter proposes a new view of participation in vowel harmony, in which neutrality results not from featural incompatibility, but rather from the vowel-specific drive to undergo harmony being too weak. Under such a perspective, vowels are neutral when they are poor targets of harmony, and so the drive for them to be targeted is insufficient to force unfaithfulness.

This view of neutrality is motivated here by an under-analyzed alternation in Hungarian between harmonic [a:] and generally neutral [e:], which challenges many analyses of Hungarian harmony. Under the standard view, Hungarian (non-low) front unrounded vowels are neutral because they lack harmonic counterparts. Such an approach leaves completely unexplained the fact that [a:] also lacks a direct harmonic counterpart, yet consistently participates in harmony and alternates with neutral [e:]. This chapter explains this behaviour, by appealing to the cross-linguistic and phonetically motivated fact that low vowels are better targets of front/back harmony than higher vowels. The concept of target quality directly incorporates the phonetic basis, within the definition, and allows for an account of the gradience of neutrality.

In doing so, the approach captures this broader typological height generalization, which goes beyond Hungarian, but also reflects the Hungarian-internal distinction in degree of neutrality for [i]/[iː], [eː] and [ɛ]. In contrast, in other formal analyses, the effects of height on harmony participation are determined arbitrarily and in a way that does not reflect the gradience of the behaviour, such as by stipulating that [ɛ] is harmonic.

This analysis has broad implications to the study of neutrality and target conditions; it has the explanatory power to account for complex patterns in a simple way, yet is appropriately constrained in its predictions by the phonetic motivations of trigger strength and target quality. Beyond Hungarian, the approach can be extended to deal with neutral paired vowels like in Mayak (Chapter 9), target asymmetries like in Kinande, trigger/target interactions like in Kirghiz and Altai, and many other patterns that require distinct, complex accounts in other frameworks. Here, instead, the central insight that captures multiple diverse phenomena is that whether harmony applies and whether disharmony matters crucially depends on the nature of the trigger and target.

.

Chapter 4: Mora-level harmony and target quality as a solution to the Votic ‘paradox’

1 Introduction

In the recent theoretical literature, Votic has been described as having a paradox, in which [i] seems to behave as unspecified for [back] in some phonological patterns, but as [-back] in others (Blumenfeld & Toivonen 2016; Hall 2018). Specifically, [i] is neutral to front/back harmony, yet participates in a pattern in which laterals surface as [l] before front vowels but [ɫ] before back vowels; the lateral is consistently [l] before [i]. These lateral alternations force a view in which [i] is [-back], which is a problem for any view in which neutrality to vowel harmony results from a lack of [back] specification.

Previous analyses have proposed round-about ways of solving this apparent paradox. Blumenfeld & Toivonen (2016) (henceforth B&T) draw a distinction between ‘strong’ and ‘weak’ specifications for [back], where [i] is only weakly specified because it is not contrastive for [back]; they suggest that harmony is relativized to strong specifications, while lateral alternations reference all specifications. As B&T note, this idea is reminiscent of work from Nevins (2010), Vaux (2004), Calabrese (2005), and others, in which individual processes/rules can be arbitrarily specified as seeing only contrastive specifications for [F], only marked values of [F], or all specifications of [F]. This analysis, as implemented in B&T, suffers from a number of problems, most notably that the definitions of ‘strong’ and ‘weak’ that are necessary to evaluate harmony reference contrast, which is an inventory property that is also determined within the grammar. Moreover, because laterals in lateral-vowel sequences are determined independently from harmony by the vowel, and because laterals themselves participate fully in harmony, the approach predicts that [i] can control harmony in cases in which it is preceded by a lateral. Hall (2018) suggests a different solution, in which [i] is not specified for [back], but for [coronal], and both [back] and [coronal] can control lateral alternations; harmony is still in the feature [back]. This approach is non-ideal, because it requires two distinct features motivating lateral-vowel agreement, and laterals do not agree with other following coronals (i.e. coronal consonants).

Both of these approaches also fail to extend more broadly to other cases in which there is an apparent mismatch between contrast and neutrality, including in front/back harmony in another Uralic language, Hungarian (see Chapter 3). In Hungarian, all front unrounded vowels are unpaired for front/back, yet there is a height effect, in which [i] is consistently neutral, [e:] is usually (but not always) neutral, and [ɛ] is usually (but not always) harmonic (e.g. Hayes & Londe 2006). The categorical distinctions drawn by B&T and Hall to explain neutrality of [i] in Votic cannot be extended to variability in harmony participation among unpaired vowels in Hungarian.

In this chapter, I propose to adapt the Hungarian analysis from Chapter 3 for Votic. In this approach, neutrality is not a result of different specification for the harmonic feature, but the result of certain vowels being poor triggers and targets of harmony, such that the drive to harmonize in combinations involving those vowels is insufficient to force unfaithfulness. Previous arguments have suggested that [i] is a particularly poor trigger and target of front/back harmony, consistent with its neutrality in Votic (and Hungarian; see e.g. Benus & Gafos 2007). In the present type of account, [i] is still specified as [-back]; it is simply that the drive for it to participate in harmony is not strong enough to make harmony actually occur. Importantly, though, this approach does not rely at all on [i] being specified differently from other front vowels; it is in fact [-back], and can behave as such in other phonological patterns like lateral alternations.

Beyond this analysis, I also develop a solution for the complexities in the behaviour of laterals in Votic. While laterals immediately preceding a vowel agree in backness with that vowel (e.g. [əɫud] 'beer', [elæ:] ‘to live’), laterals in other positions agree with the harmonic context (i.e. they are velarized/back in back words, like [na:tɫ] 'chicken pen', and plain/front in front words, like [ve:l] 'more'). The specific case where this pattern becomes problematic is in back words containing neutral [i]. In such words, the front [l] surfaces if it immediately precedes [i] (as in [lidna:] ‘town.ILL’; [əlimma] 'we were'), but the back [ɫ] surfaces otherwise, even if it immediately follows [i] (as in [iɫma] ‘without’) or if the nearest (but non-adjacent) following vowel is [i] (as in [taɫvi] ‘winter’). These examples show that the lateral must actually participate in the harmonic alternations, but does not participate if the immediately following vowel is [i]. In other words, whether laterals are excluded from harmony or not depends on whether they directly precede a vowel, and if so, whether that vowel is [i].

I propose that in Votic, harmony operates on the level of moras, including both vowels and coda consonants, while onset-nucleus [back] agreement is a separate process. As such, onset consonants are not subject to harmony, and will therefore simply agree with the following vowel. In contrast, if the lateral does not immediately precede a vowel, it is a coda, which I presume has its own mora. In such cases, the lateral can participate in harmony in the same way as vowels. Since [back] is primarily a vocalic feature, I propose that there exists a separate [back] faithfulness constraint referencing vowels, but not a similar one for consonants, which allows this analysis to capture the generalization that while consonants may be targets of vowel harmony, they do not trigger it in Votic, and generally do not trigger it in other languages either (except in exceptional cases). Specifically, in situations where harmony is required between a vowel and consonant, and no other factors intervene, the vowel will always win, because maintaining its input specification is more important than maintaining that of the consonant. This approach is therefore able to capture all of the complications of the Votic pattern, without the pathological predictions of previous accounts.

This chapter is organized as follows. In Section 2, I describe the relevant front/back phonological processes in Votic, namely harmony and lateral alternations. In Section 3, I discuss existing analyses and argue against them. Section 4 proposes a novel analysis, adopting the target-oriented approach from Chapter 3 as well as the mora-based account outlined above. Section 5 discusses and concludes.

2 Front/back phonological processes in Votic

Votic has two major front/back phonological processes of interest: vowel harmony and lateral alternations. The distinct behaviour of [i] within these patterns has resulted in an apparent paradox, if neutrality is achieved representationally: [i] is neutral in harmony, but behaves as front for the purposes of lateral alternations. This section describes these patterns, focusing in particular on the behaviour of [i].

2.1 Inventory and feature assumptions

The inventory of Votic monophthongs is given in Table 4.1. Note that all vowels also show length contrasts, and the vowel [ɨ] occurs only in Russian loans. The transcription of these vowels differ somewhat in various resources; I mostly follow Blumenfeld & Toivonen (2016), except using the IPA symbols [ø] and [æ] for their [ö] and [ä] respectively.

	Front		Back	
	Unrounded	Rounded	Unrounded	Rounded
High	i	y	(ɨ)	u
Mid	e	ø	ə	o
Low	æ		a	

Table 4.1: Votic vowel inventory (Ariste 1968)

Votic also has an extensive inventory of diphthongs composed of these vowels, listed in Ariste (1968); these will be discussed in more detail in reference to the harmony system, but are listed in (1) for convenience. Like in Ariste (1968), where the list comes from, I split them into diphthongs that can occur in main stress position (initial syllables) and those occurring further in the word (non-initial syllables); I additionally organize them into [+back], [-back], and disharmonic diphthongs. I assume that diphthongs have two moras in Votic and that each mora has its own [back] specification (as well as other feature specifications), consistent with the feature specification of the relevant vowel.

(1)

Diphthongs in main-stressed syllables:

[+back]: əa, uə, ao, uo, au, ou, əu

[-back] : ye, æi, ei, øi, yi, yæ, æy, ey, øy, iæ

Disharmonic: ai, oi, ui, əi, io, eu, iu,

Diphthongs further on in the word:

[+back]: oa, ua, əa, au, əu

[-back]: iæ, øæ, yæ, æi, ei, yi, æy, ey, iy

Disharmonic: ia, æu, eu, iu, ai, oi, ui, əi

The consonant inventory of Votic is given in Table 4.2.

		Labial	Dental	Palatal	Velar
Stops	Voiceless	p	t		k
	Voiced	b	d		g
Fricatives	Voiceless	f	s	ʃ	x
	Voiced	v	z	j ʒ	
Affricates	Voiceless		ts	tʃ ʃtʃ	
	Voiced			dʒ	
Nasals	Voiced	m	n		ŋ
Laterals	Voiced		l	ɫ[37]	
Trills	Voiced		r		

Table 4.2: Votic consonant inventory (Ariste 1968)

Aside from in diphthongs, Votic does not have glides; Ariste (1968) does notate a segment as [j], but in describing it, writes "In Votic, [j] is not a semi-vowel, as it is in Estonian, but a real fricative" (p.9). This segment moreover lacks any sort of back counterpart. In fact, with the exception of laterals, all Votic consonants can be described without the feature [back], using the place features [labial], [coronal], and [dorsal] to distinguish major place, and [anterior] to distinguish 'dentals' from 'palatals' among fricatives and affricates, although 'palatal' fricatives

---

[37] I place this lateral in the palatal column following Ariste (1968). Given that it is velarized, it might be better placed in the velar column, or in the dental column with [l], given that the difference may be a matter of secondary rather than primary place of articulation.

occur primarily in loans.[38] Ariste does not give a description of the articulation of Votic laterals, except to say that the back one is similar to Russian and the front one to Estonian. I assume that laterals are the only consonants in Votic that can take [back] specifications, and that in fact they must be specified for [back]; the formal implementation of these assumptions is beyond the scope of this chapter. The fact that only laterals can take these distinctions in Votic is perhaps related to their sonority. It will be explained below why the analysis as formulated in this thesis requires not specifying most consonants for [back], rather than having them be poor triggers/targets of harmony.

2.2 Harmony

Like many other Finno-Ugric languages, Votic has a system of front/back harmony, in which stems typically consist of either all front vowels or all back vowels, and suffixes alternate in backness depending on the stem vowels. Some examples are shown in (2). All examples in this paper are taken from Ariste (1968) or Ahlqvist (1856).

(2)

(a) sømæ-ssæ ‘eating-ELA’[39]

(b) vævy-ssæ ‘son-in-law-ELA’

(c) ro:pa-ssa ‘porridge-ELA’

(d) vasara-ssa ‘hammer-ELA’

(e) tø:-zæ ‘work-INE’

---

[38] The ‘dental’ and ‘palatal’ terminology is what is used by Ariste (1968) for distinctions like [ts] versus [tʃ]. Since the exact phonetic implementations are unclear, I follow this terminology throughout.

[39] While in Finnish, [ssa]~[ssæ] marks the inessive and the elative is [sta]~[stæ], this suffix [ssa]~[ssæ] is in fact the elative in Votic. The Votic inessive is [za]~[zæ] (Ariste 1968).

(f) ma:-za ‘country-INE’

Again as in many related languages, [i] in Votic is neutral to harmony; it can co-occur with both front and back vowels within roots, it does not alternate when it occurs in suffixes, and it is skipped over in determining the value of alternating suffixes. The examples in (3) illustrate; in all examples, the suffix value for [back] is determined by the vowel that precedes the [i]. Additionally, as shown in (c-d), the plural suffix [i] does not alternate and is skipped in determining the value of the elative suffix.[40]

(3)

(a) tæi-ssæ ‘louse-ELA’

(b) poiga-ssa ‘boy, son-ELA’

(c) lyhy-i-ssæ ‘short-PL-ELA’

(d) su-i-ssa ‘mouth-PL-ELA’

Thus, at first glance, the Votic front/back harmony system appears to be a classic case in which contrast and neutrality match: [i] has no back counterpart (in native vocabulary) and is neutral to front/back harmony.

Before turning to additional facts, it is worth briefly discussing diphthongs in relation to the harmony system. Most Votic diphthongs are harmonic, and of the ones that are disharmonic, the majority involve [i] as either the first or second vowel in the diphthong (cf. the list in (1)).

---

[40] Examples of the plural in Ariste (1968) generally involve [i] immediately after a vowel (hence forming a diphthong). Whether such examples exist with [i] on its own as a non-diphthong is unclear. This is potentially important, except that there are other cases in which [i] is obviously transparent, such as when it occurs between back vowels in [əlimma] 'we were'; it is likely that the [i] in this example is in fact the plural suffix, followed by a 1PL subject agreement suffix.

This is the behaviour we would expect given vowel harmony in the language: diphthongs involving [i] can involve disharmony, since [i] is neutral, but diphthongs involving other vowels are harmonic. Note that I assume that there is no substantive difference in diphthongs between vowels and glides, as is typical in the notations used for Votic diphthongs.

However, there are two disharmonic diphthongs in Votic that do require attention: [eu] and [æu]; the former can occur dialectally in initial (stressed) syllables, and both can occur in non-initial positions. It is unclear in which dialects [eu] occurs initially, or whether such dialects have the back counterpart of [e]. I therefore focus on the non-initial situations. It is possible that these are the result of morphological complexity, in that all of the examples in Ariste (1968) are of reflexive or intransitive imperatives ([lisæugo:] 'let it increase', [peittæuga:] 'get hidden (pl)', and [ehteuga:] 'adorn, embellish yourselves'). According to Ariste (1968), there is a verbal derivational suffix [u]~[y] in Votic that marks reflexive/intransitive, which usually occurs as [y] following front vowel stems; I suggest that perhaps exceptionally it is [u] in these cases. It is worth noting that the imperative suffix, which is also present in all of these examples, never alternates harmonically, but rather is always back (e.g. [-go:], [-ga:]); worth exploring is whether there is any connection between this and the presence of the [u] in the immediately preceding diphthongs. For the remainder of this paper, I consider these diphthongs a case of morphological exceptionality and assume that all diphthongs in the language are harmonic, except those involving [i].

2.3 Lateral alternations

Where the Votic facts differ from what is typically discussed for related languages is that in Votic, harmony is not the only phonological process that involves the feature [back] in Votic;

lateral alternations are affected by whether adjacent vowels are back or front.[41] Specifically, when there is an immediately following vowel, plain [l] occurs if that vowel is front, while velarized [ł] occurs if it is back. Examples are given in (4).

(4)

(a) elæ: ‘to live’

(b) tʃylæ-llæ 'village-ADE'

(c) poiga-łła 'boy, son-ADE'

(d) əłud 'beer'

Moreover, in situations with no following vowel, the lateral agrees with the harmonic context, such that these laterals are [l] in front words and [ł] in back words. Examples are given in (5).

(5)

(a) ve:l 'more'

(b) tʃylve-ttæ ‘to wash’

(c) næltʃ ‘hunger’

(d) na:tł 'chicken pen'

(e) kəłmaʒ ‘third’

(f) łaułłu-łł ‘sing-ADE’

---

[41] While this pattern is not often discussed, Finnish in fact does something similar with allophony of ‘clear’ versus ‘dark’ laterals depending on the vocalic context, for instance with dark [ł] in cases like [iłta] 'evening' vs. clear [l] in cases like [silmæ] 'eye' (Gunnar Ólafur Hansson, p.c.). Also note that Nevins (2010) has discussed Finnish as having an analogous problem with the fact that [i] triggers assibilation in a preceding /t/ (making it [s]), despite being neutral to harmony. Votic [i] also triggers palatalization in a preceding /k/, raising the same kind of issue of visibility of [-back] for [i] (see e.g. Hall 2018), but I focus here only on lateral alternations, and only in Votic.

Interestingly, in this phenomenon, [i] behaves distinctly as a front vowel, in that any lateral immediately preceding it surfaces as [l], even in a back vowel word. This is shown in (6); note how all examples have a plain [l] preceding [i] despite the fact that all other vowels in the word are back, and that velarized [ɫ] does occur in example (c), but not immediately preceding [i]. Critically, then, [i] behaves like other vowels in controlling alternations of immediately preceding laterals, and it behaves categorically as a front vowel in this process.

(6)

(a) əlimma 'we were'

(b) tuli-i-sə: 'fire-PL-ILL’

(c) lintu-i-ɫɫa 'bird-PL-ALL'

The lateral in a back word surfaces as [l] only if it immediately precedes [i]; if it follows [i], or if consonants intervene between the lateral and the [i], then it surfaces as [ɫ], as shown in (7).

(7)

(a) iɫma ‘without’

(b) taɫvi ‘winter’

I suggest that this distinction is about onsets versus codas: onset laterals do not participate in harmony, but instead take their [back] value directly from the vowel in the following nucleus, while coda laterals are full participators in the harmony system. This account will be formalized in Section 3.

2.4 Interim summary: a backness ‘paradox’

These facts result in a problem for any analysis of Votic harmony that claims that [i] is transparent because it is not phonologically front. Specifically, we have an apparent contradiction: [i] is required to be [-back] to explain the lateral alternations discussed in Section 2.2, but under many analyses of transparency, it also cannot be [-back] because it is skipped over by backness harmony described in Section 2.1. Several previous analyses have attempted to address this ‘paradox’; I present my own analysis in Section 3, then address the problems with previous analyses in Section 4.

3 Novel analysis

The Votic facts can be accounted for fairly simply by extending the analysis of Hungarian from Chapter 3. Specifically, I propose that the neutral status of Votic [i] is due not to non-contrastiveness for [back]; instead, like in Hungarian, it is due to the poor quality of [i] as a participator in front/back harmony. As discussed in Chapter 3, such a view is phonetically and typologically motivated, as well as independently necessary in cases other than Votic. Under this perspective, whether [i] participates in harmony is independent from whether it has the feature [-back], contrastively or otherwise, and is therefore independent from whether we predict it to participate in lateral alternations. Thus, the account that is necessary for Hungarian also provides a simpler way of understanding the apparent contradiction in the status of [i] in Votic.

In this section, I develop this type of approach for Votic. I argue that the behaviour of Votic laterals is suggestive of [back] harmony being operative on the level of moras in Votic, such that laterals immediately preceding a vowel are not subject to harmony requirements in the same way as other laterals, given that onsets are not moraic. Laterals in other contexts are moraic codas, and therefore harmonize the way vowels do. As noted in Section 2.1, I assume that laterals in Votic are the only consonants that bear [back] specifications, and that they must do so. Combining this analysis of mora-level harmony with a target-based approach to neutrality, the Votic facts are captured straightforwardly, in a way that is less stipulative and more easily extendable to other cases of harmony than previous accounts (which will be discussed in Section 4).

3.1 Target-based harmony

In order to account for the neutrality of [i] in Votic harmony, I adopt the target-based approach from Chapter 3. In this account, neutrality results from an insufficient drive for a vowel to be a trigger and/or target of harmony, rather than from differences in its featural specification. As such, [i] can be fully specified as [-back], yet the phonetic properties that make it a poor trigger and target of front/back harmony mean that the drive to apply harmony when the trigger or target is [i] is too weak to force unfaithfulness.

Following the analysis in Chapter 3, which uses constraints adapted from Pulleyblank (2004), I adopt a constraint that bans [back]-disagreeing sequences at any distance, with violations counted based on the second member of each disagreeing pair. I propose that in Votic, unlike in what was discussed for Hungarian[42], the relevant interactions between [back] features

---

[42] However, there is no obvious reason for the mora approach to not also be the case in Hungarian, given that we have no reason in Hungarian to specify [back] on any consonants, and vowels are moraic.

occur at the mora level: not specifically on vowels, but on any segment that can be specified for [back] (i.e. vowels and laterals, as well as the glides in diphthongs).[43] The details of the analysis relating to laterals will be discussed beginning in Section 3.3; for now it suffices to know that mora is the relevant level where agreement in [back] specifications is checked. The reformulated constraint is defined in (8).

(8) *[-αBack]∞[αBack]: Assign a violation for every mora that has the feature [αBack] and is preceded at any distance by a mora with the feature [-αBack].

This harmony constraint is scaled by trigger and target, in a way defined in (9) and (10), taken from Chapter 3 on Hungarian. This scaling means that the penalty for disharmony will differ depending on the specific trigger and target, where trigger scaling applies to the context of the constraint (the preceding [-αBack] mora that constitutes the environment) and target scaling applies to the focus (the following [αBack] mora for which violation or satisfaction is assessed). Every possible target is given a violation, and the trigger is defined to be the nearest possible preceding mora with a different [back] value.

(9) Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that x ≥ 1) for each degree *i* to which α is perceptually impoverished with respect to ±F.

---

[43] I treat glides as non-syllabic vowels, and they therefore behave exactly like the corresponding vowel. Notably, [j] behaves like [i]. Also, recall that no non-lateral consonants are specified for [back]; if they were, they would need the appropriate trigger/target scaling. Consonants like velars or 'palatal' affricates are often thought to be specified for [back], but I assume that the relevant feature distinguishing velars is simply [dorsal], and that 'palatal' affricates are coronal and distinguished from 'dental' ones with the feature [anterior]. In some languages there is a velar/uvular distinction as part of front/back harmony, but Votic does not have these distinctions.

(10) Scaling factor: target quality

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that x ≥ 1) for each degree *i* to which β is a quality target (defined by phonetic factors discussed in Section 3) with respect to ±F.

Since the harmony constraint here refers to moras, rather than vowels, I assume that the relevant trigger and target are the head of the mora, defined as the most sonorous element of the mora. As such, if complex codas consist of only a single mora, which I assume, then the lateral will still be the head in examples like [na:tł] 'chicken pen'. While I assume that onsets do not have a mora and attach directly to the syllable, this is not necessary; indeed, if we were to instead adopt an analysis in which onsets share the mora of the nucleus (e.g. Hyman 1985), then the behaviour of any onset-vowel sequence would still be defined by the vowel.

As argued in Chapter 3 for Hungarian, [i] must have trigger and target scaling less than or equal to that of vowels that are lower and/or rounded. In Votic, since [i] does not harmonize while other vowels do, its trigger and target scaling must be strictly less than those of other vowels, so that in combinations involving [i], the penalty for disharmony is insufficient to create an unfaithful form. Note that unlike in Hungarian, [e] does not need to be scaled lower than other vowels in Votic for the available examples, since it behaves harmonically. However, B&T note that [e] is sometimes neutral, in which case it would be possible to implement a more nuanced version of this account, similar to the one used for Hungarian.

One additional markedness constraint is necessary to prevent [i] from harmonizing in Votic: one that blocks its harmonic counterpart, namely *ɯ, which is violated for every occurrence of a [+high,+back,-round] vowel. This constraint is never violated in native Votic vocabulary, although it is violated by the vowel [ɨ] in Russian loans.[44] I assume that some sort of lexical marking of loans or a co-phonology allows loans to violate this otherwise highly weighted constraint (e.g. Itô & Mester 1995; Inkelas & Zoll 2007).

Before turning to the necessary faithfulness constraints, which are somewhat complex, I show how these considerations can broadly derive the pattern we see in Votic, namely that vowels other than [i] harmonize, while [i] does not. Following what is typically assumed about front/back harmony in native words in Uralic (and Turkic) languages, I assume that the first vowel is the trigger of harmony (in the sense of which vowel is the locus of the contrastive [back] specification in the input that ends up determining the harmony of the whole word; a given harmony constraint violation may have trigger scaling determined by a non-initial vowel); this will be discussed more in relation to the faithfulness constraints below. For the constraints dealt with so far, I will adopt the following weights and scaling factors for Votic. Note that the specific numbers are arbitrary, as long as the conditions discussed throughout the following subsections hold.

(11) Constraint weights

*ɯ: weight 200

*[-αBack]∞[αBack]: weight 1

(12) Trigger scaling for [back] harmony constraint

1 for [i]

---

[44] The choice of transcription for a high, back, unrounded vowel in languages with front/back harmony is variable, depending on narrowness of transcription or descriptive tradition on the language. As such, I use [ɯ] and [ɨ] as equivalent, with both expressing a vowel that is [+high,+back,-round].

4 for other vowels

(13) Target scaling for [back] harmony constraint

1 for [i]

4 for other vowels

In Table 4.3, we see that with these constraints and scaling factors, vowels other than [i] will harmonize. Specifically, there are currently no constraints that disprefer harmony in this case, and any kind of disharmony, whether it involves the root vowel or the suffix vowel, violates the harmony constraint. If only the suffix vowel is disharmonic, as in candidate (b), there is a single violation of *[-αBack]∞[αBack], counted for the [a] that is [+back] with a preceding [-back] vowel [æ]. On the other hand, if harmony skips the second root vowel to target the suffix only, as in candidate (c), there are two violations of this constraint, one for the [a] that is [+back] and preceded by the [-back] vowel [ø], and the other for the [-back] vowel [æ] that is preceded by the [+back] vowel [a]. A candidate in which the initial vowel is front and the others are back also violates the harmony constraint twice, once for each of the back vowels, as in candidate (d). All of these violations are scaled for the appropriate triggers and targets. As mentioned, the faithfulness constraints ensuring harmony propagates from the first vowel will be addressed shortly.

/søma-ssa/	*ɯ 200	*[-αBack]∞[αBack] 1	H
☞ a. sømæ-ssæ			0
b. sømæ-ssa		-1*4(æ)*4(a)	-16
c. søma-ssæ		-1*4(ø)*4(a) + -1*4(a)*4(æ)	-32
d. søma$_1$-ssa$_2$		-1*4(ø)*4(a$_1$) + -1*4(ø)*4(a$_2$)	-32

Table 4.3: Harmony triggered by initial vowel

In contrast to the above situation, Table 4.4 shows that [i] is neutral (and transparent) to harmony; a suffix vowel will surface in agreement with the vowel preceding [i]. Specifically, harmony that targets [i] will be ruled out by the highly weighted constraint against its harmonic counterpart, as in candidate (c). Since I assume faithfulness to the first vowel, a disharmonic form is therefore guaranteed. The question, then, is which vowel controls the harmonic value of the vowel following the [i]. Here, the distinction comes from the fact that [i] is a poor trigger of harmony, while other vowels are not. If the suffix vowel is back, as in candidate (a), then the relevant violation of *[-αBack]∞[αBack] is scaled for trigger [i] and target [a]; if the suffix vowel is front, as in candidate (b), then the violation is scaled for trigger [u] and target [æ]. Since not propagating harmony from an [u] is worse than not propagating it from an [i], the latter violation is worse, and so [i] is transparent.

/su-i-ssæ/	*ɯ 200	*[-αBack]∞[αBack] 1	H
☞ a. su-i-ssa		-1*4(u)*1(i) + -1*1(i)*4(a)	-8
b. su-i-ssæ		-1*4(u)*1(i) + -1*4(u)*4(æ)	-20
d. su-ɯ-ssa	-1		-200

Table 4.4: Transparency of [i]

3.2 Faithfulness

Turning now to necessary faithfulness constraints, I assume that in Votic, the general constraint Ident-IO[back] is so low-weighted as to have no effect in the language (i.e. has near-zero weight).[45] I therefore will not include it in tableaux; however, this constraint would of course be violated by any [back] changes from input to output. As such, even the candidates in the tableaux below that do not violate any of the constraints shown here would in fact violate that constraint, which has no effect on the harmony system in Votic.

In order to determine which [back] value wins when multiple vowels are required to harmonize, I assume positional faithfulness to the initial vowel (e.g. Beckman 1998). This constraint, Ident-IO[back]$_{V1}$, is violated for any changes in the feature [back] between input and

---

[45] Worth noting is that it is irrelevant whether this constraint has higher weight than a constraint like *y, because even in non-initial syllables, the vowel-specific faithfulness constraint to be developed later in this subsection will ensure that /y/ can surface faithfully in any position in Votic. Without such a constraint, we would need general [back] faithfulness to outweigh *y, otherwise /y/ would surface as [u] in non-initial syllables. The issues of the weights of these very low-weighted constraints in Votic is beyond the scope of this paper.

output for the first vowel. I adopt a constraint for the first vowel (defined by the feature [-consonantal]), not the initial syllable, because there is no evidence that laterals in initial syllables maintain a faithful specification; instead, they consistently alternate based on vowels, such that it is impossible to determine their underlying form. Specifically, coda laterals, which are determined by harmony, alternate harmonically, rather than remaining faithful to some underlying representation, even in initial syllables, as in [ve:l] 'more' and [miltinle:b] ‘some kind of’, versus [na:tł] 'chicken pen' and [iłma] ‘without’. If harmony were determined by initial syllable faithfulness, we would expect the lateral in all such cases to surface as contrastively back or front (depending on its underlying [back] specification), particularly in the examples where the first vowel is [i].

An analysis including V1 faithfulness is consistent with the generally progressive direction of Votic harmony, as well as other cases of front/back harmony. Moreover, it is consistent with the fact that the vowel inventory of Votic, and often related languages, is larger in initial syllables than in non-initial ones (e.g. Kiparsky & Pajusalu 2003). In particular, the vowel [ø] typically occurs only in initial syllables in Votic (at least in native roots; Ariste 1968); a ranking of IDENT-IO[back]$_{\text{V1}}$ >> *ø >> IDENT-IO[back] would achieve this result (assuming a lower ranking of harmony). Since the inventory-related asymmetries are not the focus of the present paper, I omit tableaux for this generalization. However, the existence of this asymmetry does suggest that Votic preserves the [back] value of the vowel in the initial syllable to a greater degree than non-initial syllables, supporting this as the faithfulness constraint relevant to which vowel determines harmony. Since we can assume that this constraint is never violated except for hypothetical inputs with /ɯ/ as the first vowel, it must be highly weighted, but lower weighted than *ɯ; I assign it a weight of 100. With this constraint added in, the tableaux in Tables 4.2 and

4.3 will continue to choose the correct candidates, and will rule out any candidate in which the initial vowel is changed. This will be illustrated in more detail once all the faithfulness constraints are dealt with.

Thus far, the analysis is sufficient to deal with harmony among vowels in Votic, but not with the behaviour of laterals. Specifically, if V1 is [i], then we need a way to determine which segment controls harmony between laterals and vowels, as in examples like [iłma] 'without', versus [miltinle:b] 'some kind of'.[46] As these examples suggest, it is never the lateral that controls harmony here; laterals are not contrastive for [back] in Votic, yet we see both laterals occurring even when the lateral is the first segment specified for [back] besides an [i]. Instead, a non-onset lateral agrees with the nearest non-[i] vowel, even if that vowel comes later, as was shown in (6). Such examples suggest that vowels preferentially control harmony; laterals are targets, but never triggers in Votic. B&T claim that this is always true of consonants in vowel harmony systems, though there is evidence from Turkish suggesting that this is not necessarily true, but that the exceptions are truly exceptional cases (e.g. loans; see Clements & Sezer 1982 on Turkish palatalized consonants in harmony).

I will attribute this asymmetry between laterals and vowels as triggers to a stringency hierarchy of faithfulness.[47] Given that [back] is generally a vocalic feature, it is sensible to

---

[46] It is beyond the scope of this paper to analyze which segment determines harmony in situations where the first vowel is [i] and there are multiple subsequent vowels, such as CiCaCo, where harmony is required between the non-[i], non-initial vowels. The present analysis does not have a way of establishing which vowel controls harmony, because harmony proceeds from the initial vowel. It is possible that the addition of low-weighted root faithfulness would help here: if roots tend to be bisyllabic, then when the decision does not fall to the first vowel, it will fall to the root vowel. This is a broad issue for any analysis that deals with the root-internal harmony in Uralic languages. However, it is not clear to what extent this is a problem, given that we do not know what happens to hypothetical underlyingly disharmonic roots to make them harmonic.

[47] One might think that the theory developed here provides the tools to avoid doing this, by making all consonants poor triggers but some consonants possible targets, such that are penalized appropriately by the target and trigger weightings being developed here. Such an approach is only possible if one abandons the assumption that scaling is based only on the closest possible trigger. Indeed, if all consonants are participants, then often the closest possible trigger to a target vowel is a (non-participating) consonant, and so harmony should fail to propagate. If instead the

suppose that there exists a faithfulness constraint to [back] that is specific to vowels, IDENT-IO[back]$_V$, but no such constraint specific to consonants, only a general IDENT-IO[back].[48] Regardless of the ranking of these constraints, we predict that if vowels and consonants are required to harmonize in [back], the [back] value of the vowel will always win; it is more important to be faithful to that value. We therefore automatically predict no possibilities in which the lateral can control harmony; this will be illustrated in Section 3.3 with tableaux, but essentially, if two segments are required to harmonize and there is a faithfulness constraint preserving one but not the other, harmony will be triggered by the segment that is preferentially preserved.

Clearly, the constraint IDENT-IO[back]$_V$ is regularly violated in Votic, since vowels do harmonize. The key is for its weight to be such that harmony occurs between vowels despite this constraint, so that it only becomes relevant when one of the interacting segments is a lateral. I will assign it a weight of 10; the relative weightings will be established through the tableaux.

Now that we have added the faithfulness constraints, we can return to the previous tableaux to demonstrate that we still predict the correct forms for harmony among vowels and transparency of [i]. Table 4.5 again illustrates the basic case of harmony.[49] In order to strengthen the argumentation, I assume that all vowels except the initial one are underlyingly back, and the result still surfaces as front, in agreement with the initial vowel. This result requires the penalty

---

framework were revised such that all disagreeing pairs are penalized by the harmony constraint, with appropriate scaling, then it would be possible to account for the behaviour of consonants using the trigger/target approach here.

[48] If we recognize a head versus non-head distinction at the mora level, between segments and where vowels are the head of the syllable, then the vowel-specific constraint could be reformulated as positional faithfulness, generalizing on how positional faithfulness to stressed syllables is often reformulated as referring to the head of a foot.

[49] The elative suffix /-ssA/ is perhaps not the best case for a simple illustration, since the coda of the geminate creates an intervening mora. I use the elative anyways because it is the illustrative case used in previous theoretical work on Votic. As noted, I simply assume that non-lateral coda moras are not relevant to harmony because they do not have [back] specifications.

for harmony between vowels other than [i] to be larger than that for violations of IDENT-IO[back]$_V$.

/søma-ssa/	*ɯ 200	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_V$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. sømæ-ssæ			-2		-20
b. sømæ-ssa			-1	-1*4(æ)*4(a)	-26
c. søma$_1$-ssa$_2$				-1*4(ø)*4(a$_1$) + -1*4(ø)*4(a$_2$)	-32
d. soma-ssa		-1	-1		-110

Table 4.5: Harmony triggered by initial vowel

Table 4.6 shows that [i] is neutral to harmony; a suffix vowel will surface in agreement with the vowel preceding [i]. This requires the penalty for disharmony between vowels other than [i] (e.g. u-æ, in the present example) to be larger than the penalty for the combination of IDENT-IO[back]$_V$ plus harmony in which the trigger is [i].

/su-i-ssæ/	*ɯ 200	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. su-i-ssa			-1	-1*4(u)*1(i) + -1*1(i)*4(a)	-18
b. su-i-ssæ				-1*4(u)*1(i) + -1*4(u)*4(æ)	-20
c. sy-i-ssæ		-1	-1		-110
d. su-ɯ-ssa	-1		-2		-220

Table 4.6: Transparency of [i]

3.3 Laterals alternating harmonically

The analysis thus far successfully accounts for the behaviour of vowels in the harmony system, but Votic harmony also involves laterals. As discussed in Section 2.3, there is a distinction between onset and coda laterals in Votic: the former always have an identical [back] specification to the vowel in the nucleus, while the latter take the [back] specification of the harmonic context (i.e. are back in back-harmonic words and front in front-harmonic words). In the present subsection, I deal only with coda laterals, which essentially participate in harmony the way that vowels do; discussion of onset laterals is reserved for the next subsection.

Recall that I assume (with formalization beyond the scope) that laterals are the only consonants for which [back] is specified in Votic, and that laterals must be specified for [back]. Further, the harmony constraint defined in (8) refers to moras with the feature [back], not specifically to vowels. Thus, any coda mora containing a lateral will be subject to the harmony constraint; codas with non-lateral consonants are irrelevant because they do not have the feature [back]. Unlike [i], laterals that are not onsets consistently participate in harmony in Votic, meaning that harmony involving them must be scaled sufficiently to overcome faithfulness. Votic laterals are always targets; they never control harmony themselves, and so do not behave as triggers, unlike what sometimes occurs with consonants in the better-known case of Turkish. I assume that laterals in Votic have the same target quality scaling as other vowels (i.e. 4), and will deal with trigger strength shortly. Any phonetic motivations behind the trigger and target behaviour of laterals is beyond the scope of the present paper; phonetic studies on laterals in the front/back dimension are necessary.

Table 4.7 illustrates that, with the assumptions so far, a coda lateral will agree with the harmonic context. This simply requires laterals being subject to the harmony constraint (i.e. all coda consonants are subject to it but laterals are the only ones with a [back] specification) and the existence of an IDENT-IO[back]$_V$ constraint but not one for consonants. Note that regardless of whether the two coda consonants have one or two moras, the lateral is the head of a mora, given the assumption that it is the most sonorous element in the mora that is the head. Additionally, note that the long vowel counts as two moras, and so in candidate (b), with the harmony violation, it is the second mora of the long vowel that is the trigger of the violation.

/na:tl/	*ɯ 200	IDENT-IO [back] $_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. na:tł					0
b. na:tl				-1*4(a)*4(l)	-16
c. næ:tl		-1	-1		-110

Table 4.7: Agreement with coda lateral

In the present framework, due to the way the harmony constraint is formulated and the way 'trigger' and 'target' are defined for the purposes of scaling, we also need to consider a complication posed by forms like [kəɫmaʒ] 'third'. In particular, if a form like this has disharmonic vowels underlyingly (e.g. /kəɫmæʒ/, which needs to be considered due to Richness of the Base and should come out harmonic), then the harmony should be controlled by the first vowel, meaning that the second vowel will need to change its [back] value in order to harmonize. However, for the purposes of scaling, the trigger is defined as the nearest possible segment with a different [back] value. As such, the relevant trigger/target pair to ensure harmony of the second vowel is the sequence ɫ-æ, scaled with trigger ɫ and target æ. As shown in the comparison between candidate (a) and candidate (b) in Table 4.8, the scaling of this sequence must be sufficient to outweigh vowel-specific [back] faithfulness. Otherwise, the lateral will be essentially an icy target (Jurgec 2011), undergoing harmony but not allowing harmony to propagate past it. Thus, even though laterals generally do not control harmony in Votic, they must in fact be scaled as good triggers; I assign them a value of 3 here, although it is not crucial that this value is lower than that of the harmonic vowels.[50]

/kəɫmæʒ/	*ɯ 200	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. kəɫmaʒ			-1		-10
b. kəɫmæʒ				-1*3(ɫ)*4(æ)	-12
c. kəlmæʒ				-1*4(ə)*4(l) + -1*4(ə)*4(æ)	-32
d. kelmæʒ		-1			-100

Table 4.8: Necessity of laterals being good triggers

Whether there exist reasons why laterals might be good triggers of harmony is beyond the scope of the present paper, as is the question of whether this complication can be avoided without adding unnecessary complexities to the framework. In particular, there might be some natural way of identifying the initial vowel as the trigger instead for the purposes of scaling, by revising the assumptions about which vowel counts as the trigger. Both questions are important for future research. We will see in Section 3.5, with the form [iɫma] 'without', that even allowing laterals to be good triggers does not affect the fact that they are only triggers in this context where they need to propagate harmony.

---

[50] Alternatively, we might have a pair of constraints involved, one ensuring harmony between vowels and the other between all moras. In that case, the vowel-specific harmony constraint would just need to have a higher weight, which is plausible since it involves segments that are more similar to each other.

3.4 Laterals in onsets

The remaining issue to deal with is the behaviour of laterals in onsets. Unlike for coda laterals, we cannot simply say that onset laterals participate in the harmony system; instead, they consistently agree with the following vowel. The difference can be clearly established by the fact that laterals immediately preceding [i] are consistently front, even in back harmonic words (e.g. [lidna:] 'town.ILL'; [dubli] 'healthy'). This means that onset laterals do not participate in harmony, either as triggers or targets, but instead take on the [back] value of the vowel in the nucleus of the syllable. This subsection deals with how to implement that generalization, while the following one will address [li] sequences specifically.

The statement of the harmony constraint used above already deals with half of the problem posed by onset laterals, in that it does not look at all [back] specifications: it only looks for moras that disagree in [back]. Since onsets are not moraic (as noted, I assume they attach directly to the syllable), the harmony constraint does not see them, and so the [back] specification of an onset lateral is not relevant. Here is where it is critical that the harmony constraint refer to moras rather than segments; otherwise, we would not expect to be able to make a distinction between moraic laterals, which participate in harmony, and onset laterals, which behave like the nucleus vowel of the syllable they are associated with. Thus, this particular formulation of the harmony constraint solves the problem of onset laterals not being subject to harmony; what remains is to explain why they agree with the following vowel.

The generalization that Votic laterals consistently agree with a following vowel, even if that results in disharmony between the lateral and other vowels within the word, could alternatively be expressed as an unviolated agreement relation between onset and nucleus. Hall

(2018) notes, but does not formalize, the fact that this pattern is a nucleus-to-onset relation, in explaining why laterals do not agree with following coronals, which his analysis would otherwise predict. It is worth noting that the pattern is not about following vowels in general, as laterals can be back before [i] if separated from it by a consonant (i.e. when moraic, as in [tałvi] ‘winter’). Thus, the specific point is that onsets take the [back] specification of the nucleus. This is true even immediately before [i] between back vowels, in CaliCa type words like [ma:lima] ‘paint’ or [əlimma] 'we were'.

There are several ways in which we could formalize this, such as assuming that onsets and nuclei share a single mora (e.g. Hyman 1985), and that onsets cannot have their own [back] specification independent from that of the nucleus that heads the mora. However, for simplicity, I will just assume an agreement constraint that applies between onsets and nuclei, as defined in (15). This constraint is never violated in Votic, so I assign it a high weight of 100. Note that this constraint is defined such that it cannot be violated by onset consonants that are not specified for [back] (i.e. it can only be violated when the onset is a lateral).

(15) AGREE[back]$_{ON}$: Assign a violation for every onset-nucleus pair that have disagreeing values of the feature [back].[51]

Adding this constraint ensures that an onset lateral will agree with the nucleus, as shown in Table 4.9. Due to the vowel-specific [back] faithfulness, it will again always be the vowel that controls this agreement, so that in any case of underlying disagreement between a lateral and an

---

[51] There is no particular reason for using AGREE here instead of a local version of a *AB constraint, like the harmony constraint. I adopt AGREE because Blumenfeld & Toivonen (2016) use it for the onset-nucleus lateral interactions; the only crucial thing is that the constraint be specified to onset-nucleus pairs. Note that I do not adopt AGREE as the harmony constraint because it does not have a natural way to refer to trigger and target.

immediately following vowel, the lateral will change to agree with the vowel, not vice versa. Worth noting is that the input /ełaː/ would also produce the same result, even though the winning candidate (a) would violate vowel faithfulness with that input, and the closest competitor (c) would not.

/ełæ:/	*ɯ 200	AGREE[back]$_{ON}$ 100	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. elæ:						0
b. ełæ:		-1				-110
c. eła:				-1	-1*4(e)*4(a:)	-26
d. əła:			-1	-2		-120

Table 4.9: Agreement between onset lateral and nucleus

3.5 Laterals with [i]

The final remaining issue is to ensure that this analysis adequately accounts for the behaviour of onset laterals before [i]. Recall that laterals agree with an immediately following [i] even in a back harmonic context. Table 4.10 illustrates how this analysis accounts for this fact. Candidate (a) wins because the trigger and target scaling factors for [i] are low, and because the harmony constraint cannot see the onset lateral. Candidate (b) is ruled out by violations of faithfulness to the [back] values of vowels, candidate (c) is ruled out by the highly weighted onset-nucleus agreement constraint, and candidate (d) is ruled out by the constraint against a high, back, unrounded vowel. Note in particular that there is no motivation for harmonizing the lateral to the back harmonic context, as done in candidate (c), because the lateral is not seen by the harmony constraint.

/əłimma/	*ɯ 200	AGREE[back]$_{ON}$ 100	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. əlimma					-1*4(ə)*1(i) + -1*1(i)*4(a)	-8
b. elimmæ			-1	-2		-120
c. əłimma		-1			-1*4(ə)*1(i) + -1*1(i)*4(a)	-108
d. əłɯmma	-1			-1		-110

Table 4.10: Agreement of onset lateral with [i] even in a back context

It is worth considering one further example, [lidna:] 'town.ILL', where harmonizing with the [li] sequence does not require changing the initial vowel. Specifically, we need to ensure that in cases where the [i] is the initial vowel, we still get disharmony. As shown in Table 4.11, the analysis predicts the correct form; this result requires the penalty for disharmony involving [i] to be smaller than that for violating IDENT-IO[back]$_{V}$. Note that an input /łɯdna:/ would also surface as [lidna:].

/łidna:/	*ɯ 200	AGREE[back]$_{ON}$ 100	IDENT-IO [back]$_{V1}$ 100	IDENT-IO [back]$_{V}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. lidna:					-1*1(i)*4(a)	-4
b. lidnæ:				-1		-10
c. łidna:		-1			-1*1(i)*4(a)	-104
d. łɯdna:	-1					-200

Table 4.11: Agreement of onset lateral with [i] even in a back context

A final case that is important to illustrate is [iłma] 'without', because it is a situation in which the labels 'trigger' and 'target' for the scaling factors are somewhat misleading. The [i] is relevant only because it means that the initial vowel is not harmonic. In this case, the ł-a sequence must be harmonic, and harmony is controlled by the vowel, not the lateral, because of preferential [back] faithfulness for vowels.[52] However, for the purposes of scaling, the lateral is the 'trigger' and the vowel is the 'target', because the terms are defined in relation to the context and focus of the harmony constraint, for which the context is always the first relevant segment and the focus the second. This does not cause any problems, as shown in Table 4.12, assuming a disharmonic input. The harmonic candidate (a) violates the harmony constraint twice, both with trigger [i]; notably, it does not violate onset-nucleus AGREE because the lateral is not an onset. Harmonizing

---

[52] In theory, this is a potential problem for suffix vowels. As stated so far, the analysis predicts that it is possible to get a lexical contrast in suffix vowels in roots with only neutral vowels, for instance between [i…-æ] and [i…-a]. It is unclear whether there is evidence for this in Votic, though notably in Finnish, some suffixes are back and others are front after monosyllabic neutral vowel roots, depending on the suffix vowel quality and whether the suffix is vowel or consonant initial (Kiparsky 1973). If there is no such contrast in Votic, the issue can be resolved in the same way as in Hungarian in Chapter 3. In Votic, unlike in Hungarian, there is an additional prediction that underlyingly back or front suffix vowels will trigger harmony regressively onto coda laterals in /CiLC/ roots. In other words, suffixes should behave the same way as the [a] in [iłma]. Again, it is unclear whether there is evidence for such a process, and again, it could be prevented by adding root faithfulness to the analysis.

the vowel is ruled out due to a violation of the vowel-specific [back] faithfulness constraint, as in candidate (b); the faithful disharmonic candidate is ruled out by a violation of the harmony constraint, as in candidate (c). It is beyond the scope of this paper to determine whether there is a different way of approaching ‘trigger’ and ‘target’ scaling without this issue, without creating additional unnecessary complexities in the framework.

/ilma/	*ɯ 200	AGREE[back]$_{\text{ON}}$ 100	IDENT-IO [back]$_{\text{V1}}$ 100	IDENT-IO [back]$_{\text{V}}$ 10	*[-αBack]∞ [αBack] 1	H
☞ a. iłma					-1*1(i)*4(a)+ -1*1(i)*4(ł)	-8
b. ilmæ				-1		-10
c. ilma					-1*3(l)*4(a)[53]	-12

Table 4.12: Coda laterals following [i] are back in a back word

4 Existing analyses

A number of other recent analyses have tried to solve the apparent paradox in Votic, primarily by claiming that the specification of [i] differs in some way from that of other vowels. B&T propose a distinction between ‘strong’ and ‘weak’ (effectively, contrastive and non-contrastive) specifications for [back]. This is similar in spirit to Nevins (2010) and others, not in formal implementation but in terms of different processes being parameterized with respect to what types of feature values are visible to them. Hall (2018) suggests that [i] is unspecified for [back],

---

[53] This candidate does not have an i…a violation because the violation for the [a] is assessed for only the closest [-back] segment.

but specified instead for [coronal]. In this section, I discuss the issues with these accounts, and argue that the approach presented in Section 3 improves upon them.

4.1 Blumenfeld & Toivonen (2016)

Blumenfeld & Toivonen (2016; B&T) suggest that Votic harmony refers only to ‘strong’ or contrastive features, while lateral alternations refer to all features. In this solution, [i] is specified for [back], but weakly and not contrastively, since it does not have a back counterpart. Thus, it will be exempt from harmony, but can participate in the lateral alternations. This approach therefore requires the phonology to be able to distinguish between features that are contrastive and those that are not, and different phonological processes can refer to different types of feature specifications.

This assumption adds a great deal of complexity to the phonology, and is difficult to implement within OT. Specifically, if the strength of a feature depends on whether the vowel is contrastive, then any constraints that reference strength must have a way of determining whether a vowel is contrastive. In some recent accounts of featural representation (Gradient Symbolic Representation or GSR; Smolensky & Goldrick 2016), this is possible, but in most varieties of OT, the way of determining inventory contrast is with the ranking of faithfulness versus markedness; for example, a way of establishing that [i] is not contrastive for [back] is the ranking of *ɯ >> IDENT-IO[back]. In B&T’s account, the way of knowing that [i] does not have to participate in harmony is by knowing that it is not strongly specified for [back], which comes from knowing that it is not contrastive for [back], which comes from this ranking. This fact is a substantial complication that is not consistent with most varieties of OT, which do not allow constraints to reference ranking relationships among other constraints in the grammar. While one does not have to implement such an analysis within OT, B&T's account is situated in OT, and therefore this fact is a weakness of that analysis.

Moreover, as Hall (2018) points out, in Russian loans, [i] does in fact have a [back] counterpart, [ɨ]. This fact suggests that the difference between [i] and other vowels cannot directly be encoded in whether a vowel has a harmonic counterpart in the inventory in the language as a whole. Thus, B&T would need to either say that strong and weak features are defined based on the native inventory, or else say that [i] is weakly specified for [back] despite having a potential harmonic counterpart.

More critically, B&T mention briefly that this type of analysis predicts that laterals could control harmony, but they discuss this pathology as though it is a bizarre typological problem inherent to every approach, rather than a problem in the Votic analysis itself. In particular, this approach predicts that [i] should be able to control harmony if preceded by a lateral, because the [i] controls the backness of the lateral, and the lateral participates fully in the harmony systems. Indeed, in B&T's tableaux, they do not consider candidates in which the vowels change [back] value; however, their constraints and ranking predict that a fully front candidate would win in any word with a [li] sequence. This result is shown in Table 4.13, which is the tableau in (26) in B&T, with only their winning candidate and the one in which the vowels are changed to front. Clearly, the latter candidate incorrectly wins. In this tableau, *ɯ is the same as in the present analysis, and Ident-IO-Bk is a standard input-output faithfulness constraint for [back]; Agree-LV ensures [back] agreement between a lateral and an immediately following vowel (the same as Agree[back]$_{\text{ON}}$ here). The Corr constraints enforce the formal relation of correspondence between relevant segments, notated with subscripts in the tableaux; matching subscripts indicate that segments are in correspondence. Specifically, Corr-V$_{\text{S}}$V$_{\text{S}}$ enforces this relation between ‘strong’ vowels (those that are contrastive for [back]), while CORR-L-V$_S$ enforces it between laterals and strong vowels. The constraint IDENT-BK requires [back] agreement between segments in this formal relation of correspondence; it is not a faithfulness constraint.

/əɫimmɑ/	*ɯ	AGREE-LV	CORR-V$_S$V$_S$	CORR-L-V$_S$	IDENT-BK	IDENT-IO-BK
ə$_i$l$_i$immɑ$_i$					*!	*
e$_i$l$_i$immæ$_i$						***

Table 4.13: Pathology predicted by the B&T account

The problem here is that there is no differentiation between laterals and vowels. The combined pressures for the lateral to agree with the following vowel, all back-specifiable segments (including laterals) to correspond, and all corresponding segments to agree in [back] means that the vowel with which the lateral agrees will always control harmony, even if it does not enter into the correspondence relation. This result is in direct contrast with the Votic pattern that needs to be captured.

While we might try to circumvent this problem by including [back] faithfulness for only vowels (not laterals), ranked highly enough to make vowels control harmony, such an approach applied within the existing B&T analysis would incorrectly predict that vowels do not harmonize, which would obviously be a problem. We cannot even derive the pattern by requiring the initial vowel to be faithful, as there are instances of back words with initial [li] sequences, such as [lidnaː] ‘town.ILL’; the same is true of final vowels, in cases like [dubli] ‘healthy’. A tableau with initial vowel faithfulness is shown in Table 4.14. Specifically, it will always be necessary to have the [l] stand in correspondence with the later vowels, which in turn requires making those vowels front in order to satisfy agreement; this correspondence configuration will always be preferred because of the ranking CORR-L-VS >> IDENT-IO-BK.

/lidna:/	IDENT-IO-BK$_{\sigma1}$	*ɯ	AGREE-LV	CORR-VSVS	CORR-L-VS	IDENT-BK	IDENT-IO-BK
l$_i$idna:$_i$						*!	
l$_i$idnæ:$_i$							*

Table 4.14: Pathology still present with initial vowel faithfulness

Given that disharmony between vowels is not permitted except when one of the vowels is [i], and that the reason B&T provide for the special status of [i] is that it does not need to correspond, and that laterals are required to correspond, words of this type should not exist if laterals and vowels have equal status. As such, the B&T account is not actually able to capture the most interesting and problematic pattern in Votic.

4.2 Hall (2018)

Hall (2018) proposes a different solution, in which [i] is unspecified for [back], but instead specified for [coronal]. In this approach, Votic lateral alternations are triggered by two different features, [back] and [coronal], which complicates the phonological statement of this process. Hall does not implement the analysis in detail, but there are several issues to consider. Critically, there is no reason to have lateral alternations refer to two separate features, besides a desire to leave [i] unspecified for [back] for the purposes of maintaining a specific analysis of neutrality. Given that two features are relevant, we would expect these lateral alternations to be two separate processes that could perhaps behave distinctly, but there is no evidence of that in Votic or in other related languages. As Hall notes, not all [coronal] segments immediately following a lateral trigger the front alternant; the relation is exclusively between vowels and laterals, and coronal consonants do not affect the process. While this fact could be implemented in multiple ways, this problem is avoided entirely if the relevant feature is always [back], with the caveat that it needs to be addressed why only laterals alternate and are specified with [back]. Combining lateral alternations into a single phonological process, by not requiring a solution to disharmony involving [i] even though it is [-back], is preferable to complicating the statement of the process.

4.3 General discussion of previous analyses

It is worth noting that, in addition to adding complexity to the phonology, neither of these previous solutions for Votic can be extended to complications in the neutrality/contrast relationship in other languages. In Mayak tongue root harmony (see Chapter 9), low vowels are demonstrably contrastive for [ATR] but nonetheless neutral; there is no way to account for this pattern under the idea that neutral vowels are those that do not contrastively have the feature value, or that they have a different, related feature value. Turning to Hungarian (see Chapter 3), the gradience of neutrality among front unrounded vowels shows that the solution to what is neutral cannot be as categorical as these previous analyses of Votic presume. In particular, [back] is not contrastive on any of the front unrounded vowels in Hungarian, yet they vary in neutrality. Given that this variability occurs within the harmony process, it is clear that the relevant feature on these vowels must be [back]. Ideally, any solution to why [i] is neutral in Votic, despite apparently requiring a [-back] feature for lateral alternations, should be extendable to why Hungarian front unrounded vowels are sometimes clearly front, sometimes neutral to harmony.

5 Discussion and conclusion

This chapter has shown a novel analysis for Votic harmony and lateral alternations, based on the idea that front/back harmony operates at the mora level in Votic and on the target-based harmony approach developed for Hungarian in Chapter 3. I argued that previous analyses of Votic harmony are unnecessarily complex and make incorrect predictions. In contrast, the approach presented here is correctly able to predict that [i] is neutral to front/back harmony despite triggering the front version of an immediately preceding lateral, and that laterals not immediately preceding a vowel should behave like targets of harmony, while laterals immediately preceding a vowel behave like the following vowel.

This analysis makes several novel proposals, which make predictions about harmony systems more generally. The idea that harmony can operate on the mora level predicts that if we see consonants behaving asymmetrically in a vowel harmony pattern, it should be in the direction of Votic, not the reverse (i.e. where onsets demonstrably participate in harmony, not just as an onset-nucleus relation, but codas do not participate). Specifically, we predict the possibility that all consonants that participate in vowel harmony in a language do so in the same way, or that there is a distinction like in Votic in which onsets behave like the following vowel but codas harmonize, or a case in which onsets are transparent but codas participate, but we do not predict a case in which codas behave like the preceding vowel but onsets harmonize. This idea is potentially problematic, because codas do often pattern as more tied to vowels than onsets, hence the notion of rime. The critical idea here is that, at least in Votic, consonant-vowel harmonic interactions seem to require a lesser status for onsets, in that onset laterals do not have an independent function in the harmony system, while coda laterals do. This fact is consistent with the idea that codas are moraic, while onsets are not, such that onsets do actually have a lesser status. However, understanding why onsets take on the [back] value of the nucleus (as opposed to just stipulating the relevant agreement constraint) is a difficult problem, at least without hypothesizing an onset-nucleus constituent. I leave this issue for future research; one possibility is that harmony interactions are biased this way, but that in other places in the phonology, such as in prosody, the nucleus and codas are more tightly aligned. It may also be a language-specific pattern. It is worth noting that a potential problem raised by this analysis is that, as far as we know, there is no onset/coda asymmetry in other languages in which consonants participate in vowel harmony, meaning that it is unclear whether other languages make use of the system of moraic harmony. However, given the paucity of research on many vowel harmony systems, and in particular on their phonetic properties and the participation of consonants, this analysis is worth further exploring as making predictions for other cases. In particular, we might expect that if we looked more closely at the phonetics of consonants in vowel harmony systems, we might occasionally see the onset/coda asymmetries predicted here.

Moreover, the way this analysis deals with the fact that consonants are targets of vowel harmony but not triggers is novel and worth exploring further. In particular, I have argued that because [back] is primarily a vocalic feature, there should be independent [back] faithfulness to vowels; the existence of such a constraint, and not a parallel one for consonants, correctly predicts that when a consonant and a vowel are required to harmonize for [back], it should be the vowel that controls the harmony pattern.[54] As B&T point out in briefly mentioning this pathology in their analysis, consonants that participate in vowel harmony patterns do not trigger harmony. However, in Turkish, there are exceptional cases in which palatalization on consonants

---

[54] It is worth noting that the theory here does not say anything about the feature [back] and how it characterizes consonants, but that this would be an important future step for understanding how to constraint [back] patterns involving consonants.

is contrastive and can interrupt the front/back harmony pattern (e.g. Clements & Sezer 1982). Understanding this Turkish pattern in the context of the present proposal for Votic is worth exploring in the future.

Chapter 5: A segment-specific metric for quantifying participation in harmony

1 Introduction

There is a long history of phonological research into quantifying and analyzing lexically gradient co-occurrence patterns between segments. This research has most frequently dealt with dissimilatory patterns (i.e. OCP effects) among consonants (e.g. Pierrehumbert 1993, Frisch et al. 2004), though there has also been similar work on assimilatory patterns (e.g. Rose & King 2007, Brown 2008, Arsenault 2012). However, most of this work has focused primarily on consonant co-occurrences. With respect to lexically gradient vowel co-occurrence patterns, specifically vowel harmony, the existing work focuses more on modelling variability in heteromorphemic contexts (e.g. Hayes & Londe 2006 on Hungarian), gradient carry-over of categorical harmony into other domains like compounding (e.g. Martin 2007), and quantifying overall trends towards harmony in the vowel system as a whole (e.g. Harrison et al. 2004).

Despite all of this research on segmental co-occurrence patterns, there is little research into gradient morpheme-internal vowel harmony patterns. Moreover, there remains no systematic way of calculating the gradient degree to which a segment participates in a harmony system. The standard Observed/Expected (O/E) values give a measure of the behaviour of each individual pair of segments, but a comprehensive picture of how harmonic a segment is, across its co-occurrences with all segments, is still lacking. Such a measure of overall segment-specific gradient harmony participation is a useful concept because it could illuminate gradient language-internal and cross-linguistic patterns in harmony participation that are not apparent from more categorical descriptions or entirely clear from O/E values.

In this chapter, I adopt the statistical concept of *relative risk* as a measure of participation in harmony. I compute both O/E values and the relative risk measure for vowels in corpora of three languages with front/back harmony, to show that relative risk can capture the intuitive notion of how much a vowel participates in harmony. I then consider the implications of the results, given what is known about categorical trends of participation in front/back harmony systems in other languages. This measure does in fact provide new insights into harmony participation in these languages.

2 Participation in harmony

The question of which vowels participate in harmony is a long-standing one in the literature. For example, as discussed in Chapters 2 and 3, it is well-known that non-low, front, unrounded vowels are neutral to front/back harmony in Finnish and Hungarian, and this pattern extends to a variety of other languages (e.g. Kiparsky & Pajusalu 2003). Specifically, [i] is a common neutral vowel, and neutrality of lower front unrounded vowels like [e] typologically implies neutrality of [i] (e.g. Anderson 1980, Benus 2005, Finley 2015). There are clear typological trends in which vowels tend to be neutral to a given type of harmony, and moreover these patterns often have a phonetic basis (e.g. Archangeli & Pulleyblank 1994; Chapter 3). For example, Beddor et al. (2001) argue that the tendency of [i] and [e] to be neutral in front/back harmony systems results from smaller effects of vowel-to-vowel coarticulation on them. Similarly, Benus & Gafos (2007) argue that neutrality of [i] to front/back harmony can be explained by its acoustic/perceptual stability under coarticulatory displacement along the front-back dimension; lower vowels have less of this stability. This property is used to explain the *height effect* of neutrality in Hungarian, in which, for example, [i] is more neutral than [e:], even though [e:] is still frequently neutral (e.g. Hayes & Londe 2006). As such, participation and neutrality are not clear-cut categories in Hungarian. This type of intermediate neutrality suggests that participation is more complex than categorical descriptions would suggest.

While non-participation in harmony is commonly linked to inventory structure, with neutral vowels being the ones that lack harmonic counterparts, this cannot be the entire story; at least some cross-linguistic patterns of participation are independent of contrast. For example, in Mayak (Nilotic; Andersen 1999), low vowels are generally neutral to ATR harmony, but contrastive for ATR; this corresponds to a general cross-linguistic trend for low vowels to be neutral to ATR harmony (Archangeli & Pulleyblank 1994). Moreover, vowels that lack counterparts may nonetheless participate in harmony by *re-pairing* with vowels that differ in an additional feature value (Baković 2000). For example, Maasai (Nilotic) has no [+ATR] low vowel, but its [-ATR] low vowel [a] alternates with the mid vowel [o] for tongue root harmony in certain contexts (e.g. Baković 2000). Such patterns provide clear evidence that whatever drives the typological tendencies in participation goes beyond mere contrast and inventory structure.

3 Gradient harmony and quantification

As previously mentioned, it is well-known in the literature that harmonic co-occurrence restrictions can be gradient; this occurs for instance in laryngeal consonant harmony restrictions in the Ethio-Semitic languages Chaha and Amharic (e.g. Rose & King 2007). The typical measure used to quantify gradient co-occurrence restrictions is observed over expected (O/E) values (e.g. Pierrehumbert 1993; Frisch et al. 2004; Coetzee & Pater 2008).[55] O/E indicates, for

---

[55] Note that other work, in particular Graff & Jaeger (2009) and Wilson & Obdeyn (2009), has pointed out problems with the O/E metric and argued for the rejection of its use.

each pair of segments, how over- or under-represented the co-occurrence of those segments is. This means that O/E values, as traditionally used, can only express the participation of a specific pair in co-occurrence restrictions. For example, O/E can say whether [i] followed by [a] is over- or under-represented, but does not say whether [i] in general tends to follow harmonic co-occurrence restrictions, as compared to other vowels.[56]

Beyond O/E, harmony has also been quantified overall for a language (Harrison et al. 2004). This value expresses the extent to which the percentage of harmonic co-occurrences in a corpus exceeds chance, given the relative proportions of each harmonic category in the corpus. This measure requires stipulating which vowels are neutral, and is an overall measure, not segment-specific.

Missing from the gradient quantification of harmony is a segment-specific measure of participation. Such a measure could provide interesting new insights because of the cross-linguistic patterns of participation described in Section 2. Given that these typological trends can exist independently of inventory contrast, and given that harmony can be gradient, we might expect the categorical trends in participation/neutrality to also manifest in a gradient way. For example, given that [i] is a common neutral vowel in front/back harmony, we might predict that [i] should participate gradiently less in front/back harmony even in languages in which it is not categorically a neutral vowel. However, there is no existing way of measuring this. In particular, it is insufficient to measure participation with the percentage of the segment's co-occurrences that are harmonic, because that fails to take into account the relative frequency of the two

---

[56] It is worth noting, however, that it would be possible to get a general measure of participation using O/E, by constructing a table that groups together the vowels not under consideration by category. For example, for [i] in a front/back system, such a table would have rows/columns for [i], non-[i] front vowels, and back vowels. Performing O/E on such a table could give two general measures for [i], one of how over- or under-represented it is with front vowels, and one of how over- or under-represented it is with back vowels. Such a measure is slightly different from the one pursued in this chapter, and is worth considering further in future work. One reason for using the measure in this chapter, over the O/E method, is because it combines both harmonic and disharmonic behaviour into a single measure, rather than having two separate O/E values.

harmonic categories; if one category is more frequent, then the likelihood of segments in that category being in harmonic co-occurrences is increased.

4 Relative risk measure

I propose to adopt relative risk as a way of computing segment-specific harmony participation. Relative risk (RR) is a standard statistical measure for count data, used to determine the relative probability of an 'event' in one condition versus another; it answers whether the 'event' is more likely in one condition, and if so by how much (Agresti 2013). It has previously been used in linguistics as a measure for speech errors (Tupper & Alderete in prep). The general formula for RR is as in (1), where $P(E|C_x)$ is the probability of event E given condition C_x.

(1) $$RR = P(E|C_1)/P(E|C_2)$$

Since this measure is a ratio of probabilities in different conditions, it corrects for the relative frequency of C_1 and C_2.

Interpreting the 'event' as the segment of interest, and the conditions as harmonic versus non-harmonic contexts, RR provides a measure of how much more frequent a segment is in harmonic contexts, and therefore how harmonic the segment is. In cases where all relevant counts are known, as in a corpus study of vowel co-occurrence patterns, the formula in (1) can be simplified to one that references only counts, rather than probabilities. The RR for vowel V from category C_1 is given in (2).

(2) $$RR(V) = \frac{(\#\text{ V in context } C_1) \times (\#\text{ anything in context } C_2)}{(\#\text{ V in context } C_2) \times (\#\text{ anything in context } C_1)}$$

This formula can also be re-written as in (3), as the ratio of ratios. These two formulas are equivalent.

(3) $$RR(V) = \frac{\dfrac{(\#\ V \text{ in context } C_1)}{(\#\ V \text{ in context } C_2)}}{\dfrac{(\#\text{ anything in context } C_1)}{(\#\text{ anything in context } C_2)}}$$

Note that if the vowel V belongs to category C_2, then the position of C_1 and C_2 in the formula will be reversed. Thus, this measure effectively multiplies the ratio of V in harmonic to disharmonic contexts (i.e. (# V in context C_1) / (# V in context C_2)) by a constant for each category: if V is in category C_1, the constant is as in (4); if V is in category C_2, then the constant is the inverse of this, as in (5).

(4) $$\text{constant for V in category } C_1 = \frac{(\#\text{ anything in context } C_2)}{(\#\text{ anything in context } C_1)}$$

(5) $$\text{constant for V in category } C_2 = \frac{(\#\text{ anything in context } C_1)}{(\#\text{ anything in context } C_2)}$$

If V is in category C_1, then the context C_1 also includes V. Thus, in order to avoid double-counting occurrences, it is simpler to treat this formula as directional, computed twice for each vowel: once for all co-occurrences involving V as the first of the two vowels ('V1'), and once for co-occurrences where V is the second of the two vowels ('V2').[57] In the former case, the 'context' refers to the second vowel, while in the latter case, it refers to the first vowel. This approach has the additional advantage of determining whether there are directional asymmetries

---

[57] Note that 'V1' and 'V2' are labels for relative position, and they do not necessarily reflect absolute position within the word. For example, if considering the co-occurrence between the third and fourth vowels within a given word, 'V1' would be the third vowel, as it is the first of the vowels under consideration.

in harmony participation, which we might expect to be the case, for instance, if harmony applies strictly in one direction triggered by only a single value of the feature.

One way of thinking of RR(V), as shown in (3), is that it represents the ratio of occurrences of V in harmonic contexts to disharmonic contexts, relative to the ratio of occurrences of V's harmonic versus disharmonic contexts in general. As such, the higher RR(V) is, the more harmonic V is; a high RR(V) indicates that V occurs more in harmonic than disharmonic contexts, relative to what we would expect given the relative frequency of those contexts. For the purposes of this study, I will use RR(V) as a relative measure, comparing it across vowels to establish what counts as high and low values.

To illustrate how RR(V) works, I will briefly consider several toy examples. The relevant factors to consider are whether V tends to occur in harmonic contexts or is relatively evenly distributed between contexts, and whether or not the distribution of contexts (i.e. harmonic categories) is skewed, and if so in which direction. Table 5.1 gives toy counts and RR(V) values for each of the six resulting contexts to consider.

	Symmetric category distribution	Skewed category distribution (to category of V)	Skewed category distribution (away from category of V)
V generally harmonic	# V in context C_1 = 50 # V in context C_2 = 10 # any in context C_1 = 250 # any in context C_2 = 250 ***RR(V) = 5***	# V in context C_1 = 50 # V in context C_2 = 10 # any in context C_1 = 400 # any in context C_2 = 100 ***RR(V) = 1.25***	# V in context C_1 = 50 # V in context C_2 = 10 # any in context C_1 = 100 # any in context C_2 = 400 ***RR(V) = 20***
V relatively evenly distributed between contexts	# V in context C_1 = 30 # V in context C_2 = 30 # any in context C_1 = 250 # any in context C_2 = 250 ***RR(V) = 1***	# V in context C_1 = 30 # V in context C_2 = 30 # any in context C_1 = 400 # any in context C_2 = 100 ***RR(V) = 0.25***	# V in context C_1 = 30 # V in context C_2 = 30 # any in context C_1 = 100 # any in context C_2 = 400 ***RR(V) = 4***

Table 5.1: Values for toy examples, with vowel V in category C1

As is evident, both the distribution of V and the distribution of categories affects the RR(V) value. If both V and the contexts are evenly distributed, then RR(V) = 1. This would also be true if both V and the contexts are skewed in exactly the same way (e.g. if V occurs 5 times more often in its harmonic context, and its harmonic context occurs 5 times more often than the other context). If V co-occurs with other vowels in its harmonic category less frequently than we would expect based on the relative frequency of the categories, as is the case if V is evenly distributed but the category distribution is skewed towards the category of V, then RR(V) < 1. If V occurs in its harmonic category more often than we would expect based on the distribution of categories, then RR(V) > 1, and the degree to which RR(V) exceeds 1 depends on how much more harmonic V is than we would expect. For example, if V is generally harmonic, but the contexts are skewed such that V's category occurs more frequently, then RR(V) will be lower than if the counts for V are the same but the categories are distributed symmetrically. In turn, this latter RR(V) will be lower than with the same counts for V, but contexts skewed such that the category to which V does not belong occurs more often. This property can be seen by comparing RR(V) values across a row in Table 5.1. As such, RR(V) takes into account not only how harmonic V is based on its raw counts, but also how harmonic we would expect it to be given the overall distribution of the categories; the higher the RR(V) value, the more harmonic V is, compared to what we would expect given the category distribution. Any asymmetry in the category distribution will translate accordingly into an asymmetry in RR(V) values.

It is worth noting that within a category, vowels can be compared simply by their own ratio (i.e. #V in context C_1 / #V in context C_2, for each V in category C_1), but this is not the case across categories. For example, in the toy example in Table 5.1, say that there is one vowel, V_a, that is generally harmonic, and a second vowel, V_b, that is evenly distributed between the categories. If V_a and V_b both belong to category C_1, then the ratio between their RR(V) values will be constant regardless of the distribution of their categories; in Table 5.1, this ratio is 5, as can be seen by comparing values within a column. The reason for this consistency is that the 'correction factor' for category frequency, namely the ratio of category distributions, is the same in both cases if both vowels are in the same category (see (4-5)). In contrast, if V_a belongs to category C_1 and V_b belongs to category C_2, then it is crucial to use RR(V), rather than a pure ratio, as the 'correction factor' is different. This fact becomes important when the distribution is skewed towards one of the categories. For example, if the numbers are as in Table 5.1, V_a is generally harmonic and in category C_1, V_b is evenly distributed and in category C_2, and the distribution is skewed towards category C_1, then RR(V_a) should be 1.25 and RR(V_b) should be 4. These numbers indicate that V_b is more harmonic than V_a, even though V_b is evenly distributed across categories, because the categories themselves are not evenly distributed. Not correcting for the distribution of categories in such an example would make it seem like V_a is the more harmonic vowel, which is false in this case. Thus, RR(V) allows for comparison across categories, whereas a simple ratio would not.

5 Case studies

For RR(V) to be a valid and useful measure, it needs to correspond to other notions of how much a vowel participates in harmony. While O/E does not provide a single, unified measure of how harmonic a vowel is, it does give relevant information. Specifically, using O/E values, we would consider a vowel highly harmonic if it is highly under-represented in disharmonic contexts (O/E close to 0) and over-represented in harmonic contexts (O/E greater than 1). Vowels with such patterning should have the highest RR(V) values. In contrast, we would consider a vowel not particularly harmonic if its representation in harmonic and disharmonic contexts is either approximately what we would expect from random distribution (O/E close to 1), or if it is under-represented in harmonic contexts and over-represented in disharmonic ones. Vowels of this type should have the lowest RR(V) values.

While the toy example in Table 5.1 already shows that RR(V) does correspond to the notion of how harmonic a vowel is, relative to category frequency, it is ideal to illustrate it with actual language data. Thus, to test this measure, I examine corpora of three languages with front/back harmony: Chuvash, Tatar, and Mari (Luutonen et al. 2007; Luutonen et al. 2016). All of these languages are spoken in a similar area in Russian; the former two are Turkic languages, while Mari is Uralic. Using the software Phonological CorpusTools (Hall et al. 2015), counts of all pairs of vowels adjacent on a vowel tier were obtained for all of the corpora; this resulted in 117833 co-occurrences for Mari, 59066 for Chuvash, and 106378 for Tatar. These were used to compute O/E values for all vowel pairs, as well as RR(V) for each vowel. This work is exploratory, and does not make any concrete claims about harmony in any of these languages. Notably, the corpora are dictionary lists that include derived forms and are written in orthography, not transcription. The entirety of the corpora were used in the computational analysis to get a picture of broad harmonic behaviour throughout the lexicon. However, there may be complications due to the inclusion of derived forms and loans, as well as in the orthography-to-phonology mapping, that could skew or distort the counts, and further work should examine these issues.

As mentioned, the intuitive notion of how much a vowel participates in harmony can be captured by looking at O/E values; essentially, if a vowel occurs with very low O/E values with vowels of the opposite harmonic class, then it is highly harmonic, and vice versa. To check whether RR(V) is a valid measure of participation, we can compare it to what we observe through O/E values across a vowel's co-occurrences with all other vowels. Vowels that are highly harmonic given their O/E values should have high RR(V), and less harmonic vowels should have lower RR(V).

If RR(V) is a valid measure, then we can next examine the kinds of patterns that emerge in participation. One hypothesis is that relative RR(V) values should correspond to categorical cross-linguistic participation trends. In particular, while [i] is not consistently described as neutral in the languages examined here, it is often categorically neutral in front/back harmony in related languages, like Finnish and Hungarian, in the Uralic family with Mari, and Uyghur in the Turkic family with Chuvash and Tatar (e.g. Ringen & Heinämäki. 1999 on Finnish; Siptár & Törkenczy 2000 on Hungarian; Lindblad 1990 on Uyghur).[58] Moreover, as described in Section 2, [i] is considered a good neutral vowel in front/back harmony. We might therefore predict that the RR(V) value of [i] should be low (i.e. among the lowest of the values) across these languages, and across both environments (V1 and V2 positions). Specifically, we might expect RR(i) to be lower than the RR(V) value of front vowels that are lower and/or rounded, as these are known factors for reducing the likelihood of a vowel being neutral to front/back harmony (e.g. Benus and Gafos 2007).

5.1 Chuvash

Table 5.2 gives the O/E values for all combinations of Chuvash vowels. Values above 1 indicate over-representation, while values below 1 indicate under-representation; the closer a value is to 0, the closer the co-occurrence is to categorically absent. The darker lines in the table separate the back and front categories, while the grey boxes indicate instances where harmonic co-occurrences are under-represented or disharmonic ones are over-represented. It is particularly

---

[58] Whether [i] is neutral in the languages examined here is not entirely clear. The descriptions of vowel harmony in Chuvash and Tatar are limited, but do not mention neutrality; moreover, Chuvash [i] has a back counterpart in some contexts, and there is a diphthong often taken to be a back counterpart of [i] in Tatar. (See e.g. Krueger 1961 and Clark 2015 on Chuvash; Berta 2015 on Tatar.) These facts suggest that [i] is likely not categorically neutral in Chuvash and Tatar. The case of Mari is less clear; Kangasmaa-Minn (1998) says that [i] and [e] are neutral vowels in Mari harmony, while Vaysman (2009) does not describe any neutrality in the language.

worth noting that the only O/E values of disharmonic co-occurrences (i.e. the upper right and lower left quadrants) that approach or exceed 1 involve [i], and that 5 of the 8 disharmonic co-occurrences involving [i] have high O/E values, above 0.8. These numbers are suggestive of [i] co-occurring regularly with back vowels.

V1 ↓ V2 →	a	ŏ	ɯ	u	e	ø̆	i	y
a	1.30	1.14	1.36	1.54	0.32	0.08	1.49	0.06
ŏ	1.41	2.03	0.94	0.84	0.02	0.11	0.27	0.09
ɯ	1.26	2.22	0.34	0.89	0.09	0.02	0.35	0.10
u	1.23	1.62	1.56	1.20	0.31	0.14	0.92	0.15
e	0.48	0.15	0.52	0.61	2.26	1.64	1.55	3.98
ø̆	0.04	0.04	0.62	0.26	2.70	4.43	0.24	3.18
i	0.88	0.24	0.86	0.91	1.79	1.49	1.21	0.45
y	0.03	0.02	0.00	0.17	3.03	4.24	0.15	1.95

Table 5.2: Chuvash O/E values

Table 5.3 gives the RR(V) values for all Chuvash values, in both V1 and V2 context, while Figure 5.1 shows the same information in graphical representation. Noteworthy is the fact that [a,u,i] have very low RR(V) in both contexts, while in general the more marked vowels, like [y], have higher RR(V). The one exception is [ɯ] in V2 context; however, the distribution of this vowel is restricted in non-initial position in Chuvash (Krueger 1961), suggesting that its behaviour here may be more reflective of a very small number of examples. Besides [ɯ], the vowels [a,u,i] have the lowest RR(V) in both contexts, and in particular, [i] has lower RR(V) than front vowels that are lower and/or rounded.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	2.09	2.77
ŏ	13.32	10.64
ɯ	10.31	1.94
u	3.06	2.09
e	5.13	10.26
ø̆	40.94	25.28
i	2.30	1.03
y	66.85	28.81

Table 5.3: Chuvash RR(V) values

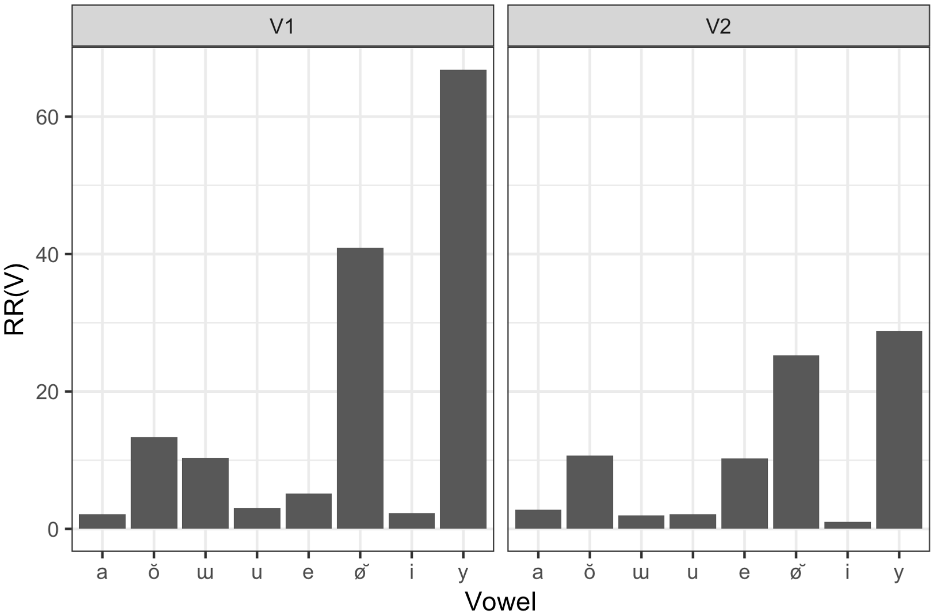

Figure 5.1: Chuvash RR(V)s

5.2 Tatar

Table 5.4 shows the O/E values for combinations of Tatar vowels. As in Chuvash, Tatar [i] occurs regularly with back vowels, and often not as regularly as expected with front vowels. In Tatar, [o] demonstrates a similar property, occurring fairly regularly with front vowels but under-represented in many back vowel contexts. For disharmonic co-occurrences (upper-right and lower-left quadrants), the co-occurrences that approach or exceed 1 are [i] with [a, o] in either order, and [e] with [o] in either order. However, [i] and [e] differ in terms of their co-occurrences with front vowels, in that the O/E of co-occurrences of [e] with front vowels consistently exceeds 1, while 5 of the 9 cases of [i] with front vowels have O/E less than 1. In terms of [o], 5 of the 7 co-occurrences of [o] with back vowels have O/E below 1, and all back-back co-occurrences with O/E less than 1 involve [o] as one of the two vowels.

V1 ↓ V2 →	a	o	ɤ	u	æ	ø	e	y	i
a	1.24	0.90	1.48	2.01	0.21	0.49	0.31	0.03	1.29
o	1.32	2.60	0.94	0.42	0.04	0.46	0.82	0.03	1.80
ɤ	1.27	0.18	2.50	2.27	0.09	1.22	0.14	0.04	0.08
u	1.94	0.44	1.76	0.61	0.08	1.23	0.32	0.03	0.84
æ	0.23	0.03	0.08	0.07	2.68	0.86	1.80	4.75	0.63
ø	0.17	0.08	0.03	0.04	3.40	3.98	2.97	0.71	0.30
e	0.61	1.34	0.27	0.21	1.52	2.17	1.88	2.12	1.09
y	0.12	0.08	0.06	0.06	4.23	1.48	2.49	0.59	0.31
i	1.04	2.05	0.24	0.27	1.50	0.39	1.43	0.57	1.23

Table 5.4: Tatar O/E values

Table 5.5 and Figure 5.2 illustrate RR(V) values for Tatar vowels. Notably, [o] and [i] show the lowest values in both contexts, and [a, u, e] are also quite low. As in Chuvash, marked vowels like [y] and [ɤ] show higher RR(V) values, though [ø] in V2 context has one of the lowest values. Worth noting again is that [i] has a lower RR(V) value than front vowels that are lower and/or rounded.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	2.91	2.35
o	1.60	0.97
ɤ	15.89	9.04
u	4.33	9.93
æ	15.52	14.81
ø	22.91	1.89
e	3.15	5.09
y	25.94	74.81
i	1.65	0.84

Table 5.5: Tatar RR(V) values

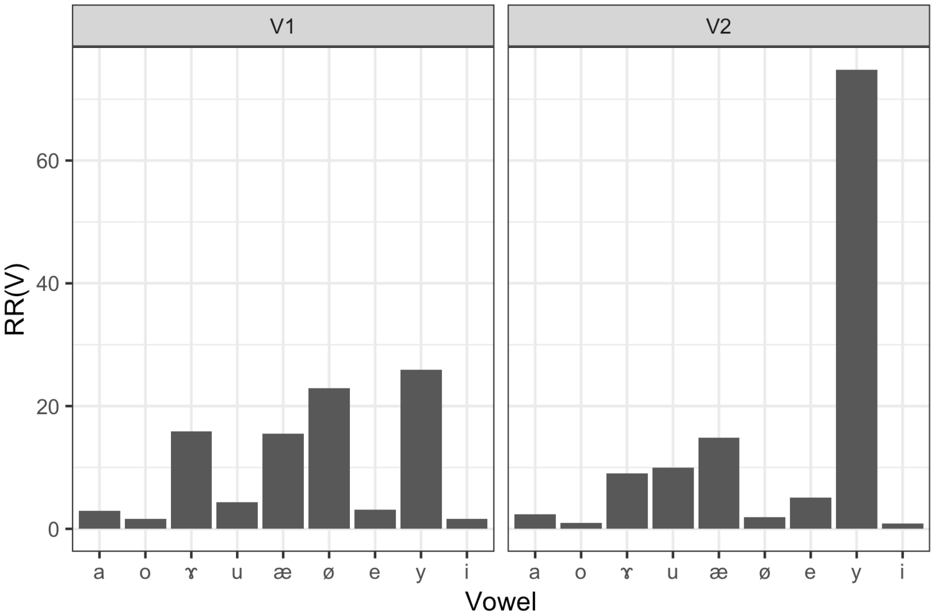

Figure 5.2: Tatar RR(V)s

5.3 Mari

Table 5.6 illustrates O/E values for combinations of Mari vowels. Again, [i] co-occurs regularly with back vowels; the same is true of [e], with most of its O/E values with back vowels nearly at 1. In fact, O/E values for [e] are consistently close to 1 across back vowels, while [i] has lower O/E values with [ɤ]. However, as in Tatar, the situation in front-front co-occurrences is different for [i] versus [e], in that [i] shows a far greater tendency than [e] to not co-occur with other front vowels (other than itself and [e]). Interestingly, in Mari, [æ] and [ə] in V1 position co-occur with each other at much higher than expected rates, while rarely occurring with any other vowels of either category. In V2 position, these vowels again occur most often with only each other. It is notable that in at least some dialects of Mari, [æ] is described as occurring only as a result of harmony (Vaysman 2009).

V1 ↓ V2 →	a	o	u	ɤ[59]	æ	ø	y	e	i	ə
a	1.32	0.71	1.17	1.12	0.06	0.16	0.93	0.85	1.26	0.05
o	0.96	2.05	0.96	1.09	0.04	0.20	0.64	1.05	1.43	0.02
u	1.01	1.50	1.36	1.38	0.04	0.37	0.77	0.84	0.99	0.04
ɤ	1.41	0.60	0.63	1.37	0.01	1.20	0.88	0.92	0.26	0.02
æ	0.06	0.10	0.24	0.08	8.46	0.07	0.84	0.26	0.27	6.14
ø	0.72	0.27	0.60	1.16	1.64	13.90	1.72	1.01	0.25	1.81
y	0.68	0.37	0.66	1.44	1.15	9.32	4.37	1.11	0.13	1.52
e	0.71	1.09	1.19	0.62	1.49	0.39	1.34	1.60	1.78	1.29
i	0.67	1.84	1.93	0.46	1.55	0.44	0.58	1.40	1.69	1.22
ə	0.02	0.11	0.18	0.02	6.50	0.43	0.60	0.60	0.12	8.63

Table 5.6: Mari O/E values

Table 5.7 and Figure 5.3 show RR(V) values for Mari vowels. What is most noteworthy in Mari is that two vowels, [æ] and [ə], have substantially higher RR(V) values than the remaining

---

[59] This vowel is often transcribed as [ə̂] in Uralic literature (e.g. Kangasmaa-Minn 1998); I have written it as [ɤ] to emphasize its status as a back mid unrounded vowel, but it should be noted that it is reduced.

vowels, for both V1 and V2 contexts. Among the remaining vowels, the differences as small, though [i] again has one of the lowest RR(V) values in both contexts.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	1.86	2.36
o	1.55	1.14
u	2.12	0.94
ɤ	2.88	2.23
æ	37.75	79.87
ø	1.74	3.82
y	1.39	1.59
e	2.05	1.30
i	1.78	1.16
ə	74.11	90.27

Table 5.7: Mari RR(V) values

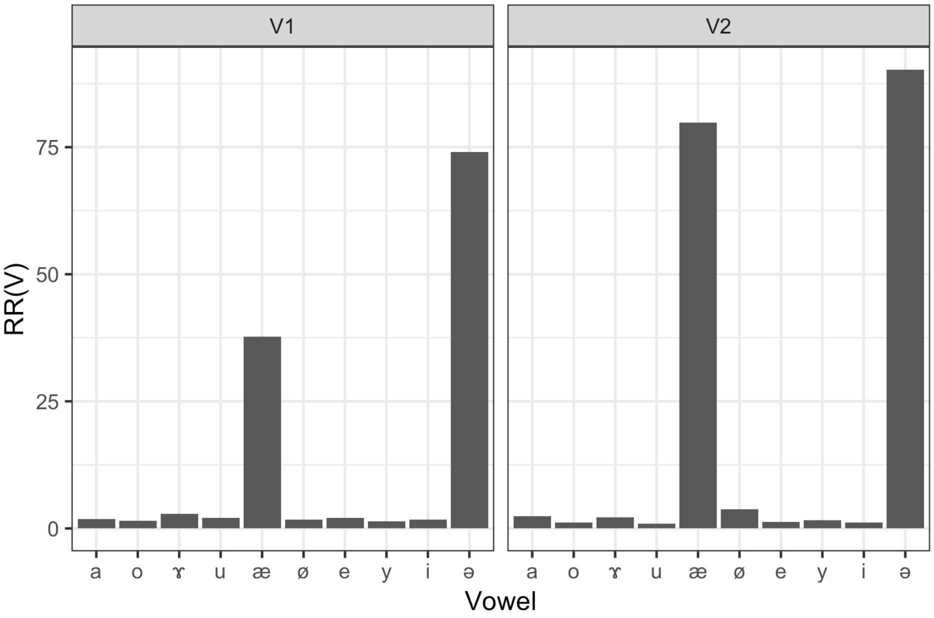

Figure 5.3: Mari RR(V)s

6 Discussion

The results of all three languages show a close correspondence between RR(V) and the notion of how much each vowel is participating in harmony, based on its O/E values. For example, in

Chuvash, the O/E values of [y] in V1 position before a back vowel are very low, ranging from 0 to 0.17; correspondingly, RR(y) for V1 is very high, at 66.85. In contrast, the O/E values for [i] in V1 position before a back vowel in Chuvash are very high, ranging from 0.24 to 0.91, with most of the cases close to expected representation; correspondingly, RR(i) is very low in the V1 condition, at 2.3. The vowel [ɯ] in Chuvash is intermediate between [i] and [y], with O/E values of 0.02 to 0.35 when it occurs in V1 position before a front vowel, and RR(ɯ) in V1 condition intermediate at 10.31. A similar correspondence holds across other vowels and other languages, suggesting that RR(V) is indeed a good measure of how much a vowel participates in harmony, where participation is viewed in terms of rarely co-occurring with vowels of the opposite harmonic class.

The concept of RR(V) allows us to look for patterns in gradient participation. Noticeable across all three languages is the fact that RR(i) is quite low, meaning it has a low degree of participation in the harmony system. In particular, in both Chuvash and Tatar, [i] has the lowest RR(V) of any front vowel in both positions, and one of the lowest overall.[60] This result is particularly interesting because, across front/back harmony systems, [i] is a common transparent vowel, and this pattern has often been connected to the lack of a back, high, unrounded counterpart to [i] in the inventory in languages with transparent [i], like Finnish and Hungarian (e.g. Vago 1973). As described above, [i] is not consistently described as neutral to harmony in any of the languages examined here, and in Chuvash it even has a back counterpart, yet RR(i) shows that it nonetheless participates less in the harmony system, in a gradient way. Specifically, [i] in both Chuvash and Tatar shows less overall participation in harmony than lower and/or rounded front vowels, which are less frequently transparent and neutral in harmony systems.

---

[60] It is unclear why some of the back vowels, in particular [o], often have lower RR(V) than [i]. It is possible that this results in part from the fact that if a front vowel is behaving neutrally, it will affect the RR(V)s of back vowels.

This result suggests that the cross-linguistic behaviour of [i] as a neutral vowel could extend beyond categorical properties and be due to reasons beyond harmonic pairing. As such, it suggests a need to re-examine more carefully the factors involved in harmony participation.

Beyond [i], it is notable that cardinal vowels in general tend to have lower RR(V), and are therefore less harmonic. For example, [a] and [u] also have low RR(V) values in Chuvash, and [a], [o], and [e] have low values in Tatar. One possible explanation for this result is that it is due to non-harmonic loans from Russian, which has neither harmony nor front rounded or back unrounded vowels. Indeed, Chuvash, Tatar, and Mari are all spoken in Russia and could be expected to have many recent, disharmonic Russian loans. Removing loans from the corpus would test this hypothesis. Alternatively, this result could illustrate a more general property about how vowels participate in harmony systems, with more marked vowels (i.e. front rounded and back unrounded ones, in these systems) being more required to harmonize. Indeed, across all three languages, particularly in Chuvash and Tatar, the more marked vowels have consistently higher RR(V) values, except in cases where their distribution is limited (e.g. V2 position of [ɯ] in Chuvash). This observation is consistent with existing ideas in the literature on categorical harmony, in which harmony has been argued to be a way to enhance the perceptibility of marked features (e.g. Kaun 1995).

The results from Mari are more complex and puzzling than those of Chuvash and Tatar; two vowels have very high RR(V) values, and the others are approximately equivalent. The vowels with high RR(V) were [æ] and [ə]; the former is perhaps expected, since it is more marked, but the latter raises questions about the nature of the front reduced vowel in Mari. However, it might be due to the fact that in at least Eastern Mari dialects, there are alternations between reduced and full vowels in particular word positions (Kangasmaa-Minn 1998), potentially masking actual vowel combinations. Specifically, there is a general relationship in Chuvash and Tatar between markedness and high RR(V) values, yet the schwa symbol typically denotes an unmarked vowel, so that relationship does not seem to hold in this case in Mari. Moreover, the other Mari vowels all behave essentially the same way, regardless of whether or not they are marked or exist in Russian, and the reasons for this result are unclear. Whether it says something about the Mari phonological system, the morphology (e.g. complex words with non-harmonizing vowels tending to avoid [æ] and [ə]), or loanwords from Russian remains to be investigated.

7 Conclusion

In summary, this chapter has explored a method of calculating segment-specific participation in vowel harmony, using relative risk. I have shown how this computation captures the notion of participation by examining relative risk in comparison to O/E values across corpora of three different front/back harmony systems: Chuvash, Tatar, and Mari. I then examined what the results imply about gradient participation. Specifically, I showed that [i] has a low degree of participation across these harmony systems, which corresponds to the fact that [i] is often neutral cross-linguistically, but is interesting because this neutrality is often considered categorical in other languages, in a way not consistently described for the languages considered here. The relative risk measure provides a way of looking at participation in harmony in a more nuanced way, and it allows for cross-linguistic comparison; this approach can therefore provide significant new insights into harmony systems.

Chapter 6: Participation in ATR harmony

1 Introduction

As in front/back harmony, a number of asymmetries are seen in participation in the typology of ATR harmony, which are the subject of the second half of this dissertation. I focus here on ATR-dominant harmony, in which combinations of ATR and RTR always harmonize to ATR, but I will also briefly discuss the differences between this type of harmony and other types of tongue root harmony. Some examples of ATR harmony in Lango are given in (1); note how the suffix alternates between the RTR low vowel [á] and the ATR low vowel [ə́] depending on whether the root is RTR or ATR.

(1) Lango ATR harmony (Woock & Noonan 1979)

ATR root vowels take ATR suffix -ə́	(a)	cíŋ	‘hand’	cíŋə́	‘my hand’
	(b)	ŋèt	‘side’	ŋètə́	‘my side’
	(c)	wót	‘son’	wódə́	‘my son’
	(d)	ŋùt	‘neck’	ŋùtə́	‘my neck’
RTR root vowels take RTR suffix -á	(e)	yíb	‘tail’	yíbá	‘my tail’
	(f)	lɛ́b	‘tongue’	lɛ́bá	‘my tongue’
	(g)	wàŋ	‘eye’	wàŋá	‘my eye’
	(h)	bwɔ́m	‘wing’	bwɔ́má	‘my wing’

ATR-dominant harmony is predominant in eastern Africa, particularly in the Nilo-Saharan family (e.g. Casali 2003; Rose 2018). The languages I discuss in detail in the rest of this section are all Western Nilotic, but the generalizations discussed in the present chapter are apparent throughout the cross-linguistics typology of ATR-dominant harmony.

Several typological overviews have provided a good understanding of the patterns within ATR harmony systems (e.g. Casali 2003; Rose 2018). I focus here on Nilo-Saharan languages; there are good descriptions of a number of harmony systems (e.g. Kutsch Lojenga 1991 on Alur; Andersen 1999 on Mayak; Woock & Noonan 1979 and Noonan 1992 on Lango; Hall & Yokwe 1981 and Steinberger & Vago 1987 on Bari; Quinn-Wriedt 2013 on Maasai; Coates 1985 on Otuho; Andersen 2007 on Kurmuk; Andersen 1989 on Päri; Reh 1996 on Anywa; Andersen 2004 on Jumjum; Lodge 1995 on Kalenjin; Andersen 1987 on Lulubo; Otero 2015 on (Ethiopian) Komo; Andersen 1986 on Lugbara; Kilpatrick 1985 on Bongo; Parker 1985 on Baka; Noske 1990, 1996 on Turkana; among many others), and for some languages (like Lango, the subject of Chapter 8) there has been a great deal of theoretical interest (e.g. Archangeli & Pulleyblank 1994; Potts et al. 2010). However, much remains unknown about ATR harmony (both ATR-dominant harmony and RTR-dominant harmony, in both Nilo-Saharan and Niger-Congo languages), since so many languages have not been studied, or at least not in detail. Even for many of the languages for which descriptions are available, there is often little phonetic work, few detailed analyses, and not the same extent of theoretical focus as has been put on better-known languages with front/back harmony like Hungarian, Finnish, and Turkish. Nonetheless, simply because so many languages have tongue root harmony, there is a relatively large body of descriptive and analytical literature about it, and certain languages have been studied in detail.

2 General patterns

2.1 ATR-dominant harmony: Definition and basic generalizations

ATR harmony in general is defined by agreement in the position of the tongue root (advanced versus retracted). Involvement of the tongue root position has been confirmed through articulatory studies of a number of languages (e.g. Jacobson 1980; Tiede 1996; Gick et al. 2006; Allen et al. 2013; among others); however, the mechanisms behind tongue root contrasts are complex, and can also involve factors like pharyngeal expansion and tongue body raising/lowering, including as the primary cues (e.g. Lindau 1987; Lindau et al. 1972; Lindau & Ladefoged 1986; Tiede 1996; see also Esling & Edmonson 2006 on aryepiglottic folds and tongue root). In ATR harmony systems, vowels typically agree for tongue root position within roots and across root-affix boundaries; harmony across heights (e.g. high vowels triggering harmony onto mid vowels and vice versa) is common. In fact, within some literature, including Stewart's (1967) definition of ATR harmony, cross-height harmony is required for a harmony to be considered ATR harmony.

ATR-dominant harmony, the focus of this section, is a particular type of tongue root harmony, in which the advanced (ATR) value is dominant, meaning that it always wins; in other words, both ATR…RTR and RTR…ATR sequences harmonize to ATR…ATR (see e.g. Casali 2003; 2008). As noted for instance by Casali (2008: 514), ATR-dominant harmony is typically distinguished by the presence of affixes (typically suffixes) that are invariantly ATR and that can trigger harmony onto RTR roots. According to Casali, such languages always also have 'recessive' affixes, which alternate in their ATR value depending on the value of the root. This distinction is often analyzed in terms of underlying feature values: ‘dominant’ invariant affixes that trigger harmony are underlyingly ATR, while ‘recessive’ alternating affixes are underlyingly RTR (Casali 2008). This type of dominance in harmony systems is common for tongue root harmony systems, but apparently rare for other types of (vowel) harmony (Casali 2008: 514), though it is not rare in nasal and other consonant harmonies (e.g. Hansson 2008). Casali (2003) notes a number of properties common in ATR dominant systems, including allophonic ATR vowels created through harmony and coalescence preserving ATR over RTR (e.g. /a+i/ → [e]). While dominant affixes were previously considered a defining feature of ATR-dominant harmony, distinguishing such systems from ‘stem-control’ harmony in which directionality of harmony is determined entirely by morphological structure (e.g. Baković 2000), it is important to note that this view is no longer held. Specifically, Casali (2008) notes that morphological structure in fact often does play a large role in ATR-dominant systems, and that dominant affixes are not the only possible indicator of ATR dominance. I assume that properties like ‘ATR dominance’ and ‘stem-control’ do not have a formal status, but rather identify common properties of patterns that fall out of certain constraints and their rankings.

Among tongue root harmony systems in African languages, ATR dominance is highly characteristic of languages in which there is a tongue root contrast in high vowels (i.e. with [i] versus [ɪ] and [u] versus [ʊ]), known as ‘2IU’ languages (Casali 2008). The majority of 2IU languages have harmony, and in most cases it is ATR-dominant (Casali 2003; 2008; Rose 2018). Languages without the contrast in high vowels (i.e. ‘1IU’ languages) may or may not have harmony (e.g. Rose 2018), but when they do, ATR is generally not dominant (e.g. Casali 2003).

The term ‘ATR-dominant’ is not without controversy, as not all literature agrees with this way of referring to a system; some instead argue that ATR is always the dominant feature in tongue root harmony (Bakovic 2000; Polgárdi 1998; van der Hulst & van de Weijer 1995; see also Casali 2003: 308 on 'universal [+ATR] dominance' theories). However, it is clear that there is a distinction between 2IU and 1IU languages, and to some extent also between Nilo-Saharan and Niger-Congo languages, in the way in which the tongue root harmony systems operate. Nilo-Saharan languages in particular are biased towards 2IU systems and it is 'characteristic' of them to show the asymmetry in which ATR acts like a dominant value (Casali 2008: 515).

2.2 Participation and neutrality in ATR-dominant harmony

As discussed above, languages with ATR-dominant harmony typically have a tongue root contrast in high vowels (i.e. 2IU). As is therefore expected, high vowels are participators in ATR-dominant harmony, both as triggers and targets. In fact, it has been argued that the presence of this contrast causes ATR to become 'active' in the phonology, due to difficulty in perceptibility of tongue root contrasts in high vowels (Rose 2018).

2IU languages may or may not have a tongue root contrast in mid vowels. In other words, they may be 1EO, with only RTR mid vowels [ɛ] and [ɔ] occurring contrastively, or they may be 2EO, with the contrasts [e] versus [ɛ] and [o] versus [ɔ]. However, it has been widely noted that in 2IU languages, mid vowels generally also participate in harmony, even when those languages are 1EO; in such cases, the mid vowels participate allophonically, such that [e] and [o] do occur on the surface in the languages, but only as the (allophonic) result of harmony (e.g. Casali 2003; 2008; Rose 2018).

It has been claimed (e.g. Archangeli & Pulleyblank 1994; Rose 2018) that mid vowels are the ideal targets of tongue root harmony, including in ATR-dominant harmony. This is based on their tendency to participate allophonically in harmony even when there is no contrast; the argument is that tongue root distinctions are more perceptible on mid vowels, so there is a perceptual advantage to spreading ATR to them (Rose 2018). However, in many of the languages with allophonic mid vowel participation (e.g. Mayak; see Chapter 9), as well as in some 2EO languages (e.g. Lango; see Chapter 8), mid vowels participate in a more limited range of contexts than high vowels. For example, in both Mayak and Kinande, high vowels are targets of harmony in both directions, but mid vowels are only targets for regressive harmony (Andersen 1999 on Mayak; Archangeli & Pulleyblank 2002 on Kinande). This is unexpected if mid vowels are the best possible targets, but expected under a view in which either high vowels are ideal targets, either on independent factors or because they agree in height with the ideal triggers (which are high).

The most significant complications regarding participation in ATR-dominant harmony come from low vowels. These languages may have either a single contrastive low vowel, the RTR [a], or a contrast in low vowels, with the ATR one typically notated as either [ʌ] or [ə] (or sometimes [ä]). Among languages that have only [a] contrastively, it may be neutral to harmony, participate allophonically, or participate by re-pairing to an ATR mid vowel ([e] or [o]). This can vary within a language: in Kinande, for instance, [a] is neutral in progressive harmony, but participates allophonically in regressive harmony (e.g. Gick et al. 2006; Archangeli & Pulleyblank 2002); in Maasai, [a] is neutral to regressive harmony, but re-pairs to [o] in progressive harmony (e.g. Baković 2000; Quinn-Wriedt 2013). Among languages with a tongue root contrast in the low vowels, they can again either participate in harmony or not, and this again varies both within and across languages. In Lango (Chapter 8), for instance, [a] harmonizes to [ə] in progressive harmony across an intervening consonant cluster only when the trigger is high (Woock & Noonan 1979; Archangeli & Pulleyblank 1994). In Mayak (Chapter 9), [a] and [ʌ] are contrastive, but [a] does not undergo harmony and [ʌ] does not trigger it (Andersen 1999). In other words, there is a large range of variation regarding the participation of low vowels in ATR-dominant harmony. Unfortunately, existing surveys of tongue root harmony systems typically focus on high and mid vowels, leaving this area relatively unexplored.

It is worth noting that more recent phonetic work has sometimes revised our understanding of which vowels participate in ATR harmony and which do not. For example, in Kinande, it was previously thought that [a] was consistently neutral to ATR harmony (e.g. Schlindwein 1987), but ultrasound evidence has shown that it in fact participates allophonically (Gick et al. 2006). This is particularly common among low vowels, where allophonic and even contrastive ATR low vowels have been missed, and even native speakers have had difficulty noting the contrasts or recognizing them as ‘important’ to the phonology (e.g. Coates 1985 on Otuho). However, it is also widely noted that RTR high vowels are perceptually confusable with ATR mid vowels (e.g. [o] and [ʊ]), such that impressionistic descriptions have in some cases had difficulty distinguishing them (e.g. Casali 2003; Rose 2018). Recent phonetic studies have therefore also revised the inventory and patterns in non-low vowels. One case in particular is Tunen, where what was described in earlier work (Mous 1986; van der Hulst, Mous, and Smith 1986) as an odd [e]~[i] alternation in the class 7 prefix, with the RTR roots taking an ATR prefix [e] and ATR roots taking the [i] alternant, has been reanalyzed as a more standard harmonic [ɪ]~[i] alternation by Boyd (2015).

2.3 Other types of tongue root harmony

While the generalizations discussed above are true of ATR-dominant harmony systems, they are not true of other types of tongue root harmony. The other (contrastive) inventory type that often has tongue root harmony is 1IU, 2EO; in other words, languages in which the contrastive inventory among high vowels consists only of ATR high vowels [i] and [u], but among mid vowels has [e] versus [ɛ] and [o] versus [ɔ]. This is true of certain dialects of Yoruba, including Standard Yoruba, which has harmony that has been described as RTR dominant (e.g. Archangeli & Pulleyblank 1994). In these types of languages, high vowels are almost always neutral to harmony; only a handful of cases of 1IU 2EO systems, such as certain dialects of Yoruba (e.g. Ijesa; Oyelaran 1973; Przezdziecki 2005), have allophonic participation of the high vowels, while most languages with the contrastive 1IU 2EO inventory type have neutral high vowels (Casali 2008; Rose 2018). Note that high vowels are neutral (either transparent or opaque) in other Yoruba dialects (Przezdziecki 2005).

In these 1IU cases, then, mid vowels are very evidently the ideal participators in the harmony system; often, there are at least some contexts in which harmony affects only mid vowels. In Ogori/Oko, for example, mid vowels must be harmonic within roots, but high and low vowels can co-occur disharmonically in roots (Chumbow 1982). Like for ATR-dominant harmony, there is often neutrality of the low vowel in these 1IU systems[61], and as mentioned, high vowels are also typically neutral.

Beyond differences in inventory structure and which segments participate in harmony, these other types of tongue root harmony systems are not ATR-dominant. They are often root-controlled, in that the ATR/RTR value of the root determines harmony of the entire word; some systems like Yoruba have been described as RTR-dominant (e.g. Archangeli & Pulleyblank 1994; Pulleyblank et al. 1995; Leitch 1996), which is a controversial claim, as it has been suggested by others (e.g. Bakovic 2000; Polgárdi 1998; van der Hulst & van de Weijer 1995)

---

[61] Unfortunately, most typological work focuses on contrasts in high and mid vowels only. It is unclear whether there are 1IU languages with a phonemic contrast in low vowels, or to what extent allophonic participation of low vowels in root-controlled 1IU languages occurs.

that ATR is always the active feature. (See also Casali 2003 on this issue.) Given that I focus solely on ATR dominant harmony here, I do not pursue these arguments or assess them.

An account of the differences between ATR-dominant harmony and other types of tongue root harmony, and all of the relevant correlations among inventory, dominance, and participation, is beyond the scope of the present dissertation. However, it is a particularly important question for future research to look at why a contrast in the high vowels appears to be so critical to determining how harmony operates, and why certain vowels (in particular high vowels) appear to be good participators in one type of tongue root harmony (ATR-dominant) but poor participators in others (root-controlled or RTR-dominant), when the same features are involved in all types. Casali (2008: 527) notes that 'a fully satisfactory solution to the problem of explaining system-dependent [ATR] dominance has yet to be found', although attempts have been made in both purely theoretical and more phonetically based research (e.g. van der Hulst 2018 for a theoretical analysis; Rose 2018 for a perceptual explanation).

3 Harmony in the proto-language

Since all the languages I will discuss in detail in the following chapters are Nilotic, I will briefly discuss harmony in Proto-Nilotic and what it can suggest about how certain patterns in harmony participation developed. The Nilotic family tree is generally assumed to involve three distinct branches: Western, Southern and Eastern Nilotic, with the possibility that the Southern and Eastern branches form their own sub-branch separate from Western Nilotic (Köhler 1948; Dimmendaal 2002). Substantial reconstructive work has been done on both Southern and Eastern Nilotic, with less on Western Nilotic (e.g. Rottland 1982 on Southern Nilotic; Vossen 1982 on Eastern Nilotic; Andersen 1989; 1990 on Western Nilotic; see also Dimmendaal 2002 on all branches). Both Southern Nilotic and Western Nilotic are reconstructed as having a ten-vowel system, with all vowels paired for ATR (Rottland 1982 on Southern Nilotic; Andersen 1990 on Western Nilotic); Eastern Nilotic is reconstructed with uncertainty as to whether there are nine or ten vowels, with the ATR low vowel possibly missing from the proto language (Vossen 1982).[62] The harmony systems have developed in a number of distinct ways, with some descendant languages losing harmony and even tongue root contrasts entirely, including entire branches of Western Nilotic (Dinka-Nuer, which instead contrasts breathiness; e.g. Dimmendaal 2002; Storch 2005). Among the languages with harmony, one of the primary distinctions is in whether the ATR low vowel was kept, and if not, what development led to its loss and what effect that had on the harmony system. Note that the ATR low vowel is often notated as [ä] in Nilotic transcription traditions.

In Southern Nilotic, in the Elgon and some other clusters of Kalenjin, *[ɛ] merged with [a] and *[ä] with [e], such that both *[ɛ]~*[e] and *[a]~*[ä] in the proto-language harmony system became an [a]~[e] alternation in the modern harmony system (Rottland 1982). In other words, this diachronic change led to the historical appearance of re-pairing in harmony. In the Päkoot cluster of Kalenjin, it appears that the former merger also occurred, but the latter occurred in the opposite direction (i.e. *[e] to [ä]), such that the *[ɛ]~*[e] alternation from the proto-language merged with the [a]~[ä] alternation (Dimmendaal 2002). These languages therefore maintain the low vowel alternation.

In Eastern Nilotic, the Bari subgroup has the ATR low vowel, while the Teso-Lotuko-Maa group typically does not (Vossen 1982), though [ä] is maintained after glides in some cases

---

[62] However, given that the other two branches are assumed to have had the low ATR distinction, we can likely assume that an ATR low vowel existed in Proto-Nilotic, which means that the question is really just about when the loss of the distinction happened for Eastern Nilotic. In other words, the question is whether Proto-Eastern-Nilotic as a whole lost the ATR low vowel from Proto-Nilotic, and the Bari subgroup regained it, or whether the merger happened at a later stage. This issue is not crucial for the present purposes.

in Lotuko (Muratori 1938; Dimmendaal 2002). However, even in Bari, [a] can in some instances be preceded or followed by ATR vowels in roots (Spagnolo 1960), suggesting some degree of neutrality or that it was previously ATR there. This is interesting because it suggests that some of the neutrality patterns of [a] might be reconstructable as a neutrality effect in the proto-language, perhaps in the same way that front/back harmony only affected low vowels in Proto-Finnic (see Chapter 2). Bari [ä] typically corresponds to [a] in the other subgroup of Eastern Nilotic, but seems to have merged with [ɔ][63] in Teso-Turkana, and as a result there are re-pairing harmonic alternations of [a]~[o] in several Eastern Nilotic languages (Dimmendaal 2002).

In Western Nilotic, the Lwoo subgroup is the one with full harmony systems, with the individual languages varying in whether they have nine vowels (i.e. without the ATR low vowel) or ten (e.g. Dimmendaal 2002). In Päri, which has [ä] contrastively, this vowel can optionally be realized as [o] near labial consonants and [e] elsewhere (Andersen 1989). In languages lacking the ATR low vowel, like Alur, an individual morpheme with *[ä] in the proto-language will have either [e] or [o] in the daughter language, depending on environment; however, in harmonic alternations, *[ä] as the ATR counterpart of [a] (in the same morpheme) consistently became [e], never [o], regardless of contextual factors (Andersen 1989; Dimmendaal 2002). In other words, [a] alternates with [e] in Western Nilotic languages that lack [ä]; the 'elsewhere' environment from root effects was the one generalized. Interestingly, in some Western Nilotic languages, specifically Acholi and Lango, the original *[ä] was lost, merging with [e] and [o], but then [ä] was reintroduced through harmony and now exists contrastively (Dimmendaal 2002). Such a development is highly suggestive of the instability of low vowels in ATR harmony systems.

---

[63] Apparently, it did in fact merge with an RTR vowel, not with the ATR [o]. The details of this merger and its effects on the harmony systems are beyond the scope of the present work, but interested readers are referred to Dimmendaal (2002) and references therein for a general overview of Nilotic reconstruction, and to Noske (1990; 1996) on Turkana.

In general, the most studied, and perhaps most interesting, harmony-related developments from Proto-Nilotic to the modern languages consist of what has happened to the low vowels. A large number of distinct patterns in low vowel behaviour exist in Nilotic languages, including both participation and neutrality for contrastively paired low vowels, as well as neutrality, re-pairing to [e], and re-pairing to [o] when [a] lacks an ATR counterpart in the inventory. Many of these patterns can be traced to specific mergers involving the ATR low vowel. It is particularly interesting that these mergers occurred in different ways and directions depending on language (or subgroup), and that there is no clear correlation between the presence of [ä] and the behaviour of low vowels in the harmony system.

4 Why we might want to rethink how participation gets defined

As discussed in Section 2, the factors affecting participation in ATR harmony are highly complex. Mid vowels typically participate in ATR-dominant harmony, whether they are contrastively paired or not, although their participation may be limited in ways in which that of high vowels is not (e.g. by directionality). The behaviour of low vowels, on the other hand, is much more variable, both in the options for participation when there is no contrast (allophony versus re-pairing), and in the fact that absolute neutrality (i.e. non-participation in any context) seems to be far more common among low vowels than mid vowels. Particularly interesting is the case of Mayak (see Chapter 9), in which the mid vowels do not contrast for ATR, yet participate in regressive harmony, while the low vowels do contrast for ATR, yet are neutral (Andersen 1999).

It is clear, then, that any relationship between contrast and participation in ATR-dominant harmony is non-trivial. Among non-high vowels in particular, participation is predicted more readily by vowel quality, specifically height (mid versus low), than by contrast. This is also the case when comparing high versus mid vowels, including in 2IU, 2EO languages, in some of which high vowels are targets of harmony in contexts in which mid vowels are not. In other words, there is an evident height effect to participation in ATR-dominant harmony, which cannot be captured through a simple consideration of which vowels are contrastive.

5 Structure of the rest of Part 2

The rest of Part 2 (the section on ATR harmony) is structured as a series of chapters on three Western Nilotic languages: Alur, Lango, and Mayak. First, in Chapter 7, I present an analysis of Alur, in which an unpaired [a] participates in some contexts by re-pairing with [e]. This analysis does not require the target scaling factors developed through this dissertation, but is consistent with it. Moreover, it is an important illustration of one of the mismatches between participation and contrast. Next is Chapter 8 on Lango, which has a complex system of interactions among triggers, targets, and distance; I show that the framework developed for front/back harmony, in Chapters 3 and 4 on Hungarian and Votic, is also able to deal with Lango in a simpler way than previous analyses. Finally, Chapter 9 focuses on Mayak, in which the low vowels are contrastive for ATR, yet nonetheless neutral to the harmony system. The same way of conceptualizing participation as developed throughout this dissertation can deal with this pattern, while any approach drawing an equivalence between neutrality and lack of contrast cannot.

Chapter 7: An analysis of ATR harmony in Alur

1 Introduction

Alur is an understudied Western Nilotic language spoken primarily in Uganda. One of the few descriptions of the phonology of the language is a short paper on the harmony system by Kutsch Lojenga (1991). In this paper, she outlines the complex ATR harmony pattern, which depends on the combination of whether a root is a noun or a verb, whether the potential target vowel is [a] or not, and whether a suffix is of form V or CV. She does not provide an analysis of the pattern.

In this chapter, I analyze the Alur harmony pattern as described by Kutsch Lojenga (1991). I show that both nouns and verbs in fact behave fundamentally the same way, and that the discrepancies between their behaviour arise from the fact that verb roots are of form CVCV, whereas noun roots are of form CVC. I argue that the domain of harmony in Alur is a binary foot aligned to the left edge of the root, similar to an argument made about the closely related language Lango by Pulleyblank (2001). Harmony is ATR-dominant, and within the foot domain; whether it occurs depends on a combination of directionality, the number of intervening consonants, and whether an additional feature change is required to create an ATR vowel. Specifically, regressive harmony is exceptionless within the foot domain, whereas progressive harmony is blocked in the presence of a coda consonant or if the target vowel is [a], which must also change the feature [low] in order to become [ATR]. I propose a straightforward analysis that captures the complex patterning of both nouns and verbs.

The proposals made in this dissertation predict a range of patterns, including one where trigger and target scaling produce the kind of pattern in which inventory properties correlate with neutrality. I illustrate this using Alur, showing how the proposals made here and more conventional accounts can capture the Alur pattern. Unlike the other languages analyzed in this dissertation, an analysis of Alur does not require trigger and target scaling factors; the only vowel quality-based distinction in behaviour is regarding [a] as a target, but that can be attributed to inventory asymmetries. In Alur, trigger and target scaling are simply set as equivalent for all vowels. As such, while the scaling factors are not necessary to account for Alur, the analysis here is entirely consistent with the Alur pattern and can account for it. This evidently is important for an account meant to deal with all typologically attested types of relationships between lack of contrast and neutrality.

The chapter is organized as follows: Section 2 describes the Alur harmony pattern and my informal analysis, Section 3 formalizes the account, and Section 4 discusses the implications.

2 Alur vowel harmony pattern

2.1 Inventory

Alur has a nine-vowel inventory given in Table 7.1. High and mid vowels are paired for ATR, while the low vowel [a] is unpaired.

	ATR			RTR		
	Front	Central	Back	Front	Central	Back
High	i		u	ɪ		ʊ
Mid	e		o	ɛ		ɔ
Low					a	

Table 7.1: Vowel inventory of Alur (Kutsch Lojenga 1991)

2.2 Absence of harmony with prefixes

Alur has a complex system of ATR harmony, triggered by ATR vowels and targeting RTR vowels. This harmony occurs only between roots and suffixes, as will be shown in the following subsections. However, prefixes are not involved in harmony, as the examples in (1) show. All data in this paper comes from Kutsch Lojenga (1991). Regardless of the ATR values of prefix or root vowels, prefixes never alternate and never trigger alternations in roots; the examples in (a) and (d) are disharmonic. I propose that prefixes are outside of the domain of harmony, which I suggest begins at the left edge of the root. I will mark this domain with parentheses in the examples, and argue further about this domain throughout the rest of Section 2.

(1) (a) ì-(nɛ́nɔ̀) ‘2SG-saw’
(b) ɛ́-(nɛ̀nɔ̀) ‘3SG-saw’
(c) ì-(ŋéyò) ‘2SG-knew’
(d) ɛ́-(ŋèyò) ‘3SG-knew’

2.3 Suffix-to-root harmony

In Alur, harmony between root and suffix is described as depending on the interaction of a number of factors: whether the root is a noun or a verb, whether the suffix is of form V or CV, whether the potential trigger vowel is in the root or suffix, and whether the potential target vowel is low or not. I argue that both nouns and verbs behave fundamentally the same way, and that the complexities arise from a combination of domains, directionality, distance, and re-pairing, all of which are well-known phonological factors in the harmony literature.

The examples in (2) show the behaviour of suffix-to-root harmony in verb and roots respectively. ATR suffixes of form V in Alur always trigger harmony to underlyingly RTR roots, as seen in (2a) versus (2b) for verbs and (3a) versus (3b) for nouns. In verbs, ATR CV suffixes trigger no alternation (2c). In contrast, ATR CV suffixes do trigger harmony to noun roots (3c).

Note that verbs in the dataset are CVCV, with root-final elision with a V suffix, explaining the presence of the [ɔ̀] in (2c) but not (2a,b).

(2) (a) ɛ́-(nɛ̀n-á) ‘3SG-saw-1SG’
(b) ɛ́-(nèn-í) ‘3SG-saw-2SG’
(c) ɛ́-(nɛ̀nɔ̀)-wú ‘3SG-saw-2PL’[64]

(3) (a) (tʃɔ̀ng-á) ‘knee-1SG’[65]
(b) (tʃòng-í) ‘knee-2SG’
(c) (tʃòng-wú) ‘knee-2PL’

I argue that this distinction between nouns and verbs is due to domains: similar to what was argued for the closely related language Lango by Pulleyblank (2001), I suggest that the domain of Alur harmony is a bisyllabic foot aligned to the left edge of the root, marked with parentheses in the examples. Noun roots contain only a single vowel, meaning that both V and CV suffixes fall within this domain of harmony. In contrast, verb roots are of form CVCV, where the final root vowel gets elided with a V suffix, but not with a CV suffix. As such, with verb roots, only V suffixes can be incorporated into the harmony domain; CV suffixes are outside the binary foot, and therefore the result in (2c) is predicted. While harmony in some Nilotic languages, such as Lango, can be sensitive to the type of suffix, regardless of the size of the stem, here it seems to be about the foot domain, since object suffixes of form V do trigger harmony in verbs (2b), and

---

[64] The [ɔ̀] in this example is part of the verb root, present in all examples except those with V suffixes. Several facts suggest this is deletion of a root-final vowel rather than epenthesis. First, different vowels appear in different verb roots, in a way that does not appear predictable from the context. Second, noun roots are consistently C-final, even when unsuffixed and when preceding C-initial suffixes, unlike verb roots which are vowel-final in those contexts and C-final only before V-initial suffixes. The behaviour of nouns suggests that codas are permitted, giving no reason for epenthesis. In contrast, there are no VV sequences in the data, so that a deletion analysis could be motivated by hiatus-avoidance.

[65] Kutsch Lojenga (1991) writes these forms with [ng]; it is unclear whether the nasal there is actually pronounced as alveolar or whether perhaps the IPA transcription is [ŋg]. Moreover, she writes these as sequences, not as prenasalized stops. I follow her transcriptions throughout.

what appear to be the same series of suffixes attached to nouns (as possessive suffixes) do trigger harmony even when of form CV (3c).

2.4 Root-to-suffix harmony

As expected by the foot domain analysis, RTR V suffixes undergo harmony with ATR verb roots, as shown in (4a) versus (4b), while CV suffixes do not (4c). However, despite being in the domain of harmony, RTR CV suffixes do not undergo harmony with ATR noun roots (5c), even though RTR V suffixes do, as seen in (5a) versus (5b). Note that whether the trigger is the root or suffix depends on which morpheme has an underlyingly ATR vowel.

(4) (a) ɛ́-(gùd-é) ‘3SG-hurt-3SG’
(b) ɛ́-(nɛ̀n-ɛ́) ‘3SG-saw-3SG’
(c) ɛ́-(gùdò)-gí ‘3SG-hurt-3PL’

(5) (a) (lìmb-é) ‘cheek-3SG’
(b) (tʃɔ̀ng-ɛ́) ‘knee-3SG’
(c) (lìmb-gí) ‘cheek-3PL’

I analyze this effect of suffix onset in nouns in (5c) as being due to distance and directionality. I propose that the NC sequences written by Kutsch Lojenga (1991), like in [lìmb] ‘cheek’, are behaving phonologically as prenasalized stops, not sequences. Under this analysis, all nouns in the data are of form CVC. As such, with V suffixes, there are no codas, while with CV suffixes, the root-final C is a coda. This structure is illustrated in the examples in (6), repeated from (5), with periods as syllable breaks.

(6) (a) (lì.^{m}b-é) ‘cheek-3SG’

(b) (tʃɔ̀.^{n}g-ɛ́) ‘knee-3SG’

(c) (li^{m}b.-gɪ́) ‘cheek-3PL’[66]

While regressive harmony in Alur is exceptionless within the foot, I suggest that progressive harmony cannot cross multiple consonants, or alternatively cannot cross a mora. Similar moraic distance effects have been documented in related languages (e.g. it has been claimed for Lango, as in Archangeli & Pulleyblank 1994; note however that geminates and non-geminate clusters behave differently in Lango, so it may not just be about moras). Thus, regressive harmony in Alur is strong enough to overcome the distance effect, while progressive harmony is not. Since codas do not appear in verb roots, this analysis remains consistent with the behaviour of verbs.

2.5 Behaviour of [a]

A final issue concerns the behaviour of [a], which has no ATR counterpart in the inventory. In suffixes, [a] is neutral to progressive harmony, in that onsetless [a] suffixes do not harmonize in the presence of ATR roots, unlike other onsetless RTR suffixes. This is shown in (7); compare (4a-b) and (5a-b) above.

---

[66] Note that the set of suffixes in Kutsch Lojenga’s (1991) Alur dataset is limited, such that the only V suffixes are /a/, /ɛ/, and /i/, while the only CV ones are /wa/, /wu/, and /gɪ/; as such, there is ambiguity as to whether the distinction here is in terms of distance or targeting of high/low versus mid vowels in progressive harmony. Given the pattern in the closely related language Lango, I adopt a distance-based analysis here, but additional data from other types of suffixes is necessary to disambiguate the possibilities in Alur.

(7) (a) ɛ́-(gùd-á) ‘3SG-hurt-1SG’

(b) (lìmb-á) ‘cheek-1SG’

In contrast, [a] in roots does undergo regressive harmony in the presence of an ATR suffix, by re-pairing to the mid ATR vowel [e]. This is shown in (8) and (9); the root surfaces with [a] when the suffix is RTR, in the (a) forms, but with [e] when the suffix is ATR, in the (b) forms. Note that regressive harmony is strong enough to trigger harmony onto [a] even when there is an additional intervening consonant, as in (9c).

(8) (a) ɛ́-(tʃàk-á) ‘3SG-chose-1SG’

(b) ɛ́-(tʃèk-í) ‘3SG-chose-2SG’

(9) (a) (wàŋ-á) ‘eye-1SG’

(b) (wèŋ-í) ‘eye-2SG’

(c) (wèŋ-wú) ‘eye-2PL’

Similar to the fact that CV suffixes trigger but do not undergo harmony in nouns, I analyze this effect as resulting from a stronger tendency for regressive than progressive harmony in Alur. Alur does not tolerate ATR low vowels as the product of harmony; there is independent evidence of this from the fact that RTR…ATR sequences with [a] harmonize by changing the feature [low], as in (8-9). Given this fact about the grammar, we can understand the lack of harmony with [a] in ATR…RTR scenarios in (7) as upholding a ban on ATR low vowels in the language (i.e. structure-preservation), but opting for disharmony over re-pairing. Regressive harmony is strong enough to overcome the additional feature change required for re-pairing, while progressive harmony is not. This is consistent with reports of progressive harmony being weaker in a number of ATR languages, which can manifest in various ways, including non-iterativity (failure to extend more than one or two vowels), lack of re-pairing, and targeting a subset of vowels.

2.6 Summary of facts

To summarize, Alur has ATR-dominant vowel harmony between roots and suffixes. Within verbs, only V suffixes trigger and undergo harmony; within nouns, only V suffixes undergo harmony, but both V and CV suffixes trigger it. Within roots (both verbs and nouns), /a/ harmonizes to [e] in the presence of an ATR suffix; in contrast, suffix /a/ never harmonizes. I suggest that this complex pattern emerges simply through the interaction of straightforward phonological factors: a foot domain, stronger regressive than progressive harmony, a distance effect, and a re-pairing versus neutrality effect. In the subsequent sections, I will formalize this analysis in Harmonic Grammar and discuss broader implications.

It is worth noting some additional facts beyond the harmony system that these data show. Specifically, verb roots are CVCV, but the root-final vowel deletes in the presence of a V-initial suffix. Noun roots are CVC, and when there are two root-final consonants, they always consist of an NC sequence; these sequences seem to be behaving phonologically as prenasalized stops, and I analyze them as such.

3 Analysis

3.1 Analysis background

With respect to the domain, I follow Pulleyblank (2001) on Lango, and I propose that the domain of harmony in Alur is the foot. This foot is binary (in terms of syllables) and aligned to the left edge of the root.[67] All vowels within the foot are required to harmonize; vowels outside the foot are not. Since the foot is aligned to the left edge of the root, prefixes are always outside of the domain of harmony; suffixes may or may not be within the domain, depending on the number of vowels in the root. In CVCV verb roots, deletion of the root-final vowel with V suffixes, but not with CV suffixes, means that V suffixes are part of the foot domain, but CV ones are not. Thus, we expect harmony exactly between roots and V suffixes in verbs. In CVC noun roots, both V and CV suffixes belong to the foot domain; we would therefore predict harmony between the root and all suffixes.

There are two remaining points to explain: that CV suffixes do not undergo harmony despite being in the domain for nouns, and that suffix [-a] does not undergo harmony, even though root [a] does. These can be generalized. ATR-[a] sequences are permitted within the foot, even though [a]-ATR sequences are not, and ATR-CC-RTR sequences are permitted within the foot, even though RTR-CC-ATR sequences are not. The latter is only apparent with noun roots, but there is no contradictory data in verb roots, since the only place where VCCV sequences can arise is with a noun root (…VC-CV). Both of these facts are indications of the stronger regressive than progressive harmony, with progressive harmony limited by vowel quality ([a])

---

[67] Given the lack of available data on Alur, it is unclear whether there is any independent evidence, from tone or stress, for this type of foot structure in Alur. However, it is worth noting that Lango, which is fairly closely related and much better documented, is described as having root-initial stress (Noonan 1992); if true in Alur, that would provide additional evidence for this foot-based analysis. Moreover, Dimmendaal (2012) notes that metrical structure is understudied in Nilotic languages, but there is evidence for its relevance to phonological processes.

and by distance (CC). Table 7.2 summarizes the permitted and forbidden foot-internal vowel sequences, with rows indicating the first vowel and columns the second.

V1 ↓ V2 →	[a]	Non-low RTR vowels	ATR vowels
[a]	OK	OK	*
Non-low RTR vowels	OK	OK	*
ATR vowels	OK	OK across a coda * elsewhere	OK

Table 7. 2: Permitted and forbidden foot-internal vowel sequences

3.2 Formalization: Foot structure

Before dealing with harmony, I formalize my analysis of foot structure and vowel deletion. In order to assign the correct foot structure and delete the root-final vowel in verbs before vowel-initial suffixes, we require the constraints defined in (10).

(10)

ALIGN-FT-RT(L,L): assign a violation for every foot that is not left-aligned to the left edge of a root

FT-BIN: a foot must be binary (at the syllable level)

PARSE-SYLL: all syllables should be parsed into feet

*VV: assign a violation for every VV sequence in the output

REALIZE-MORPHEME: assign a violation for every morpheme with null phonological exponence in the output (Kurisu 2001)

I first show that these constraints achieve the desired foot structure and root-final vowel deletion, and then omit them from further tableaux for space reasons; they do not interact with harmony except to define the domain.

The most interesting case is that of a verb root with a V suffix, shown in Table 7.3. Either the root-final vowel or the suffix vowel must be deleted because of *VV; the former is chosen because the suffix vowel needs to be maintained to satisfy Realize-Morpheme.[68] The choice of foot placement and size is established through a combination of Align-Ft-Rt(L,L), Ft-Bin and Parse-Syll, requiring a single binary foot at the left edge of the root in both Table 7.3 and Table 7.4. In Table 7.4 in particular, Ft-Bin keeps the suffix out of the domain of harmony.

	/ɛ́-nɛ̀nɔ̀-í/	Align-Ft-Rt(L,L)	*VV	Realize-Morpheme	Parse-Syll
	ɛ́-nɛ̀nɔ̀-í				**!**
	ɛ́-(nɛ̀nɔ̀)-í		*!		**
☞	ɛ́-(nɛ̀n-í)				*
	ɛ́-(nɛ̀nɔ̀)			*!	*
	(ɛ́-nɛ̀n-í)	*!			
	(ɛ́-nɛ̀n)-í	*!			*

Table 7.3: Deletion of root-final vowel in verbs with a V suffix

---

[68] Alternatively, this constraint could be thought of as Anchor-L; the suffix has only a single segment, and so deleting that segment means deleting the first one, thereby violating Anchor-L. As an additional alternative, it may be a preference for V1 deletion, which could be disambiguated if we had data with VC suffixes.

	/ɛ́-nɛ̀nɔ̀-wú/	FT-BN	PARSE-SYLL
☞	ɛ́-(nɛ̀nɔ̀)-wú		**
	ɛ́-(nɛ̀nɔ̀-wú)	*!	*

Table 7.4: Foot structure for verb with CV suffix

Thus, this combination of constraints, with the only crucial ranking being that PARSE-SYLL ranks below all other constraints, can correctly predict the proposed foot structure that functions as the harmony domain. I now assume the correct foot structure and vowel deletion patterns in all further tableaux.

3.3 Harmony constraints

I adopt two harmony constraints to explain the Alur pattern, both of which apply only within the domain of the foot. Since regressive and progressive harmony pattern differently, Alur requires separate constraints for enforcing harmony in ATR…RTR (11a) and RTR…ATR (11b) sequences; the latter should be weighted higher, since harmony is stronger regressively than progressively. Presumably, there are additional similar harmony constraints, not restricted to the foot domain, that are too low-weighted to show an effect in Alur. This style of harmony constraint is adopted from Pulleyblank (2004).

(11) Harmony constraints

(a) $*[\text{ATR}]\infty\underline{[\text{RTR}]}_{\text{Ft}}$: within the foot, an [RTR] vowel may not be preceded by an [ATR] vowel at any distance.

(b) $*\underline{[\text{RTR}]}\infty[\text{ATR}]_{\text{Ft}}$: within the foot, an [RTR] vowel may not be followed by an [ATR] vowel at any distance.

Since the number of intervening consonants affects harmony, Alur requires a way of differentiating harmony applying across a single consonant versus across any number of consonants.[69] As such, I adopt distance scaling defined in (12). The relevant unit in Alur, as in Lango in Chapter 8, is the number of consonants; the effect of harmony will therefore be weaker across multiple consonants than across a single one. It is worth noting that consonants are known to block ATR harmony in other languages, such as Assamese (Mahanta 2008).

(12) Scaling factor: distance

For a trigger α and a target β, multiply the penalty earned by a constant x (such that $0<x<1$) for each unit (e.g. syllable) of distance d intervening between α and β.

Since Alur treats all triggers and targets equally, we can set trigger and target scaling to be the same for all vowels. These concepts are defined as in (13-14). As in previous chapters, these scaling factors apply equally to both harmony constraints, but to no other constraints.

(13) Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that $x \geq 1$) for each degree i to which α is perceptually impoverished with respect to ±F.

---

[69] Since feet are maximally bisyllabic, the constraint across any number of consonants could alternatively be conceptualized as one limited to adjacent syllables (regardless of non-syllabic intervening material) or even one that is completely unbounded. Moreover, assuming that the first consonant of a CC cluster is a coda and that Alur codas are moraic, these differences could alternatively be expressed as harmony applying without an intervening mora versus across any intervening material. I express the patterns in terms of consonants for notational simplicity.

(14) Scaling factor: target quality

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant *x* (such that *x* ≥ 1) for each degree *i* to which β is a quality target (defined by phonetic factors discussed in Section 3) with respect to ±F.

In addition to the harmony constraints, I also adopt the faithfulness constraints in (15). Two are related to changes in tongue root, with separate faithfulness for [ATR] versus [RTR] to explain why harmony is ATR-dominant; a higher weighting of IDENT-IO[ATR] than IDENT-IO[RTR] ensures that when a vowel is required to change to harmonize, ATR will always be preserved at the expense of RTR. A third faithfulness constraint to [low] is necessary to deal with the fact that /a/ sometimes changes its [low] value to harmonize to [e], but sometimes does not.[70]

(15) Faithfulness constraints

IDENT-IO[ATR]: for a segment x in the input and its correspondent x’ in the output, assign a violation if x is [ATR] and x’ is not

IDENT-IO[RTR]: for a segment x in the input and its correspondent x’ in the output, assign a violation if x is [RTR] and x’ is not

IDENT-IO[low]: for a segment x in the input and its correspondent x’ in the output, assign a violation if x and x’ do not have the same value for the feature [low]

To prevent direct harmony of [a] to an ATR low vowel, we require the markedness constraint in (16). In other words, we want to prevent an allophone [ʌ] of /a/ from occurring even in ATR

---

[70] For ease of exposition, I use only a single symmetric faithfulness constraint for [low], penalizing both [+low] to [-low] and the reverse, unlike the tongue root faithfulness constraints. However, it would be equally possible to view this constraint as two separate ones, and in practice, only faithfulness to [+low] is necessary in the present analysis.

contexts, which is the way harmony operates in a number of other languages, such as Kinande, but not in Alur. The constraint in (16), weighted highly, ensures that [ʌ] does not exist in Alur. Note specifically that *ʌ is necessarily weighted higher than IDENT-IO[ATR] due to Richness of the Base, to ensure that hypothetical underlying ATR low vowel does not surface.

(16) *ʌ: assign a violation for any occurrence of an ATR low vowel in the output

For space reasons, I will omit constraints that are irrelevant to a given candidate.

The following constraints are unviolated in Alur, and must therefore be at the top of the weighting:

$*\underline{[\text{RTR}]}\infty[\text{ATR}]_{\text{Ft}}$

*ʌ

Additionally, IDENT-IO[ATR] is violated only to satisfy *ʌ. Other weightings will be discussed throughout the derivations.

3.4 Derivations

First, Table 7.5 shows an RTR verb root with an ATR V suffix. Candidate (a), the faithful candidate, violates the highly weighted harmony constraint $*\underline{[\text{RTR}]}\infty[\text{ATR}]_{\text{Ft}}$. Candidate (b), the winner, in which the root harmonizes, violates only low-weighted IDENT-IO[RTR]. Candidate (c), in which RTR harmony occurs to change the suffix, violates high-weighted IDENT-IO[ATR]. Candidate (d), in which the prefix also becomes ATR, has a second, fatal violation of IDENT-

IO[RTR]; there is no motivation for the prefix to change, because it occurs outside the harmony domain. This tableau demonstrates that weight(IDENT-IO[ATR]), weight(*[RTR]∞[ATR]$_{Ft}$)*scaling(trigger)*scaling(target) > weight(IDENT-IO[RTR]). I assign them weights of 12, 25, and 2 respectively. Since all triggers and targets behave the same way in Alur, I assign all scaling factors at 1.

		/ɛ́-nɛ̀nɔ̀-í/	*[RTR]∞ [ATR]$_{Ft}$ W=25	IDENT-IO[ATR] W=12	Ident-IO[RTR] W=2	H
	(a)	ɛ́-(nɛ̀n-í)	-1*1(ɛ)*1(i)*1(C)			-25
☞	(b)	ɛ́-(nèn-í)			-1	-2
	(c)	ɛ́-(nɛ̀n-ɪ́)		-1		-12
	(d)	é-(nèn-í)			-2	-4

Table 7.5: Suffix-to-root harmony in a verb root with a V suffix

Table 7.6 shows the same verb root with a CV suffix. In this case, the faithful candidate (a) wins, because the suffix is outside of the foot, and therefore outside of the harmony domain. Candidate (b), in which the root harmonizes to the suffix, violates IDENT-IO[RTR]. Notably, here is where it is crucial that the relevant harmony constraints are bounded to the foot; with the more general constraint *[RTR]∞[ATR], as shown in the tableau, a harmonized candidate would be preferred. The higher weight of IDENT-IO[RTR] than the general *[RTR]∞[ATR] is what ensures that harmony is bounded to the foot. Even though the constraint *[RTR]∞[ATR]$_{Ft}$ does not do anything in this tableau, I include it because it is relevant that the disharmonic candidate does not violate it.

		/έ-nɛ̀nɔ̀-wú/	*[RTR]∞[ATR]$_{\text{Ft}}$ W=25	IDENT-IO[RTR] W=2	*[RTR]∞[ATR] W=1	H
☞	(a)	έ-(nɛ̀nɔ̀)-wú			-3	-3
	(b)	έ-(nènò)-wú		-2	-1	-5

Table 7.6: Suffix-to-root harmony does not occur in a verb root with a CV suffix

Table 7.7 shows re-pairing of /a/ to [e] in a verb root with an ATR V suffix. Candidate (a), the faithful candidate, is ruled out by a violation of highly-weighted *[RTR]∞[ATR]$_{\text{Ft}}$. Candidate (b), the re-pairing candidate, violates IDENT-IO[low], but wins because this constraint is lower-weighted. Candidate (c), in which the vowel [ʌ] surfaces, is ruled out by *ʌ. This tableau adds the weight(*[RTR]∞[ATR]$_{\text{Ft}}$)*scaling(trigger)*scaling(target), weight(*ʌ) > weight(IDENT-IO[low]). The reason why /a/ re-pairs to [e], rather than [o], depends on the feature assumptions; here I assume that Alur specifies vowels for rounding but not backness, such that /a/ differs from [e] in only the features [ATR] and [low], while it differs from [o] in those plus [round].

		/ɛ́-tʃàk-í/	*ʌ W=25	*[RTR]∞ [ATR]$_{\text{Ft}}$ W=25	Ident-IO[ATR] W=12	Ident-IO[low] W=8	Ident-IO[RTR] W=2	H
	(a)	ɛ́-(tʃàk-í)		-1*1(a)*1(i)*1(C)				-25
☞	(b)	ɛ́-(tʃèk-í)				-1	-1	-10
	(c)	ɛ́-(tʃʌ̀k-í)	-1				-1	-27
	(d)	ɛ́-(tʃàk-í)			-1			-12

Table 7.7: Re-pairing of /a/ in verb roots

Table 7.8 shows root-to-suffix harmony from an ATR verb root to a V suffix. The faithful candidate (a) is ruled out by highly-weighted *[ATR]∞[RTR]$_{\text{Ft}}$; candidate (b) wins. This tableau demonstrates the weighting weight(*[ATR]∞[RTR]$_{\text{Ft}}$)*scaling(trigger)*scaling(target) > weight(Ident-IO[RTR]).

		/ɛ́-gùdò-ɛ́/	Ident-IO[ATR] W=12	*[ATR]∞ [RTR]$_{\text{Ft}}$ W=3	Ident-IO[RTR] W=2	H
	(a)	ɛ́-(gùd-ɛ́)		-1*1(u)*1(ɛ)*1(C)		-3
☞	(b)	ɛ́-(gùd-é)			-1	-2
	(c)	ɛ́-(gʊd-ɛ́)	-1			-12

Table 7.8: Root-to-suffix harmony to a V suffix from a verb root

Table 7.9 demonstrates that this harmony does not apply to suffix /a/. Candidate (a), the faithful candidate, wins despite violating *[ATR]∞$\underline{\text{[RTR]}}_{\text{Ft}}$. Candidate (b), the re-pairing candidate, violates high-weighted IDENT-IO[low], while candidate (c) violates high-weighted *ʌ. This tableau therefore illustrates the ranking weight(IDENT-IO[low]) > weight(*[ATR]∞$\underline{\text{[RTR]}}_{\text{Ft}}$)*scaling(trigger)*scaling(target). In particular, the difference between /a/ re-pairing in roots, as in Table 7.7, and not in suffixes here, comes from the fact that the regressive harmony constraint outweighs IDENT-IO[low], which in turn outweighs the progressive harmony constraint.

		/ɛ́-gùdò -á/	*ʌ W=25	IDENT-IO[ATR] W=12	IDENT-IO[low] W=8	*[ATR]∞ $\underline{\text{[RTR]}}_{\text{Ft}}$ W=3	IDENT-IO[RTR] W=2	H
☞	(a)	ɛ́-(gùd-á)				-1*1(u)*1(a)*1(C)		-3
	(b)	ɛ́-(gùd-é)			-1		-1	-10
	(c)	ɛ́-(gùd-ʌ́)	-1				-1	-27
	(d)	ɛ́-(gʊd-á)		-1				-12

Table 7.9: Lack of re-pairing of /a/ in suffixes

The case of RTR CV suffixes with ATR verb roots is identical to the case of ATR CV suffixes with RTR verb roots in Table 7.6: harmony will not occur because the CV suffix occurs outside of the domain of harmony. Moreover, the patterns with V suffixes for nouns will be identical to those with verbs in Table 7.5; noun roots are CVC, and verb roots become CVC in the presence of V suffixes due to hiatus resolution. What remains is to illustrate the patterns for CV suffixes with nouns, which differ from those with verbs; nouns with CV suffixes are of form CVC-CV, while verbs are of form CVCV-CV. Thus, while CV suffixes were not part of the foot with verbs, they are with nouns. Here is where the separation of constraints based on whether a single consonant or multiple consonants intervene becomes crucial, as it is the first occasion in which the root and suffix vowels are separated by multiple consonants.

Table 7.10 illustrates suffix-to-root harmony with an RTR noun root and an ATR CV suffix. The root and suffix are both within the foot, meaning that the faithful candidate violates *[RTR]∞[ATR]$_{\text{Ft}}$. Candidate (b), with harmony, wins despite violating IDENT-IO[RTR], since it is lower-weighted. This result requires weight(*[RTR]∞[ATR]$_{\text{Ft}}$)*scaling(trigger)*scaling(target)*scaling(CC) > weight(IDENT-IO[RTR]).

		/tʃɔ̀ng-wú/	*[RTR]∞ [ATR]$_{\text{Ft}}$ W=25	IDENT-IO[RTR] W=2	H
	(a)	(tʃɔ̀ng-wú)	-1*1(ɔ)*1(u)*0.5(CC)		-12.5
☞	(b)	(tʃòng-wú)		-1	-2

Table 7.10: Suffix-to-root harmony in noun roots with a CV suffix

Table 7.11 demonstrates the lack of root-to-suffix harmony with an ATR noun root and an RTR CV suffix. Candidate (a), the faithful candidate, wins despite violating *[ATR]∞$\underline{[RTR]}_{\mathrm{Ft}}$, which is low-weighted when the distance involves multiple intervening consonants. Crucially, this candidate does not violate *[ATR]∞$\underline{[RTR]}_{\mathrm{Ft}}$ for a single intervening consonant, as there are multiple consonants intervening between the vowels. Candidate (b), with harmony, fatally violates IDENT-IO[RTR]. Thus, we require weight(IDENT-IO[RTR]) > weight(*[ATR]∞$\underline{[RTR]}_{\mathrm{Ft}}$)*scaling(trigger)*scaling(target)*scaling(CC). Specifically, progressive harmony behaves differently depending on whether there is only a single consonant intervening or multiple consonants. Crucially, while progressive harmony across a single intervening consonant outweighs IDENT-IO[RTR], the more general progressive harmony across any number of consonants does not. In other words, weight(*[ATR]∞$\underline{[RTR]}_{\mathrm{Ft}}$)*scaling(trigger)*scaling(target)*scaling(C) > weight(IDENT-IO[RTR]) > weight(*[ATR]∞$\underline{[RTR]}_{\mathrm{Ft}}$)*scaling(trigger)*scaling(target)*scaling(CC) is the weighting that ensures progressive harmony occurs across one consonant, but not more than that.

		/lìmb-gí/	*[ATR]∞ $\underline{[RTR]}_{\mathrm{Ft}}$ W=3	IDENT-IO[RTR] W=2	H
☞	(a)	(lìmb-gí)	-1*1(i)*1(ɪ)*0.5(CC)		-1.5
	(b)	(lìmb-gí)		-1	-2

Table 7.11: Lack of root-to-suffix harmony with a noun root and CV suffix

The full weighting requirements needed for dealing with Alur harmony are given in (17).

(17) Weighting for Alur harmony

$*\underline{[\mathrm{RTR}]}\infty[\mathrm{ATR}]_{\mathrm{Ft}}$, *ʌ >>

IDENT-IO[ATR] >>

IDENT-IO[low] >>

$*[\mathrm{ATR}]\infty\underline{[\mathrm{RTR}]}_{\mathrm{Ft}}$ * scaling (C) >>

IDENT-IO[RTR] >>

$*[\mathrm{ATR}]\infty\underline{[\mathrm{RTR}]}_{\mathrm{Ft}}$ * scaling (CC)

These weightings can capture all of the complex behaviour seen in Alur harmony, for both nouns and verbs.

4 Discussion and conclusion

Alur has a complex harmony system, described by Kutsch Lojenga (1991) as being distinct for nouns versus verbs, for V suffixes versus CV suffixes, and for [a] versus other vowels. In this paper, I have analyzed the first property as being due to the difference in canonical structure of nouns versus verbs (CVC versus CVCV), the second property as being due to a combination of whether the suffix is integrated into the foot domain and whether multiple consonants intervene between the root and suffix vowels, and the third property as due to whether changes in an additional feature (here, [low]) are required to make a vowel [ATR]. This analysis allows all aspects of the Alur pattern to be expressed through purely phonological motivations; the phonology does not have to refer to nouns versus verbs, or to V versus CV suffixes except in how suffix structure affects the distance between the vowels (in terms of the number of intervening consonants). As mentioned, alternatively the V versus CV suffixes may relate to which vowels are targets of progressive ATR spreading, since the harmonizing V suffix is mid, while the non-harmonizing CV one is high.

Regressive and progressive harmony behave differently in Alur, with regressive harmony applying in more situations than progressive harmony does, which is a pattern that has been noted in a wide variety of languages (e.g. Hyman 2002). This fact is captured in the analysis by separating out progressive and regressive harmony constraints.

Within progressive harmony, behaviour is different depending on whether there is only a single consonant intervening or multiple consonants, with the latter blocking progressive harmony. This pattern has also been noted for the related language Lango (see Chapter 8), and has been analyzed as harmony being blocked by an intervening mora (e.g. Archangeli & Pulleyblank 1994). In the present analysis, I have expressed this distance effect using two harmony constraints, one applicable with a single intervening consonant, and the other applicable with any number of intervening consonants. Alur instantiates both types of patterns predicted by this choice: regressively, harmony applies regardless of the number of intervening consonants, while progressively, it applies only across a single consonant. This choice of using the number of consonants (one versus any number) rather than the number of moras (zero versus any number) in the constraints is not crucial, in particular because given that Alur harmony is limited to a syllable binary foot, so the only possible intervening moras are coda consonants. Note that, because regressive harmony does cross coda consonants, the foot domain must crucially be binary based on syllables, not moras; as such, the explanation for why progressive harmony fails to affect CV suffixes cannot fall to the domain. Note that a number of harmony systems are known to be constrained by foot domains, as in Lezgian (Haspelmath 1995 on vowel harmony;

Ozburn & Kochetov 2018 on consonant harmony), Kera (Pearce 2006), and Assamese (Mahanta 2008).

Further research into Alur should examine three aspects of the present analysis. First, the analysis of noun patterns required NC sequences to be prenasalized stops, rather than actual clusters. This analysis was what allowed consistent patterning with respect to how the number of consonants (or moras) affects whether harmony occurs. Specifically, if NC represents two separate consonants, then some nouns are CVCC, while others are CVC, and so it becomes more difficult to distinguish the behaviour of V versus CV suffixes in progressive harmony with nouns based solely on distance.[71] However, it is unclear whether there is any additional phonological or phonetic evidence in Alur supporting NC sequences being prenasalized stops; this should be examined further. Second, the domain of harmony was established to be the foot, which is binary and aligned to the left edge of the root. A similar proposal has been argued for Lango (Pulleyblank 2001), but the existence of foot structure has not been established in Alur. Research should look for further evidence from tone or stress for a binary foot at the left edge of the root. Third, for it to be sensible to have the number of intervening consonants affect whether harmony occurs, it would make sense for those consonants to be moraic; again this has been proposed for Lango (e.g. Archangeli & Pulleyblank 2001), but has not been established for Alur; additional evidence about moraicity would be useful to examine.

As shown above, the pattern in Alur can fit within the type of analysis proposed in this dissertation. Alur treats all vowels equally in terms of targets, and there is no reason to assume that the behaviour of [a] comes from anything besides its status in the inventory. Yet, because treating all vowels equally is an option in the analysis developed through this dissertation, Alur

---

[71] This is not necessarily impossible, though, if the relevant unit is a mora rather than just a consonant; it would depend in that case on the moraic/syllabic analysis.

does fit within this framework. Specifically, in Alur, the trigger and target scaling factors are all the same regardless of vowels; the distance scaling can be set such that a distinction is made between one consonant and more than one. Thus, in this analysis, neutrality is in fact arising from the same factors that shape the inventory.

As such, it is worth noting that the neutrality of /a/ in Alur harmony could be formulated in traditional terms, relying on the context-free markedness constraint *ʌ, ranked above IDENT-IO[low] and IDENT-IO[ATR]. This ranking simultaneously gives both the re-pairing/neutrality behaviour under harmony (depending on the ranking with the harmony constraint) as well as the general absence of an ATR low vowel in the inventory. Thus, in this analysis, the neutrality of /a/ does arise from the same factors that shape the inventory. This is exactly the type of analysis that does not work for Hungarian and the other languages discussed in this dissertation. However, such an analysis would be fully consistent with the Alur facts. Tables 7.12 and 7.13 repeat Tables 7.8 and 7.9 above, except using a standard Optimality Theory account (with separate constraints for harmony across a single consonant versus multiple consonants). Evidently, the same result is predicted.

		/ɛ́-gùdò-ɛ́/	IDENT-IO[ATR]	*[ATR]C [RTR]$_{\text{Ft}}$	IDENT-IO[RTR]
	(a)	ɛ́-(gùd-ɛ́)		*!	
☞	(b)	ɛ́-(gùd-é)			*
	(c)	ɛ́-(gʊd-ɛ́)	*!		

Table 7.12: Root-to-suffix harmony to a V suffix from a verb root

		/ɛ́-gùdò -á/	*ʌ	IDENT-IO[ATR]	IDENT-IO[low]	*[ATR]C [RTR]$_{\text{Ft}}$	IDENT-IO[RTR]
☞	(a)	ɛ́-(gùd-á)				*	
	(b)	ɛ́-(gùd-é)			*!		*
	(c)	ɛ́-(gùd-ʌ́)	*!				*
	(d)	ɛ́-(gʊd-á)		*!			

Table 7.13: Lack of re-pairing of /a/ in suffixes

As such, the approach used in the other chapters is not necessary for Alur, and there is no reason to think that learners adopt different trigger and target scaling depending on vowel. Nonetheless, as shown in the preceding sections, the approach in the other chapters is still consistent with the Alur data. This is crucial, because we do not want to lose the ability to account for cases in which there is a match between inventory structure and participation in harmony, like in Alur progressive harmony. Moreover, having the option of making /a/ less good as a harmony target allows for the prediction of languages similar to Alur in which low vowels do contrast in some but not all environments and /a/ is neutral. This does occur in other Western Nilotic languages, such as Mayak (Chapter 9).

Chapter 8: Simplifying the analysis of Lango ATR harmony

1 Introduction

Lango (Western Nilotic; Uganda) has a complex ATR harmony system, in which harmony depends on the interaction of directionality (progressive vs. regressive), trigger vowel quality (high front vs. high back vs. non-high), target vowel quality (high vs. non-high), and distance (across a single consonant vs. across multiple consonants). This complicated pattern has been the subject of a number of analyses, including that of Archangeli & Pulleyblank (1994), who adopt multiple rules based on grounded conditions; that of Smolensky (2006), with conjoined constraints; and that of Potts et al. (2010), who use conditions on heads of harmony domains and cumulative constraint interaction (gang effects) in Harmonic Grammar.

As Potts et al. (2010) note, analyses of Lango previous to theirs have suffered from loss of generality and/or typological over-generation. The Potts et al. (2010) analysis itself requires substantial representational complexity. Specifically, it requires identifying certain vowels as heads, with the stipulation that heads must be faithful and only exist in ATR domains that are minimally binary. This particular version of heads therefore requires all head-related constraints to see both the input, to know which vowel was faithful to ATR, and the output, to know whether the domain is minimally binary. As such, it cannot be accounted for without 'two-level' constraints that are a combination of both markedness and faithfulness. Kaplan (2008) provides an analysis of the non-iterativity of Lango harmony, but not the complex asymmetries related to triggers, targets, and distance.

In this paper, I propose a novel way of analyzing these facts about Lango harmony, from an interaction of simpler, more generalizable constraints and principles. Specifically, only basic faithfulness and harmony constraints are required, but the penalty for disharmony depends on the specific vowels interacting and their distance apart. This analysis builds on the one for front/back harmony in Chapters 3 and 4, adapting it for an ATR system. I argue that this approach maintains the restrictive nature of the Potts et al. analysis, while also avoiding the problems inherent in the use of heads to account for the Lango pattern.

2 The facts on Lango harmony

This section presents the Lango harmony pattern in detail. All facts on Lango come from Woock & Noonan (1979). Like Smolensky (2006) and Potts et al. (2010), the present paper covers the same set of facts as Archangeli & Pulleyblank (1994), based on Woock & Noonan (1979).[72]

The vowel inventory of Lango is shown in Table 8.1. There is a full set of ten vowels, each paired for tongue root.

	ATR			RTR		
	Front	Central	Back	Front	Central	Back
High	i		u	ɪ		ʊ
Mid	e		o	ɛ		ɔ
Low		ə			a	

Table 8.1: Vowel inventory of Lango

Lango primarily shows ATR-dominant harmony, in which ATR vowels control harmony regardless of whether they occur in the root or in the suffix. There is additionally a single case of

---

[72] It is worth noting, as Potts et al. (2010) do, that other descriptions of Lango harmony differ. For example, while Woock & Noonan (1979) have examples of RTR spreading and describe certain instances of harmony as being blocked by consonant clusters, Noonan (1992) describes harmony as being only ATR spreading and claims that it is suffix-initial consonants that block harmony, as opposed to clusters. Like Potts et al. (2010), I use the Woock & Noonan (1979) facts, as characterized by Archangeli & Pulleyblank (1994), to make the present analysis comparable to previous ones. However, I consider the facts from Noonan (1992) in the discussion in Section 6.

RTR harmony, ignored in many previous analyses, that will be discussed briefly below. Focusing first on Lango ATR harmony, it can be divided into four sub-patterns, depending on directionality and distance. The conditions are given as a summary in Table 8.2, and will be discussed through the next subsections.

	Across 1 C	Across 2 Cs
Progressive	Always	High trigger
Regressive	High trigger	Trigger [i], or trigger [u] + high target

Table 8.2: Conditions on Lango ATR harmony

2.1 Progressive ATR harmony

The strongest type of harmony in Lango is progressive and across a single consonant; in this case, harmony can be triggered by any vowel and target any vowel. An example occurs with the inalienable version of the 1SG suffix, which alternates between [-á] and [-ə́] depending on the ATR value of the root. As shown in (1), this suffix consistently harmonizes regardless of which vowel is in the root; (a-d) have ATR root vowels and the ATR suffix, while (e-h) have RTR root vowels and the RTR suffix. As noted above, all examples in this paper are taken from Woock & Noonan (1979).

(1)

ATR root vowels take ATR suffix -ə́	(a)	cíŋ	‘hand’	cíŋə́	‘my hand’
	(b)	ŋèt	‘side’	ŋètə́	‘my side’
	(c)	wót	‘son’	wódə́	‘my son’
	(d)	ŋùt	‘neck’	ŋùtə́	‘my neck’
RTR root vowels take RTR suffix -á	(e)	yíb	‘tail’	yíbá	‘my tail’
	(f)	lɛ́b	‘tongue’	lɛ́bá	‘my tongue’
	(g)	wàŋ	‘eye’	wàŋá	‘my eye’
	(h)	bwɔ́m	‘wing’	bwɔ́má	‘my wing’

The suffix marking alienable possession in Lango also has an alternation between [-á] and [-ə́], but it also triggers gemination of the root-final consonant. As such, the potential trigger and target of harmony are at a greater distance apart: two consonants, one of which is a coda. In this case, the ATR version of the suffix [-ə́] occurs only if the root-final vowel is both ATR and high, as in (2a-b). ATR root vowels that are non-high, as in (2c-e), do not trigger harmony in this situation, and the suffix vowel surfaces as RTR [-á]. The [-á] alternant of course also occurs after RTR root vowels, as in (2f-j).

(2)

High ATR root vowels take ATR suffix -ə́	(a)	píg	‘juice’	píggə́	‘my juice’
	(b)	òpúk	‘cat’	òpúkkə́	‘my cat’
Non-high ATR root vowels take RTR suffix -á	(c)	gwèn	‘chickens’	gwènná	‘my chickens’
	(d)	ɲə̀ŋ	‘crocodile’	ɲə̀ŋŋá	‘my crocodile’
	(e)	dòk	‘cattle’	dòkká	‘my cattle’
RTR root vowels take RTR suffix -á	(f)	àtín	‘child’	àtínná	‘my child’
	(g)	bɛ̀l	‘wheat’	bɛ̀llá	‘my wheat’
	(h)	càl	‘picture’	càllá	‘my picture’
	(i)	kɔ̀m	‘chair’	kɔ̀mmá	‘my chair’
	(j)	lʊ̀t	‘stick’	lʊ̀ttá	‘my stick’

From these sets of examples, we can extract two of the crucial factors in the application of ATR harmony in Lango: number of intervening consonants, and height of the trigger. Specifically, harmony is more likely across a single consonant than across two, and more likely with a high trigger than with a non-high one. These factors interact, such that progressive harmony across two consonants occurs only with a high trigger. Section 3 will discuss factors that make high vowels better triggers of this type of harmony.

Worth noting now, and explored further in Section 6, is that geminates are relevant to whether harmony applies; Noonan (1992) notes that in some situations, harmony applies across multiple consonants only in the case of geminates. Notably, other possessive CV suffixes like

/-gɪ/ and /-wa/ never undergo harmony, even from high vowels (Noonan 1992). These facts will be left for discussion in Section 6.

2.2 Regressive ATR harmony

Turning now to regressive harmony, we again see a distinction between high and non-high triggers, and between harmony across a single consonant versus two. Additionally, there is a distinction in regressive harmony between high front and high back triggers, and between high and non-high targets.

Across a single consonant when the trigger is high (front or back), harmony occurs, regardless of the target. Some examples are shown in (3); the RTR vowels in the unsuffixed roots are ATR in the suffixed form.

(3)

RTR root vowels become ATR with [ú] in suffix	(a)	pí	‘for’	pìwú	‘for you (pl.)’
	(b)	lɛ̀	‘axe’	lèwú	‘your (pl.) axe’
	(c)	jɔ̀	‘people’	jòwú	‘your (pl.) people’

Regressive harmony also occurs across single as well as multiple consonants to any vowel if the trigger is high front [i], as shown in (4); the RTR vowels in the unsuffixed roots in (f-j) are ATR in the suffixed forms, while the ATR root vowels in (a-e) do not alternate. This pattern is consistent with the progressive facts discussed earlier. Note that the lack of gemination for (b) and (f) is due to alienable versus inalienable possession. Additionally, the alienable possessive suffix is /-ní/, but the /n/ assimilates after consonant-final stems, creating the gemination.

(4)

ATR root vowels remain ATR with í in suffix	(a)	píg	‘juice’	píggí	‘your (sg.) juice’
	(b)	ém	‘thigh’	émí	‘your (sg.) thigh’
	(c)	ɲə̀ŋ	‘crocodile’	ɲə̀ŋŋí	‘your (sg.) crocodile’
	(d)	dòk	‘cattle’	dòkkí	‘your (sg.) cattle’
	(e)	búk	‘book’	búkkí	‘your (sg.) book’
RTR root vowels become ATR with í in suffix	(f)	yíb	‘tail’	yíbí	‘your (sg.) tail’
	(g)	dɛ̀k	‘stew’	dèkkí	‘your (sg.) stew’
	(h)	màc	‘fire’	mə̀ccí	‘your (sg.) fire’
	(i)	kɔ̀m	‘chair’	kòmmí	‘your (sg.) chair’
	(j)	lʊ̀t	‘stick’	lùttí	‘your (sg.) stick’

In contrast, if the trigger is high back [u] and separated from the target by multiple consonants (including if the suffix is [-wu], in which the consonant does not assimilate), some speakers apply harmony only if the target is also high. Some examples are shown in (5). Note how in examples (a-b), there is harmony, but there is no harmony in (c-e); this is in contrast to what occurs with the same root vowels (and, in some cases, exactly the same root) in the examples in (4) when [i] is the trigger. Note further that while the generalization stated in Woock & Noonan (1979), and analyzed in many subsequent papers, is about high front versus high back triggers, Noonan (1992) suggests that in fact the distinction is about -V versus -CV suffixes and intervening geminate versus non-geminate clusters. I adopt the trigger distinction here for comparability with previous analyses, but discuss the intervening material facts in Section 6.

(5)

High RTR root vowels become ATR with suffix -wú	(a)	nìŋ	‘name’	nìŋwú	‘your (pl.) name’
	(b)	lʊ̀t	‘stick’	lùtwú	‘your (pl.) stick’
Non-high RTR vowels remain RTR with suffix -wú	(c)	dɛ̀k	‘stew’	dɛ̀kwú (some speakers); dèkwú (most speakers; Noonan 1992)	‘your (pl.) stew’
	(d)	kàl	‘millet’	kàlwú	‘your (pl.) millet’
	(e)	kɔ̀m	‘chair’	kɔ̀mwú	‘your (pl.) chair’

As before, we can analyze the distinction between harmony across a single consonant versus across more than one as a distance effect. Additionally, this data suggests a need to further separate preferred triggers, in that harmony is more likely with [i] than with [u], and so [i] is a better trigger. We also observe in (5) that harmony is more likely with high targets than with non-high targets, suggesting that high vowels are, in at least this context, better targets of harmony than non-high vowels. Finally, these examples are indicative of a directionality effect: [u] triggers progressive but not regressive harmony across multiple consonants to non-high vowels (see [òpúkkə́] 'my cat' in (2b) versus [kàlwú] 'your millet' in (5d)), meaning that progressive harmony in Lango is stronger than regressive. This directionality distinction is the opposite of what is seen in many languages (see e.g. Hyman 2002 on regressive bias), and will be discussed further in Section 3, along with these further trigger and target asymmetries.

The fact that progressive harmony is stronger than regressive can additionally be seen by the fact that non-high vowels do not trigger regressive harmony in Lango, as shown in (6); the last of these examples shows that this absence of harmony holds even across a single consonant. Like with progressive harmony, the fact that high vowels trigger harmony in cases where non-high ones do not can be analyzed as high vowels being preferred triggers. Note that while ATR mid vowel suffixes do sometimes undergo RTR harmony according to Woock & Noonan (1979), as will be discussed in Section 2.4, they fail to trigger ATR harmony.[73] Note that examples (6a-d) involve ATR roots, while examples (6e-g) involve RTR roots.

---

[73] As noted, the data discussed here is based on Woock & Noonan (1979). Noonan (1992) describes the ventive suffix /-o/ as triggering regressive harmony across geminates, and the plural suffix /-e/ as triggering harmony; according to this, the RTR harmony discussed in Section 2.4 is unusual. This later data on mid vowels is discussed in Section 6.

(6)

(a)	riŋŋo	‘to run’
(b)	ketto	‘to put’
(c)	pwoddo	‘to beat’
(d)	rucco	‘to entangle’
(e)	lɪmmo	‘to visit’
(f)	nɛnno	‘to see’
(g)	wayo	‘to pull’

2.3 Summary of ATR harmony

The conditions on Lango harmony can be summarized as in Table 8.3, repeated from Table 8.2 above. Progressive harmony across a single consonant applies regardless of the vowels involved; other contexts require the trigger to be high, while regressive harmony across two consonants also requires either the trigger to be front or the target to be high.

	Across 1 C	Across 2 Cs
Progressive	Always	High trigger
Regressive	High trigger	Trigger [i], or trigger [u] + high target[74]

Table 8.3: Conditions on Lango ATR harmony (according to Woock & Noonan 1979; see Section 6 on the characterization by Noonan 1992)

---

[74] As noted, it could instead be a high trigger with conditions on intervening consonants, instead of high back versus high front. Additionally, according to Noonan (1992), the trigger could also be a mid vowel in these cases, depending on morphology and suffix form. As the data from Noonan (1992) differ from the description by Woock & Noonan (1979), they are discussed in Section 6.

From these facts, we can conclude that there are four interacting factors that affect when ATR harmony occurs. First, harmony is more likely across one consonant than two. Second, harmony is more likely progressively (i.e. in the sequence ATR…RTR) than regressively (i.e. RTR…ATR). Third, harmony is affected by the 'strength' of the trigger, in that it is more likely with a high front trigger than a high back trigger, and with any high trigger than with a non-high trigger. Finally, harmony is more likely with high than non-high targets. All of these factors must be incorporated into the analysis.

2.4 Progressive RTR harmony

Before concluding this section, I briefly discuss the one instance of RTR harmony in Lango. This harmony is triggered only by back RTR stems, and apparently targets suffix /-o/, as shown comparing examples with this suffix in (7) versus (8) below. These examples contain the infinitive suffix, but a similar pattern is seen with the transitive object suffix in Woock & Noonan (1979). However, with the latter suffix, some roots with front RTR vowels also take the RTR version of the suffix; the choice of [-o] or [-ɔ] after roots with front RTR vowels is lexically determined. Most existing analyses, with the exception of Archangeli & Pulleyblank (1994), ignore RTR harmony as exceptional, and it is beyond the scope of the present paper because of problems in the interaction of this harmony with ATR-dominant harmony, but it is important to consider in future work. In particular, the conditions on this type of harmony in Lango are not particularly well understood, and should be investigated further.

(7)

(a)	riɲɲo	‘to run’
(b)	ketto	‘to put’
(c)	pwoddo	‘to beat’
(d)	rucco	‘to entangle’
(e)	lɪmmo	‘to visit’
(f)	nɛnno	‘to see’
(g)	wayo	‘to pull’

(8)

(a)	lwɔkkɔ	‘to wash’
(b)	lʊbbɔ	‘to follow’

3 Explaining the patterns

In Section 2, we saw that ATR harmony in Lango is affected by a combination of triggers, targets, distance, and directionality. This section will explore those generalizations in more detail, justifying and explaining them within the context of cross-linguistic and phonetic evidence.

3.1 Triggers

In Lango, harmony is more likely with a high front trigger than a high back one, and with any high trigger than a non-high one. This is evident particularly in regressive harmony across two consonants, in which high front vowels trigger harmony to any vowel, high back vowels trigger harmony only to high vowels, and non-high vowels do not trigger harmony at all.

This type of patterning is common cross-linguistically in ATR-dominant harmony. As noted by Casali (2003; 2008), there is a strong correlation between languages having an ATR contrast in high vowels and the presence of ATR-dominant harmony. Surveys by both Casali and Rose (2018) show that virtually all languages with RTR high vowels [ɪ] and [ʊ][75] (in addition to ATR ones [i] and [u]) have harmony, and these systems are typically ATR-dominant. This is true regardless of what other contrasts exist in the language; in cases like Kinande, for instance, the only contrastive ATR vowels are high (see e.g. Archangeli & Pulleyblank 2002 and references therein). Such surveys also indicate that languages without an ATR contrast in high vowels (but with such a contrast in mid vowels) may have harmony, but it is almost never ATR-dominant; an example is Yoruba, in which high vowels do not contrast, but mid vowels do, and harmony is RTR-dominant (Archangeli & Pulleyblank 1994). As such, there is a strong correlation between just the presence of this high vowel contrast and the existence of ATR-dominant harmony.

Unsurprisingly, then, in languages with ATR-dominant harmony, high vowels appear to consistently act as triggers; other vowels may also be triggers, but do not necessarily need to be. For example, in Mayak, another Western Nilotic language (to be discussed in Chapter 9), low vowels are contrastive for ATR versus RTR, but the ATR low vowel does not trigger ATR harmony in most contexts, while high vowels do (Andersen 1999). The Lango pattern, in which high vowels are preferred triggers of ATR harmony, is cross-linguistically quite robust.

Rose (2018) proposes that this tendency for high vowels to trigger ATR harmony may relate to perceptual difficulties in distinguishing ATR from RTR among high vowels.[76] The idea that harmony might be triggered preferentially by perceptually impoverished triggers is common,

---

[75] This is presumably only true if contrasts like [i,u] versus [ɪ,ʊ] in languages like English (and Germanic languages in general) do not involve tongue root, but rather laxing.

[76] This may be a problem for high vowels being optimal targets, but this will be discussed in the next subsection.

and has been proposed originally by Suomi (1983) for front/back harmony in languages like Finnish; it was also adopted by Kaun (1995; 2004) for rounding harmony, as well as by Kimper (2011) for harmony more broadly. As noted in surveys by Becker-Kristal (2010) and Starwalt (2008) on a number of languages with various types of tongue root contrasts, the difference in F1 between ATR versus RTR high vowels tends to be smaller than that between ATR versus RTR mid vowels. While Rose (2018) notes that other factors such as voice quality may help to distinguish the contrast in high vowels, she suggests that these smaller acoustic differences may translate into perceptual difficulties that result in the strong correlation between ATR contrast in high vowels and ATR-dominant harmony.

Some evidence for this claim comes from a perceptual study by Fulop et al. (1998) with speakers of Degema, which has ten contrastive vowels paired for ATR. They found that speakers could only reliably distinguish ATR contrasts in mid vowels; these are well-separated in acoustic space, whereas the other vowels show a great deal of overlap such that even native speakers could not reliably distinguish them. This result suggests that there is not a substantial perceptual (or acoustic) contrast between ATR versus RTR high vowels, which potentially motivates their behaviour as preferred triggers of harmony, given the tendency for perceptually impoverished contrasts to trigger harmony. However, it is important to note that this study used manipulated stimuli, and found that listeners reliably use F1, rather than F2, to distinguish vowels; it is unclear whether they might use other cues which were absent from the manipulated signal.

Nonetheless, this perceptual-motivation hypothesis suffers from some problems, including lack of support from other experiments on other languages. Rose et al. (2019) report on a perceptual study conducted with speakers of Akan, using stimuli produced by a bilingual Akan/Gua speaker; both languages have ATR-dominant harmony and an ATR contrast among high and mid vowels. Participants were asked whether pairs of stimuli were the same or different, and they had over ninety percent accuracy on all pairs that differed only in ATR. In particular, participants were very good at discriminating the ATR contrast among high vowels, and not worse at this discrimination than they were with mid vowels. This result could be explained by phonetic facts of vowel productions in Akan and Gua, in which mid ATR vowels have lower F1 than high RTR vowels (i.e. mid ATR vowels are actually 'higher' than high RTR vowels in terms of the acoustic vowel space), meaning that the tongue root pairs are well separated in acoustic space. However, it raises questions about the idea that perceptual difficulties with the ATR contrast in high vowels is a motivating factor for harmony.

An additional potential problem for this view is the fact that ATR contrasts are also reportedly difficult to perceive among low vowels, given the many languages in which they have been overlooked, yet low vowels are typically the worst possible triggers in ATR-dominant harmony. In the Degema study described above, speakers also had difficulty with the contrast in low vowels. Additional anecdotal evidence in the literature supports this conclusion; for instance, Otuho speakers report that the ATR contrast in low vowels is 'not important', despite being contrastive; this is suggestive of more difficulty differentiating the low vowels, although it could be because of factors like low functional load (Coates 1985).[77] Other languages in which allophonic or even contrastive ATR low vowels have been missed by researchers and native speakers include Kinande (allophonic; Gick et al. 2006) and Dagaare (contrastive; Ozburn et al. 2018). In the case of allophonic alternations, it is possible that researchers may have heard the alternations but did not consider them important; in the case of contrastive vowels, it could be that there are few examples of the contrast. Starwalt (2008) suggests that F1 measurements for

---

[77] It is not clear whether speakers tend to be literate in the language, though it is worth noting that the standardized orthography does not distinguish the low vowels, although apparently that orthographic decision was made because it was felt that the low vowel distinction was not important (Coates 1985).

low ATR pairs are close in some language and well-separated in others, and that Nilotic 10 vowel systems like Bari tend to show wide F1 separation in low vowels. As such, it is not consistent by language whether we expect low ATR distinctions to be difficult to perceive; however, the same is true of high vowels, with some languages like Akan showing much wider separation of high ATR contrast than other languages (see e.g. Rose et al. 2019). Noonan (1992) provides average formant measurements for Lango vowels, which showed that the low ATR/RTR pair and the front ATR/RTR pairs are well-separated in F1, whereas the back (non-low) ones are not; this distinction does not appear to pattern with the way these vowel pairs behave in the harmony systems.

It is therefore unlikely that perception forms the entire basis for high vowels triggering harmony; articulation may also play a role, in terms of the degree of tongue root differences in high versus non-high vowels, or perception may be relevant in ways yet to be tested. What is clear, however, is that high vowels are the best triggers across languages with this type of harmony; Lango is not at all unusual in this regard.

The fact that high front vowels are preferred triggers over high back vowels is less thoroughly typologically and phonetically supported, but still evident. Archangeli & Pulleyblank propose a condition ATR/Bk that says ‘if [+ATR], then [-back]’, which they use to capture asymmetries between high front and high back vowels in both Lango and Kinande. It is reasonable to expect that a fronted position of the tongue body would have additional effects on tongue root advancement that would make front high vowels better triggers of ATR harmony than back high vowels, though the exact mechanisms of this distinction remain to be determined. Crucially, though, the front versus back asymmetry is not unique to Lango.

3.2 Targets

In the case of regressive harmony across two consonants with the trigger [u], we see a distinction between potential targets in Lango, where potential targets that are high undergo harmony but non-high ones do not. Such a pattern is in fact attested in other ATR-dominant harmony systems, with several languages showing harmony that only targets high vowels in one direction. This is true of progressive harmony in Kinande (e.g. Archangeli & Pulleyblank 2002) and Mayak (Andersen 1999; see Chapter 9); in both these languages, regressive harmony targets both high and mid vowels, and in Kinande also the low vowel. This generalization about high versus non-high targets can be conceptualized in two possible ways: either high vowels are better targets of ATR harmony independent of other factors, or else targets that agree in height with the trigger are better. For the purposes of the analysis in Section 4, either of these views could correctly predict the Lango pattern. However, I will adopt the former account here, in part because it substantially reduces the complexity of the analysis. Evidence on whether high vowels are ideal targets of ATR harmony is somewhat contradictory; given that neither view has clear evidence for its accuracy, I consider both of them here.

The idea that target preferences in Lango are a case of height agreement would make it an instance of parasitic harmony, which is a well-established pattern, though not in ATR harmony. Rounding harmony is quite often parasitic on height, as noted by Kaun (1995; 2004), and it has been proposed that this property is motivated by the fact that rounding is articulatorily implemented differently in vowels of different heights. Consonant harmony is also often parasitic on agreement in non-harmony features, such as on manner and/or voicing in sibilant harmony (Hansson 2001; Rose & Walker 2004). Given the relationship between tongue root position and tongue height, it would be reasonable to hypothesize that the precise mechanisms for realizing ATR differ somewhat between high and non-high vowels, motivating parasitic harmony. As such, harmony being parasitic on [+high] could potentially make sense as an explanation for why high vowels are often targets of ATR-dominant harmony triggered by high vowels.

In contrast, at first glance, viewing high vowels as ideal targets of ATR harmony, independent of the trigger, is potentially in conflict with a number of proposals about the motivations for harmony more broadly and certain properties of ATR harmony in particular. Specifically, proposing that high vowels are both ideal triggers and targets is contrary to perceptual arguments about trigger and target preferences, and high vowels being ideal targets fails to explain why almost all languages with ATR-dominant harmony without an ATR contrast in the mid vowels produce allophonic ATR mid vowels through harmony.

As noted for rounding harmony by Kaun (1995; 2004) and considered for ATR harmony patterns by Rose (2018), targets of harmony tend to be those with a strong perceptual contrast for the feature. In rounding harmony, preferred targets tend to be high, which Kaun (1995; 2004) argues is due to rounding being more perceptible on high vowels, meaning these vowels are better for extending the perceptibility of the feature. A similar pattern can be seen in metaphony (e.g. Walker 2005), in which the target is the stressed vowel, a strong position in which contrasts are more perceptible. However, as described in detail above and noted by Rose (2018), high vowels are specifically the ones with a weak perceptual contrast in ATR harmony; it is the mid vowels that are typically well-separated acoustically and perceptually. Thus, while they are predicted to be ideal triggers based on perception, they should not be ideal targets.

Considering the situation of mid vowels in more detail, it is worth noting that Rose (2018: 9) says that languages with an ATR contrast in high but not mid vowels 'consistently' create ATR mid vowels allophonically through harmony; this situation occurs, for example, in Kinande (Archangeli & Pulleyblank 2002) and Mayak (Andersen 1999; see Chapter 9), and has been identified as a characteristic of ATR dominant languages (Casali 2003; 2008; Rose 2018). In contrast, it is quite rare for languages with the contrast in mid but not high vowels to create allophonic RTR high vowels; it does occur in a few cases, including in certain dialects of Yoruba (Przezdziecki 2005), but Rose (2018) identifies only 3 cases out of the 82 languages in her database that have the relevant inventory.[78] This distinction in behaviour of non-contrastive vowels is suggestive of the fact that mid vowels are ideal targets of ATR harmony; they harmonize even at the expense of creating vowels that do not otherwise exist in the language. High vowels, on the other hand, typically harmonize only when contrastive. In other words, harmony targeting high vowels is typically always structure-preserving.

However, looking at ATR dominant harmony more carefully, it is worth noting that there are a number of languages that behave like Lango, in which there are instances of only high vowels harmonizing, yet to my knowledge, no cases in which mid vowels harmonize to the exclusion of high vowels. The latter type of pattern occurs regularly in other types of tongue root harmony, like in Yoruba (RTR dominant), but not in ATR dominant harmony. This asymmetry is precisely the opposite of what we would expect if the motivation for harmony is to extend the feature to positions where it is perceptually most salient, and if mid vowels are such a position. Instead, we would expect ATR dominant harmony to be like rounding harmony, in which we do see the ideal perceptual targets behaving as targets to the exclusion of other vowels in some languages (Kaun 1995; 2004). Specifically, we would expect cases of ATR dominant harmony in which only mid vowels are targeted. Thus, there remains the possibility that high vowels are in

---

[78] Note that 'lack of ATR contrast' means different things in mid versus high vowels, in terms of which vowels are in the inventory. For mid vowels, a lack of contrast means that only RTR ones are in the contrastive inventory, whereas for high vowels, a lack of contrast means that only ATR ones are present.

fact the best targets of ATR dominant harmony, as this appears to be most consistent with the typological facts.

Despite the perceptual properties that might make them non-ideal targets of ATR harmony, high vowels are generally accepted to be the ideal ATR vowels, as well as better as ATR than as RTR. Based on the inventory structures described in Casali (2016) for languages with tongue root contrasts, in languages with only a single set of high vowels, those high vowels are ATR, whereas languages with only one set of mid vowels typically have RTR mid vowels, and those with only a single low vowel consistently have an RTR one. It is known that ATR high vowels are easier to produce than RTR ones, since tongue root advancement also necessarily involves tongue body raising (Hall & Hall 1980; Archangeli & Pulleyblank 1994). To encode this generalization, Archangeli & Pulleyblank (1994: 176) propose a condition ATR/HI, which states "If [+ATR], then [+high]". This implication also goes the other way; Archangeli & Pulleyblank (1994: 174) also propose a HI/ATR condition, "If [+high], then [+ATR]".

Given these facts, we could argue that high vowels are ideal targets of ATR harmony: they are the best possible ATR vowels. This issue reduces to what we consider the point of harmony to be, in terms of what kinds of phonetic and markedness pressures drive harmony. If the reason harmony occurs is to extend articulatory gestures or, in particular, to make a feature more perceptible, then whether a vowel is 'good' at being ATR is not necessarily correlated to whether it is a good target of ATR harmony; good targets are generally thought to be those vowels on which the harmony feature is more perceptible, or which are particularly susceptible to coarticulation due to articulatory factors. As discussed above, high vowels are not the best according to these criteria. Nonetheless, a vowel that is 'bad' at being ATR might have some degree of resistance to becoming ATR as the result of any phonological processes that create

ATR vowels, including harmony. In this way, ‘goodness’ of the particular type of vowel may in fact indirectly correlate to its quality as a target.

In relation to this argument, it is worth noting that certain cases of harmony specifically do not create the best possible vowels.[79] For instance, in (Ethiopian) Komo (Otero 2015), in which the only contrastive ATR vowels are high, RTR-ATR sequences harmonize to ATR-ATR if the RTR vowel is mid or low, as in /hám-úk/ surfacing as [hə́m-úk] ‘yawn-PFV’, but to RTR-RTR if the RTR vowel is high, as in [yìl-ʊ́k] ‘see-PFV’ with the same perfective suffix /-úk/. In both cases, the ATR vowel must be high, since the only contrastive ATR vowels are high in Komo. In other words, harmony always creates the worst option in terms of vowel quality: in the former case, harmony creates ATR mid and low vowels, which do not otherwise exist in the language, instead of creating RTR high vowels that do; in the latter case, it creates RTR high vowels, instead of the articulatorily easier ATR high vowels. While this pattern is difficult to analyze under any framework, it strongly suggests that the goal of harmony is not creating ‘good’ vowels. As such, it is not clear whether the fact that high vowels are ‘good’ at being ATR is relevant to whether they are good targets.

All of these facts suggest that we do not fully understand the motivations and typology of ATR dominant harmony well enough to know what factors are relevant to dealing with the height asymmetries in targets in Lango; it could be parasitic harmony or a preference for high targets, and both ideas have advantages as well as problems. I will assume that the pattern results from a preference for high targets here, since it seems most consistent with the typological facts and is simpler to formalize. However, as noted, the other choice is also consistent with the present analysis. Specifically, target scaling could either be based on agreement in height with

---

[79] This connects to Casali’s point that RTR high vowels are not underattested; they occur in the majority of languages with ATR contrasts, but given that 2IU languages are associated with ATR dominance, the RTR vowels are never created by harmony, but instead are underlying.

the trigger or on high versus non-high target (without reference to the trigger); either way the same principles apply and the analysis works, but the latter option is formally much more straightforward.

3.3 Distance

Within both progressive and regressive harmony in Lango, vowels that are separated by only a single consonant are more likely to harmonize than vowels at a greater distance. For example, progressive harmony can be triggered by non-high vowels across a single consonant, but not across more than one. This effect can be attributed to distance: vowels further apart are less likely to harmonize.

Distance effects can be found in many types of harmony systems across languages. In terms of cases in which the presence of multiple consonants blocks vowel harmony, a pattern similar to the Lango one can be found in another Western Nilotic language, Alur (Kutsch Lojenga 1991; see Chapter 7); in Alur, CV suffixes never undergo harmony, even though V suffixes sometimes do. It is unclear whether such cases are best described in these terms, as presence versus absence of suffix onset, instead of number of intervening consonants. Distance effects can also be seen in other ways, including in the number of intervening syllables; an example is the count effect in Hungarian, in which multiple ‘transparent’ vowels can block harmony when a single one cannot (Hayes & Londe 2006; see Chapter 3). This effect can be analyzed as due to the trigger of harmony being separated by too many syllables from the potential target.[80]

---

[80] A similar analysis could possibly apply to some instances of opaque vowels in harmony, particularly those in which the vowels bear the harmonic feature but harmony does not propagate past them (e.g. Tungusic rounding harmony; Li 1996); the vowel that is able to trigger harmony is made too distant from the potential target by the intervening opaque vowel. However, in many of these cases, the rounded high vowels block harmony, while the

All of these examples are cases that illustrate the importance of proximity between triggers and targets in harmony systems. The effect in Lango is precisely as expected: vowels that are further apart are less likely to be required to harmonize, where distance is measured in intervening consonants (or moras, see Archangeli & Pulleyblank 1994; whether we count consonants or moras is not critical to this analysis, except that if geminates are a single consonant, then the distinction between geminates and non-geminate clusters could be analyzed straightforwardly in terms of intervening consonants, but not in terms of intervening moras). This distance generalization needs to be incorporated into the theory, preferably with a mechanism that is also able to generalize to languages that count distance in other ways (like Hungarian). This will be achieved with a distance scaling factor in the present analysis.

3.4 Directionality

In terms of directionality, according to the facts in Woock & Noonan (1979), progressive harmony is stronger than regressive harmony in Lango[81], as seen for instance by the fact that progressive harmony can be triggered by non-high vowels, whereas regressive harmony cannot be. This pattern is precisely the opposite of what we expect based on cross-linguistic tendencies: in general, there is a right-to-left bias in both vowel and consonant harmonies (e.g. Baković 2000; Hyman 2002; Hansson 2010). Even in other Western Nilotic languages, including Alur (Kutsch Lojenga 1991; see Chapter 7) and Mayak (Andersen 1999; see Chapter 9), harmony occurs more in regressive than in progressive contexts.

---

non-rounded high vowels are transparent. As such, a distance-based account would also need to be able to consider the nature of the intervening segment. In theory, this should be possible because of the (non-important) disharmony created by the intervening vowels, but an implementation is beyond the scope of this paper.

[81] However, based on Noonan (1992), regressive harmony may in fact be stronger, given that it is less restricted in terms of which vowels participate and intervening consonants, whereas progressive harmony applies only to vowel-initial suffixes and across geminates. See Section 6 for more discussion.

The restrictions on triggers and targets are not the only sense in which progressive harmony in Lango is stronger; regressive harmony is also non-iterative, which is not the case for progressive harmony, as evidenced by the middle voice suffix /-ɛrɛ/ harmonizing to [-ere] (Noonan 1992; Archangeli & Pulleyblank 1994; Kaplan 2008). This non-iterativity is often ignored in previous analyses, though it is stipulated in the analysis of Archangeli & Pulleyblank (2004) and is the basis of the analysis by Kaplan (2008), who views it as a licencing condition. Specifically, he argues that ATR in a suffix must be licenced by ATR in a root, but changes to the root in order to enforce that licencing are minimal, and therefore affect only the final root vowel. Given that high vowels are apparently the ones in which the ATR contrast is more difficult to perceive, it makes sense that the contrast in suffixes might need to be licenced for high but not mid vowels, though it is strange under this argumentation that low suffix vowels are not subject to the same requirement.

Given this idea, a possible way of understanding the directionality distinction in Lango is root faithfulness; regressive harmony involves changes to roots, and therefore might be expected to be more restricted, given that roots are a preferred position for faithfulness (e.g. Beckman 1998; Baković 2000). Progressive harmony, on the other hand, affects suffixes, and is therefore not restricted by root faithfulness. Note that prefixes never participate in Lango harmony; they neither trigger nor undergo it (Noonan 1992).

For the purposes of this paper, I follow some previous analyses of Lango (e.g. Potts et al. 2010) in setting aside non-iterativity. I treat the differences in progressive versus regressive harmony as due to the strength of the harmony constraint; for reasons that are not clear, Lango simply has a greater requirement to harmonize ATR…RTR sequences than the reverse. However, I will show briefly at the end of Section 4.3 that a root faithfulness view would also work within the present account. In terms of non-iterativity, this fact could be due to domain restrictions; Pulleyblank (2001), for instance, argues that Lango regressive harmony is limited to a binary foot domain, and provides an account of this domain. Such a view is consistent with the analysis I present in Section 4.

4 Analysis

The previous sections outline the Lango harmony system and explain it in the context of broader cross-linguistic facts. These generalizations have been noted and analyzed before, as by Archangeli & Pulleyblank (1994), Smolensky (2006), and Potts et al. (2010). The novel aspect of the present paper is incorporating the generalizations into the analysis in a different way. The present section explains and illustrates my approach, while a comparison to previous accounts is left for Section 5.

The analysis presented here accounts for the facts in Table 8.3 in Section 2, which is also the set of facts accounted for in Archangeli & Pulleyblank (1994), Smolensky (2006), and Potts et al. (2010).[82] I adopt a framework in which weighted harmony constraints are scaled by trigger, target, and distance, as motivated independently for other languages by Kimper (2011) and Chapter 3 on Hungarian. Scaling factors are determined by articulatory and perceptual properties that make certain vowels better at triggering or undergoing harmony; languages have a limited set of phonetically motivated options. For example, in ATR harmony, high vowels are ideal triggers, as discussed in Section 3. In all ATR-dominant harmony languages, trigger scaling factors must therefore be set such that harmony triggered by high vowels is either more likely

---

[82] I will set aside RTR harmony as exceptional, as has been done in previous analyses except for Archangeli & Pulleyblank (1994). I will also set aside non-iterativity, again following previous analyses like Potts et al. (2010); one option consistent with the current analysis would be to follow the proposal of Pulleyblank (2001) for a domain restriction. While there are some analysis (e.g. Kaplan 2008) that deal with these facts, they do not deal with the broader harmony system, in particular with the trigger/target/distance asymmetries that are the present focus.

than or equally likely as harmony triggered by other vowels.

4.1 General framework

This analysis is couched in a Harmonic Grammar framework (HG; Legendre et al. 1990a; b; Smolensky & Legendre 2006; Pater 2009), in which constraints are weighted rather than ranked. Like Potts et al. (2010), among many others, I use negatively formulated constraints with weights that are positive integers.

Previous HG analyses of other types of harmony allow for multiplying violations of the harmony constraint by a scaling factor, so that the weight of the violation can be influenced by properties like trigger, target, and distance (e.g. Kimper 2011; Ozburn 2019a). A scaling factor can either increase or decrease the penalty incurred for a specific violation. For example, it was discussed above that high vowels are better triggers of ATR harmony than non-high vowels. Giving high vowels a larger trigger strength scaling factor compared to non-high vowels will mean that the penalty for disharmony will be greater if the potential trigger is high than if it is non-high. Additional details will be presented below.

A locus of violation of a constraint will have penalty $-W*S_1*\ldots S_n$, where W is the constraint weight and the S_i are the relevant scaling factors (if any). The penalties a candidate incurs for all constraints are then added to determine its ‘harmony score’, and the candidate with the highest such score wins (e.g. Potts et al. 2010). Note that the harmony score is always negative here, since constraints are negatively formulated, so the winning candidate is the one with a score closest to zero.

As noted in Chapter 3, the nature of the harmony-driving constraint does not matter, as long as trigger and target can be distinguished. Following the approach there, I adopt no-disagreement-style constraints (Pulleyblank 2004), with the focus of the constraint identified as the target and the environment as the trigger. Specifically, *[RTR]∞[ATR] encodes a demand for regressive harmony, in the sense of banning the disharmonic co-occurrence of an RTR target vowel with a following ATR trigger (at any distance), while *[ATR]∞[RTR] will enforce progressive harmony; these constraints will be defined formally in detail in Section 4.2. Note, however, that these constraints do not themselves determine the repair; they simply ban particular disagreeing sequences, and the repair is determined by faithfulness (and/or other markedness constraints that might be violated more by harmony in one direction than the other). As such, even though the RTR vowel is identified as the 'target' in both of these constraints, and will be scaled as such, these constraints allow for the possibility of the RTR vowel actually acting as the trigger. It is higher-ranked faithfulness to ATR than to RTR that will establish that ATR is the dominant feature in Lango.

As mentioned, the harmony constraints will be scaled based on distance, triggers, and targets. The relative scaling factors, for instance high vowels being scaled as better triggers than non-high vowels, are as motivated in Section 3 in reference to cross-linguistic and phonetic factors about ATR and other harmony systems. I assume that both harmony constraints must be subject to identical scaling factors, in that learners assume that constraints for a given type of harmony must be subject to identical scaling factors, in order to reduce the computational complexity. As such, differences in patterns between directionalities must be due to differences in the weights of the harmony constraints; this choice reduces the computational complexity of the analysis. The scaling factors will also be defined in the Section 4.2. Scaling factors are variables that a learner sets, independent of constraints and their ranking, but which are referenced by constraints; a given constraint is tagged for a particular set of scaling factors. I assume that the set of scaling factors for which a specific constraint is tagged is consistent cross-linguistically; languages may vary only in setting the values of the scaling factors and constraint weights.

4.2 Constraint definitions, weights, and scaling factors

This subsection formally defines all necessary constraints and provides their weights and scaling factors. Note that for both weights and scaling factors there are a variety of options for the specific numbers; what is crucial is for the relationships among the numbers to fulfill certain criteria. These criteria will be discussed in detail in the presentation of the tableaux in Section 4.3.

As discussed above, I consider ATR dominance to be derived by a separation of faithfulness constraints preserving ATR and RTR. These faithfulness constraints are defined formally in (9).

(9)

IDENT-IO[ATR]: Assign a violation for every segment that is [ATR] in the input but [RTR] in the output.

IDENT-IO[RTR]: Assign a violation for every segment that is [RTR] in the input but [ATR] in the output.

In order to make ATR the dominant feature, IDENT-IO[ATR] must outweigh its counterpart for RTR, so that ATR vowels will never change. I assign it a weight of 100, which is sufficient to ensure that vowels will never change from ATR to RTR; I will therefore omit this constraint from the tableaux. IDENT-IO[RTR], on the other hand, needs to be weighted low enough that vowels will sometimes change their RTR value in order to harmonize. I assign it a weight of 10; the exact number is arbitrary, and the requirements on the relative weight of this constraint compared to the harmony constraints will be developed through the tableaux in Section 4.3.

In terms of markedness constraints, as discussed above, we require constraints for progressive and regressive harmony, which simply ban RTR vowels preceded by and followed by ATR vowels respectively. The directionality of harmony is enforced by the combination of these constraints and the high weighting of IDENT-IO[ATR]. The formal definitions of the harmony constraints are given in (10), adapted from Pulleyblank (2004). Both of these constraints consider RTR to be the focus, meaning that violations are counted based on the RTR vowels. Additionally, target scaling will be based on the RTR vowels, while trigger scaling is based on the nearest ATR vowel.

(10)

*<u>[RTR]</u>∞[ATR]: Assign a violation for every RTR vowel followed at any distance by an ATR vowel.

*[ATR]∞<u>[RTR]</u>: Assign a violation for every RTR vowel preceded at any distance by an ATR vowel.

Since regressive harmony is weaker than progressive harmony in Lango, it should have lower weight. I adopt a weight of 2 for *<u>[RTR]</u>∞[ATR], and a weight of 6 for *[ATR]∞<u>[RTR]</u>. Both of these constraints, however, are scaled by identical values based on the quality of the trigger and target and the distance between them, as discussed in Section 4.1. These scaling factors, defined in (11-13), are taken from Chapter 3, where they were adapted from the trigger strength and distance scaling in Kimper (2011).

(11) Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that $x \geq 1$) for each degree i to which α is perceptually impoverished with respect to ±F.

(12) Scaling factor: distance

For a trigger α and a target β, multiply the penalty earned by a constant x (such that $0<x<1$) for each unit (e.g. syllable) of distance d intervening between α and β.

(13) Scaling factor: target quality

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that $x \geq 1$) for each degree i to which β is a quality target (defined by phonetic factors discussed in Section 3) with respect to ±F.

As discussed above, Lango treats high vowels as preferential triggers, particularly high front vowels, and also has a preference for targets either being high or agreeing in [high] with their triggers. To produce the situation of height agreement, a modification would be required in the target quality scaling proposed by Chapter 3, in that the degree to which β is a quality target depends on α. This possibility was noted in relation to parasitic rounding harmony in that paper, but not developed in any detail. Note, however, that as discussed below, the Lango analysis will work out regardless of whether we assume that the preference is for high targets as such or height agreement between trigger and target. Since it is far more straightforward analytically, I assume the former route in the present analysis. Further, harmony is enforced more strongly across a single consonant than across more than one in Lango, and so the unit of distance in that scaling factor is about the number of consonants. The relevant numerical values for all of the scaling factors are given in (14). Note that these relationships among the numbers, for instance in high vowels having higher trigger scaling than non-high vowels, were all justified in the preceding section. Scaling factors apply solely to the harmony constraints in (10). In other words, the harmony constraints are tagged for these particular scaling factors, set in this particular way in Lango, and no other constraints see this type of scaling.

(14a) Trigger scaling

8 for high front ([i])

4 for high back ([u])

1 for non-high vowels

(14b) Target scaling[83]

3 for high target (or agreement in [high] with trigger)[84]

2 for non-high target (or no agreement in [high] with trigger)

---

[83] This is where the two possible analyses of target behaviour discussed in Section 3.2 are relevant. Here, I scale based simply on high versus non-high, but the analysis would also work if this were based on agreement in [high]. However, having a target scaling factor that needs to look at the trigger in order to determine whether there is agreement is clearly more complicated, and the exact formal implementation of this would be difficult. If height agreement, as opposed to high versus non-high target status, is in fact the relevant factor, this problem would need to be explored further.

[84] Given how low the trigger scaling is for non-high vowels, higher target scaling due to agreement in [high] would have no effect on combinations of mid and low vowels; they would still not be required to harmonize, despite agreeing in [high].

(14c) Distance scaling

1 for across a single consonant

0.5 for across multiple consonants

4.3 Derivations

Given the high weight of IDENT-IO[ATR], I will omit it from the tableaux, along with any candidates that change ATR vowels to RTR. These candidates are always dispreferred to those that change RTR vowels to ATR, given the much lower weight of IDENT-IO[RTR].

Looking first at progressive harmony across a single consonant, it will correctly apply regardless of triggers and targets, given the constraint weights and scaling defined in Section 4.4. In this case, the distance scaling will be 1 (across a single consonant), the lowest possible trigger scaling is 1 (if the trigger is non-high), and the lowest possible target scaling is 2 (if the target is non-high). The progressive harmony constraint has weight 6, so when scaled by even the lowest possible combination of scaling factors, a violation will incur a penalty of 6*1*1*2 = 12. The faithfulness constraint IDENT-IO[RTR] has weight 10, so since -10 > -12, it is preferable to change RTR than to have disharmony. The tableau in Table 8.4 illustrates. The condition on weights and scaling required here is that weight(IDENT-IO[RTR]) < weight(*[ATR]∞[RTR]) * scaling(across a C) * scaling(worst trigger) * scaling(worst target).

	/wót + á/	IDENT-IO[RTR] 10	*[ATR]∞[RTR] 6	H
	wódá		-1*1(o)*2(a) =-2	-12
☞	wódə́	-1		-10

Table 8.4: Progressive harmony across a single consonant

Turning now to progressive harmony across multiple consonants, it will correctly apply only for high triggers. If the trigger is high, the lowest possible scaling will be 4 (trigger [u]) * 2 (non-high target) * 0.5 (across multiple Cs) = 4. Since the weight of the progressive harmony constraint is 6, the total penalty incurred is 4*6 = 24; since -10 > -24, the candidate that violates IDENT-IO[RTR] wins. The result is shown in Table 8.5. The condition we require here is weight(IDENT-IO[RTR]) < weight(*[ATR]∞[RTR]) * scaling(across multiple Cs) * scaling(worst high trigger) * scaling(worst target).

	/òpúk + ká/	IDENT-IO[RTR] 10	*[ATR]∞[RTR] 6	H
	òpúkká		-1*4(u)*2(a) *0.5(dist) =-4	-24
☞	òpúkkə́	-1		-10

Table 8.5: Progressive harmony across multiple consonants from a high trigger

In contrast, with a non-high trigger across multiple consonants, the highest possible scaling is 1 (non-high trigger) * 3 (high target) * 0.5 (across multiple Cs) = 1.5, creating a total penalty of 1.5*6 = 9. Since -9 > -10, it is preferable here to maintain disharmony than to violate IDENT-IO[RTR]. Table 8.6 illustrates; the condition required is weight(IDENT-IO[RTR]) > weight(*[ATR]∞[RTR]) * scaling(across multiple Cs) * scaling(non-high trigger) * scaling(best target). Due to available examples, the table does not show a high target; further discussion of Lango morphology and the effect of geminate versus non-geminate clusters is reserved for Section 6.

	/dòk + ká/	IDENT-IO[RTR] 10	*[ATR]∞[RTR] 6	H
☞	dòkká		-1*1(o)*2(a) *0.5(dist) =-1	-6
	dòkkə́	-1		-10

Table 8.6: No progressive harmony across multiple consonants from a non-high trigger

The situation for regressive harmony differs because the constraint enforcing it is weighted lower than the one enforcing progressive harmony. Looking first across a single consonant, harmony will occur with a high trigger: the lowest possible scaling is 4 (trigger [u]) * 2 (non-high target) * 1 (across a C) = 8. Given that the weight of *[RTR]∞[ATR] is 2, the total penalty is 8*2 = 16. Since -10 > -16, it is preferable to violate IDENT-IO[RTR] in order to harmonize. This is shown in Table 8.7; the required weighting relationship is weight(IDENT-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(across a C) * scaling(worst high trigger) * scaling(worst target).

	/jɔ̀ + wú/	IDENT-IO[RTR] 10	*[RTR]∞[ATR] 2	H
	jɔ̀wú		-1*4(u)*2(ɔ) = -8	-16
☞	jòwú	-1		-10

Table 8.7: Regressive harmony from a high vowel across a single consonant

When the trigger is non-high, on the other hand, harmony will not occur. The highest possible harmony scaling in this case is 1 (non-high trigger) * 3 (best target) * 1 (across a C) = 3, so the total penalty for violating harmony is 3*2 = 6. Since -6 > -10, disharmony is preferred in this case to a violation of IDENT-IO[RTR]. Table 8.8 illustrates, and the required weighting is weight(IDENT-IO[RTR]) > weight(*[RTR]∞[ATR]) * scaling(across a C) * scaling(non-high trigger) * scaling(best target). Given the available examples across a single consonant, I illustrate with a non-high target; however, with a high target, the disharmonic candidate will still win, as described above.

	/way + o/	IDENT-IO[RTR] 10	*[RTR]∞[ATR] 2	H
☞	wayo		-1*1(o)*2(a)	-4
	wəyo	-1		-10

Table 8.8: No harmony from a non-high trigger across a single consonant

Of course, since regressive harmony does not occur from a non-high trigger (in the data as described by Woock & Noonan 1979) across a single consonant, it will also not occur across multiple consonants, since the highest possible scaling across multiple consonants is less than that across a single consonant. Specifically, the highest scaling for a non-high trigger across multiple consonants is 1 (non-high trigger) * 3 (best target) * 0.5 (across multiple Cs) = 1.5, so the total penalty for violating harmony is 1.5*2 = 3. Since -3 > -10, disharmony is preferred in this case as well. This requires weight(IDENT-IO[RTR]) > weight(*[RTR]∞[ATR]) * scaling(across multiple Cs) * scaling(non-high trigger) * scaling(best target).

Regressive harmony across multiple consonants is where the separation of high front from high back triggers becomes critical, as does target scaling.[85] Looking first at cases where [i] is the trigger, we will see harmony regardless of the target; the lowest possible scaling is 8 (trigger [i]) * 2 (worst target) * 0.5 (across multiple Cs) = 8, so the penalty for disharmony is 8*2 = 16. Since -10 > -16, it is preferable to harmonize by violating IDENT-IO[RTR]. Table 8.9 illustrates; we require weight(IDENT-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(across multiple Cs) * scaling(trigger [i]) * scaling(worst target).

	/kɔ̀m + mí/	IDENT-IO[RTR] 10	*[RTR]∞[ATR] 2	H
	kɔ̀mmí		-1*8(i)*2(ɔ) *0.5(dist) =-8	-16
☞	kòmmí	-1		-10

Table 8.9: Regressive harmony from [i] across multiple consonants

In contrast, when the trigger is [u], the result depends on the target. If the target is high, then the scaling of harmony will be 4 (trigger [u]) * 3 (high target) * 0.5 (across multiple Cs) = 6. Since 6*2 = 12, and -10 > -12, harmony will occur in this case, because disharmony is worse than a violation of faithfulness. This is shown in Table 8.10, and requires weight(IDENT-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(across multiple Cs) * scaling(trigger [u]) * scaling(high target).

---

[85] As noted above and discussed in more detail in Section 6, the distinction between high back and high front triggers might in fact be about intervening geminates versus non-geminate clusters.

	/lʊ̀t + wú/	IDENT-IO[RTR] 10	*[RTR]∞[ATR] 2	H
	lʊ̀twú		-1*4(u)*3(ʊ) *0.5(dist) = -6	-12
☞	lùtwú	-1		-10

Table 8.10: Regressive harmony from [u] across multiple consonants to high target

On the other hand, if the target is non-high, then harmony will not occur; the scaling will be 4 (trigger [u]) * 2 (non-high target) * 0.5 (across multiple Cs) = 4. The total penalty for disharmony is therefore 4*2 = 8, and -8 > -10, so in this situation, disharmony is preferable to violating IDENT-IO[RTR]. Table 8.11 illustrates; we require weight(IDENT-IO[RTR]) > weight(*[RTR]∞[ATR]) * scaling(across multiple Cs) * scaling(trigger [u]) * scaling(non-high target). Given the weighting condition for Table 8.10, this means that scaling(high target) > scaling (non-high target) is necessary.

	/kɔ̀m + wú/	IDENT-IO[RTR] 10	*[RTR]∞[ATR] 2	H
☞	kɔ̀mwú		-1*4(u)*2(ɔ) *0.5(dist) =-4	-8
	kòmwú	-1		-10

Table 8.11: No regressive harmony from [u] across multiple consonants to non-high target

Now that the analysis of harmony has been presented, it is worth returning to the issue of root faithfulness. Adding in a root faithfulness constraint and weighting the progressive harmony constraint lower would clearly be an alternative way of dealing with the preceding tableaux. Specifically, consider a constraint IDENT-IO[RTR]$_{\text{ROOT}}$, which assigns a violation for every root vowel that is RTR in the input and ATR in the output. In all the tableaux for regressive harmony (Tables 8.7-8.11), this constraint is violated in exactly the same cases that IDENT-IO[RTR] is violated, meaning that their total weight only needs to be 10. Say, then, that the weight of root-specific faithfulness is 7, and that of the general faithfulness is 3. As mentioned, this affects nothing for regressive harmony, given that all vowels that change in RTR are within the root. In contrast, for progressive harmony (Tables 8.4-8.6), the alternating vowels are in the suffix, where root faithfulness is not relevant. As such, a lower weight of general faithfulness means that the weight of the harmony constraint can be reduced to equal to that of regressive harmony, while still consistently predicting the correct form. The tableaux in Tables 8.12-8.14 illustrate. In other words, an analysis of Lango in this framework does not actually require that progressive harmony be stronger than regressive harmony, which is useful given that cross-linguistically, regressive harmony is thought to be the default.

	/wót + á/	IDENT-IO[RTR] 3	*[ATR]∞[RTR] 2	H
	wódá		-1*1(o)*2(a) =-2	-4
☞	wódə́	-1		-3

Table 8.12: Progressive harmony across a single consonant

	/òpúk + ká/	IDENT-IO[RTR] 3	*[ATR]∞[RTR] 2	H
	òpúkká		-1*4(u)*2(a) *0.5(dist) =-4	-8
☞	òpúkkə́	-1		-3

Table 8.13: Progressive harmony across multiple consonants from a high trigger

	/dòk + ká/	IDENT-IO[RTR] 3	*[ATR]∞[RTR] 2	H
☞	dòkká		-1*1(o)*2(a) *0.5(dist) =-1	-2
	dòkkə́	-1		-3

Table 8.14: No progressive harmony across multiple consonants from a non-high trigger

4.4 Summary of analysis

This approach, with weights and scaling factors defined by the criteria presented throughout this subsection, can account for all of the data on Lango ATR harmony dealt with in previous approaches. It uses only a combination of faithfulness and harmony constraints, but with the insight that the penalty for disharmony should depend on the particular interacting segments and their distance apart. Lango harmony arises fairly straightforwardly as interactions of triggers, targets, and distance, with these factors motivated by cross-linguistic and phonetic facts.

It is worth noting that, while Harmonic Grammar makes this approach much easier to illustrate, it is not actually necessary. It would alternatively be possible to have multiple harmony constraints for each trigger/target/distance combination, such as *[+high,ATR]CC[RTR] versus *[-high,ATR]CC[RTR] to reflect the difference between high and non-high triggers with multiple intervening consonants and any target. These constraints could be ranked in a fixed-ranking hierarchy (Prince & Smolensky 1993), to mirror the universal hierarchy in the scaling factors used here. Alternatively, a stringency relationship of constraints could be used. For example, while the 'context' in the constraints used here is simply ATR, with trigger scaling factors reflecting the distinctions in trigger behaviour, we could instead have separate constraints for each trigger, one with context [+high,-back,ATR], one with context [+high,ATR], and one with context [ATR]. No matter how constraints with such contexts are ranked, we only get the possibilities reflected in the way that scaling factors are used here. Similarly, the distance effect could be captured using a distinction between C and C_0 in place of the ∞, and the high versus non-high target preferences could be analyzed using the focus <u>[+high, RTR]</u> versus simply <u>[RTR]</u>. This would give a total of 3 (contexts) * 2 (intervening material) * 2 (focus) = 12 constraints for each directionality that would operate in standard OT and predict the same set of possibilities as the scaling factors adopted here. Further details of such an approach and a comparison between it and the present account, including any possible differences in prediction, remain to be determined in future work.

4.5 A comment on RTR harmony

Like previous analyses, the present account has set aside the RTR harmony with the infinitive suffix, which was illustrated in (8). Dealing with this pattern, even considering it to be exceptional, could be complicated. Exceptional faithfulness and markedness constraints (e.g. Pater 2007) could help, in that an exceptional version of IDENT-IO[RTR]$_{\text{ROOT}}$ could be ranked higher to derive dominance reversal, or an exceptional harmony constraint could be higher ranked.

The potentially problematic part is explaining why back vowels are the only triggers. One possibility is that there is a difference between preferred triggers in ATR versus RTR dominance, but this option raises the question of whether dominance can really just be about faithfulness. Another possibility is that there is a requirement for featural agreement between the trigger and target, which we have seen possible evidence for elsewhere, but that there is a different use of features in creating the scale. Specifically, backness here is used instead of height. Elsewhere in Lango, for ATR harmony, front vowels are the better triggers, which means it might be sensible to have back vowels be the better triggers of RTR harmony. However, incorporating backness into the scaling without predicting potential effects on the rest of the pattern is problematic. In particular, it would mean that in some form of an *RTR…ATR constraint, the RTR must be scaled such that back vowels should be more likely to undergo or trigger harmony. However, such a constraint is also applicable in other instances, like [kɔ̀mwú] ‘your chair’, which do not harmonize; ensuring there would be no effect on such forms is complicated because of the fact that the good triggers/targets of the RTR harmony do not in any way interact with the ATR harmony, and so this is beyond the scope of the present paper.

Moreover, while height clearly affects the typology of ATR harmony systems, backness does not to the same extent, in that there are not the same types of significant, cross-linguistic asymmetries, such as inventory gaps and determination of the type of harmony, in front versus back vowels as there are in high versus non-high vowels. There are other languages in which front and back vowels behave differently with respect to ATR harmony, such as Kinande (e.g. Archangeli & Pulleyblank 1994; 2002), but front/back asymmetries are far less widespread and critical than height asymmetries in ATR harmony. In particular, while it is widely noted that the presence or absence of tongue root contrasts in high and mid vowels affect the existence and type of tongue root harmony (Casali 2003; 2008; Rose 2018), there are no such claims about backness.

It is worth noting that the analysis here does not insist that the 'trigger' and 'target' actually behave as such. In the case of RTR progressive harmony, as in /lʊbb-o/ → [lʊbb-ɔ] 'to follow', the actual trigger is /ʊ/ and the target /o/, but for the purposes of scaling, it could be the reverse, with back vowels in the root like [ʊ] scaled as 'targets', and the suffix as the 'trigger'. In such a case, it would be exceptional faithfulness forcing the opposite change. Again, however, such an approach would raise the question of why the mid vowel /o/ in the infinitive suffix causes this change, when mid vowels are not good 'triggers' elsewhere. Instead, it does seem preferable to consider the mid vowel as the target.

Overall, the exceptional RTR harmony in Lango is problematic. It has not been treated by previous OT analyses, although it was treated by Archangeli & Pulleyblank (1994) in a pre-OT framework, and the goal of the present paper is to simply show an alternative approach to those analyses. However, working out the details of how backness affects Lango harmony is an important future direction, beyond the scope of the present paper. That being said, it is worth noting that while the infinitive suffix behaves this way, other /-o/ and /-e/ suffixes mostly do not (Noonan 1992), so this is a morphologically conditioned exception, rather than a fact about mid ATR suffixes more generally in the language.

5 Comparison to previous approaches

In this section, I argue for the analysis in Section 4 over previous ones: it maintains generalizations without over-generating, but does not require the complexities of the Potts et al. (2010) account. Instead, this analysis uses only interactions between harmony and faithfulness, with the insight that how harshly harmony is penalized should depend on the specific interacting segments and their distance apart.

5.1 Analyses prior to Potts et al. (2010)

Potts et al. (2010) provide important criticisms of analyses of Lango preceding theirs, specifically Archangeli & Pulleyblank (1994) and Smolensky (2006). In this subsection, I overview their objections to these analyses, and argue that the present approach does not suffer from these issues.

A major problem described by Potts et al. (2010) for previous analyses of Lango is that they over-generate, because they are unable to capture certain generalizations and relationships. Potts et al. (2010) note that the analysis of Archangeli & Pulleyblank (1994) uses inviolable constraints on rule applications, which means these are not true of ATR harmony in the language overall. Moreover, any combination of rules with any conditions can exist in such an analysis, meaning it loses the generalization of ‘uniform strength’, for instance that progressive harmony in Lango is uniformly stronger than regressive harmony. In the parametric rule account employed by Archangeli & Pulleyblank (1994), it is possible to derive uneven applications of rules, in which for instance, regressive harmony is subject to conditions not enforced on progressive harmony, and progressive harmony is subject to conditions not enforced on regressive harmony.

Like the Potts et al. (2010) analysis, the present analysis does not suffer from these issues. Being formulated in HG, this analysis uses violable constraints that apply consistently to all data, unlike the parametric approach. Further, scaling factors here are taken as determined once for a given type of harmony in a given language and applied to all harmony constraints to which they are relevant.[86] This assumption allows for the generality and restrictiveness that Potts et al. (2010) argue is lacking in previous accounts. Specifically, the difference in directionalities reduces to a difference in the weight of the progressive versus regressive harmony constraints. These weights get multiplied by the same scaling factors, so if the regressive constraint is weighted lower than the progressive one, then any combination of trigger, target, and distance that is sufficient to force regressive harmony must also be sufficient to force progressive harmony.

It would be tempting to suggest that the present analysis does, to some extent, suffer from the loss of generality that Potts et al. (2010) argue exists for the Archangeli & Pulleyblank (1994) and Smolensky (2006) accounts. This is because the information about triggers, targets, and distance gets incorporated in two separate places: both the regressive and progressive harmony constraints. However, such an argument fails to consider both the meaning and

---

[86] If a language has multiple types of harmony, such as both tongue root and rounding harmony, then scaling factors can be set separately for each type of harmony. It is only within a type of harmony that scaling factors are set a single type. For instance, all ATR harmony in a language will have one set of scaling factors. See Section 4.1 for more details.

implementation of scaling factors. The language chooses just once to, for instance, scale high vowels as better triggers than non-high ones; it is simply the case that this choice is applicable in two separate places in the grammar. This fact is in contrast to what Potts et al. (2010) argue is problematic in previous accounts, where the language chooses multiple times. For example, in the Smolensky (2006) analysis, the constraint enforcing high triggers appears both on its own and in two conjoined constraints; in the Archangeli & Pulleyblank (1994) account, high triggers are imposed separately on multiple rules. In both cases, the loss of generality comes from the fact that the language needs to decide multiple times to impose the condition of high vowels being preferential triggers; it is a separate decision for every constraint conjunction or parameter on a rule. Crucially, in the present analysis, it is a single decision, in the scaling factors, that simply happens to be applied in multiple places.

Despite these arguments, it is worth noting that Archangeli & Pulleyblank's (1994) analysis is quite similar to the present one, in that they posit rules of varying strengths, which to apply need to be substantively boosted by particular triggers and targets. This is similar to the idea of different weightings of constraints, which are boosted by scaling of triggers and targets. The basic insight is similar, but the present analysis is formulated in a more modern framework that provides greater restrictiveness. While Archangeli & Pulleyblank (1994) treat the retraction cases that are not treated here, the additional data in Noonan (1992) suggests that those cases are in fact exceptional, related only to a specific suffix.

5.2 Potts et al. (2010)

Like the current analysis, the one in Potts et al. (2010) does not suffer from the previously outlined problems in the accounts of Archangeli & Pulleyblank (1994) and Smolensky (2006).

However, in this subsection, I argue that the Potts et al. (2010) analysis is unnecessarily complex, does not capture the main insights of the Lango pattern, and cannot actually be formulated given standard assumptions about Optimality Theory (and derivative frameworks like HG).

The primary issues with the Potts et al. (2010) analysis come from the use of a particular notion of heads. In their account, the head of an ATR domain in an output representation must definitionally satisfy two requirements: it must have been underlyingly ATR, and it must be part of an ATR domain that consists of more than a single element. Using this notion, the account crucially relies on a number of constraints about heads: HEAD[high], HEAD-L, and HEAD[front], which penalize heads that are not high, leftmost in their ATR domain, and front respectively. Given the definition of head, these constraints are essentially penalizing instances in which ATR has spread from an underlyingly ATR vowel that is non-high, not leftmost, and not front, respectively. In addition to these head constraints, LOCAL-C and ATR[high], which penalize clusters within an ATR domain and non-high ATR vowels respectively, also conflict with harmony. Potts et al. (2010) adopt a SPREAD[ATR] constraint to motivate harmony; this constraint must apply in a broader prosodic domain than the domain considered for the head constraints. In all instances in which the winning candidate is not harmonic (i.e. does not satisfy SPREAD[ATR]), it is gang effects among the other constraints that rule out the harmonic candidate. For example, in their T6 in Table I, expanded here in Table 8.15, a gang effect of HEAD[high] and HEAD-L rules out the harmonic candidate, because the head [e] in the harmonic candidate is neither high nor on the left. Notably, there is no head in the winning (disharmonic) candidate, because there is only a single element ([e]) in the ATR domain, and Potts et al. (2010) define heads as being in minimally binary domains. In other words, there is no head in the candidate [ɪCe], and the head in the candidate [iCe] is the underlyingly ATR vowel.

	/ɪCe/	SPREAD [ATR] 11	HEAD[high] 8	HEAD-L 4	LOCAL-C 4	HEAD[front] 2	ATR[high] 2	H
☞	ɪCe	-1					-1	-13
	iCe		-1	-1			-1	-14

Table 8.15: Potts et al.'s (2010) T6 from Table I, expanded

In terms of their function in the analysis, the HEAD[high] constraint behaves like the scaling factor that prefers high triggers, while HEAD[front] adds the additional preference for high front triggers. LOCAL-C encodes the distance effect, which in the present analysis is done through the distance scaling factor. The effect of HEAD-L in the Potts et al. (2010) analysis is captured in the present analysis through the difference between regressive and progressive harmony constraints, and ATR[high] is represented in this analysis as a preference for high targets, through target scaling factors. As such, both analyses consist of components that do equivalent work in capturing the asymmetries of Lango harmony, yet in critically different ways.

In order for the violations of the head constraints in Potts et al. (2010) to be assigned correctly, heads must be identified somehow either in the representation or by the constraint. Beyond the properties that the constraints themselves enforce, heads are required to satisfy a number of characteristics that are not independently encoded in constraints. First, they must be ATR, since these constraints refer only to heads of ATR domains. Second, they must be faithful to their underlying representation, which is an assumption that Potts et al. (2010) note, and means that heads must be underlyingly ATR. Third, they must be associated to an ATR feature that has other dependents, since Potts et al. (2010) stipulate that ATR domains are minimally binary, and that a representation with an ATR feature associated to a single vowel does not have a head. This last property is a sharp departure from the assumptions adopted about headedness of spans in McCarthy's Span Theory (2004), as spans can consist of just a single segment in that approach.

The heads in this analysis must be of ATR domains, and it cannot be the case that the relevant head constraints apply to RTR as well. This fact is clear from T6 in Table I from Potts et al. (2010: 97), expanded as Table 8.15 above, in which the input ɪCe surfaces faithfully. Potts et al. (2010) do not consider a candidate with RTR harmony, like ɪCɛ, but this candidate vacuously satisfies SPREAD[ATR], since there is no ATR to spread. Additionally, since the head constraints refer to heads of ATR domains, and there is no ATR domain in this candidate, it vacuously satisfies all of those as well. It would even satisfy the head constraints if they can refer to either tongue root value, since the 'head' of the RTR domain in this example would be a high vowel on the left. As such, on the constraints that Potts et al. (2010) consider, this candidate is preferable to the disharmonic one, which is supposed to win. This is shown in Table 8.16 below.

	/ɪCe/	SPREAD [ATR] 11	HEAD[high] 8	HEAD-L 4	LOCAL-C 4	HEAD[front] 2	ATR[high] 2	H
	ɪCe	-1					-1	-13
☞	ɪCɛ							0

Table 8.16: Potts et al.'s (2010) T6 from Table I, comparing disharmony to RTR harmony

While Potts et al. (2010) do mention that they are leaving out the faithfulness constraint IDENT-

IO[ATR] violated by harmony, they do not mention that, in fact, this constraint needs to be differentially weighted depending on whether it penalizes ATR vowels becoming RTR or the reverse. Indeed, the only way to rule out harmonizing ɪCe by making both vowels RTR is to require faithfulness to the /e/, above harmony and above faithfulness to the /ɪ/. The tableau in Table 8.16 above illustrates that the candidate with RTR harmony would be expected to win for T6 without some additional faithfulness preventing ATR to RTR changes; that faithfulness must have a weight greater than 13 to prefer the disharmonic candidate. Table 8.17 below illustrates one of their cases in which ATR harmony does occur, where the differential between candidates is 9. If the same faithfulness that must be used to prevent RTR harmony in Table 8.16, of weight greater than 13, were also adopted for the ATR change in Table 8.17, then the harmonic candidate would be penalized by an amount greater than 15, and so the disharmonic candidate would win. Thus, changes from ATR to RTR must be penalized more heavily than changes in the opposite direction.

	/iCε/	SPREAD[ATR] 11	HEAD[high] 8	HEAD-L 4	LOCAL-C 4	HEAD[front] 2	ATR[high] 2	H
	iCε	-1						-11
☞	iCe						-1	-2

Table 8.17: Potts et al.'s (2010) T1 from Table I, showing that RTR to ATR changes must be permitted

Note in particular that the failure of /e/ to become [ε] in the /ɪCe/ example cannot be about heads being faithful, since the [e] in the candidate [ɪCe] is crucially not a head. Instead, the Potts et al.

(2010) account necessitates the same type of separated faithfulness to ATR versus RTR as used in the present approach. While this might therefore seem like a point on which the analyses are equal, faithfulness is the only place in the current account in which ATR has a special status over RTR. In contrast, the Potts et al. (2010) analysis also requires having many constraints on heads of ATR domains, but none at all on heads of RTR domains.

The fact that heads must be faithful raises major issues in evaluating which vowel is the head. Specifically, it requires either some form of representational stipulation or for constraints referencing heads to be markedness constraints that are able to see underlying forms. For example, take T6 in Table I from Potts et al. (2010: 97), in which the input /ɪCe/ surfaces faithfully as [ɪCe], winning over the harmonic candidate [iCe], as shown in Table 8.15 above. This result critically requires the candidate [iCe] to violate HEAD[high] and HEAD-L in this instance, which is true because [e] is the head; it is the ATR vowel that is faithful to the input. In contrast, in T1 in the same table (Table 8.17 above), the input /iCɛ/ surfaces as [iCe], winning over the faithful candidate; in this case, neither candidate violates the aforementioned head constraints, since the faithful ATR vowel is the [i], which is both high and on the left. Critically, then, the same output candidate [iCe] violates HEAD[high] and HEAD-L in T6, but not in T1. However, there is no way to know that without referencing the underlying form; these candidates appear surface identical. We cannot say that there is an independent underlying 'head' status that is referenced by these constraints, because whether a vowel acts as a head depends on whether it has dependents in the output. For example, [e] is a head in [iCe] in T6 (/ɪCe/), but not in its faithful competitor [ɪCe]. As such, whether an ATR vowel will be a head in a candidate cannot be determined entirely from the input. It must therefore be the case that headedness gets constructed in the output, yet this cannot be a typical markedness situation, because of the aforementioned reference to underlying form. The only way to get around this problem is to add representational complexity, for instance by having underlying ATR vowels project an extra level, which then makes them a head in exactly the cases in which the ATR is additionally associated to another vowel. However, there is no independent evidence for this kind of complexity; it is simply a stipulative way to get around the problem of requiring markedness constraints to refer to which vowels are faithful to an underlying ATR value.

The stipulation that ATR domains consisting of a single vowel do not have a head is critical to the account; without it, the wrong candidate would win in all cases Potts et al. (2010) consider that involve gang effects including HEAD[high] or HEAD[front], specifically T6, T9, and T11 of their Table I. For example, in T6, if [e] were a head in the candidate [ɪCe], then this candidate also violates HEAD[high], meaning it would tie with the candidate [iCe] on this constraint; as such, the gang effect would not occur, and the harmonic candidate would win. This is shown in Table 8.18 versus Table 8.19; the former fleshes out this example from the Potts et al. (2010) table, with their assumptions about heads; the latter is the same example under the more standard assumption about heads, in which domains with a single vowel also contain heads, such that [e] is a head in the faithful candidate. As can be seen, the wrong candidate wins in the latter case, showing that this assumption is critical to their analysis.

	/ɪCe/	SPREAD[ATR]	HEAD[high]	HEAD-L	LOCAL-C	HEAD[front]	ATR[high]	H
		11	8	4	4	2	2	
☞	ɪCe	-1					-1	-13
	iCe		-1	-1			-1	-14

Table 8.18: Potts et al (2010) T6 from Table I fleshed out

	/ɪCe/	SPREAD[ATR]	HEAD[high]	HEAD-L	LOCAL-C	HEAD[front]	ATR[high]	H
		11	8	4	4	2	2	
	ɪCe	-1	-1				-1	-21
☞	iCe		-1	-1			-1	-14

Table 8.19: Potts et al (2010) T6 from Table I, without the stipulation about heads being of minimally binary domains

The importance of this binarity to heads is unusual, because it means that any reference to heads must be able to see whether an ATR domain consists of more than one vowel. While Potts et al. (2010) simply stipulate that an ATR domain is minimally binary by definition, this assumption cannot be encoded straightforwardly into constraints. It is clearly possible to get outputs like [ɪCe], where an ATR feature is associated with only a single vowel; in other words, it is not that ATR must always be associated with multiple vowels (as we might expect if the issue were solely about domains), but that a vowel can only count as a head if it has a dependent. This assumption is non-standard (e.g. it differs from Smolensky 2006), and it requires heads to have some sort of representational status, which obviously adds complexity to the phonological representation. Given the definition of head in this approach, this representation would essentially be equivalent to annotating output ATR vowels as to whether they are faithful or not, and then having the markedness constraints able to access those ‘annotations’ of faithfulness. This is a problem given that markedness constraints should only be able to reference output information in most versions of OT, and means that the Potts et al. (2010) account requires ‘two-level’ constraints. Beyond the issue of knowing which vowel was underlyingly ATR, the head-related constraints must also be able to check not only whatever specific property of the head they are checking (e.g. whether it is high), but also whether the ATR domain is minimally binary. Otherwise, these constraints cannot establish whether a head exists. In any case, there is no way to resolve the fact that the constraints HEAD[high], HEAD-L, and HEAD[front], which are markedness constraints, must also be able to see input representations in some way.

Overall, while the Potts et al. (2010) analysis looks straightforward, the numerous head constraints in fact encompass some major theoretical issues, because the notion of head they adopt is non-standard and impossible to account for with the assumptions inherent in most varieties of OT. Specifically, in order to evaluate the violations of those constraints, it is necessary to establish which vowel is the head, and knowing which vowel is the head requires both knowing which was faithful to an ATR input and whether the output candidate has multiple ATR vowels in the domain. Thus, there is a large amount of complexity hidden in the analysis, and it is unclear how the fundamental assumptions about heads could be incorporated within the analysis.

6 Discussion

6.1 Additional data from Noonan (1992)

As noted by Potts et al. (2010), the Lango generalizations as stated in Noonan (1992) differ from those in Woock & Noonan (1979); it is from the latter article that Archangeli & Pulleyblank (1994), Smolensky (2006), Potts et al. (2010), and the present chapter draw the data that they analyze. In the present section, I consider the major differences in the data from Noonan (1992), and comment on how it could be incorporated into the current analysis.

Woock & Noonan (1979) describe the suffix /-wú/ as triggering harmony only onto high vowels (when the stem ends in a consonant), as analyzed here; however, Noonan (1992) notes that this is the case only for some speakers. For instance, [dɛ̀kwú] 'your stew' is the form for such speakers, while Noonan (1992) notes that [dèkwú] is the form for most speakers. These speakers described in Noonan (1992) could be accommodated into the present analysis by not differentiating high versus non-high vowels in the target scaling. In other words, Lango harmony appears to have target restrictions for some but not all speakers, either through the speakers with the form [dɛ̀kwú] developing new restrictions, or through older target restrictions being lost for most speakers through simplification/generalization. The theory outlined here can accommodate the speaker differences straightforwardly in terms of different settings of target scaling factors.

An additional difference in Noonan (1992) is the discussion of some ATR mid vowel suffixes that trigger harmony onto the root. Examples include the ventive infinitive /-ô/, which unlike the transitive/intransitive infinitive /-ò/ discussed above (as exceptionally undergoing RTR harmony; see Section 2.4), triggers ATR harmony onto roots, as in [cwàllò] 'to send' versus [cwə̀llô] 'to send toward me' (Noonan 1992). The plural suffix /-ê/ also triggers harmony, as in [kwàc] 'leopard' versus [kwə̀c-ê] 'leopard-PL' (Noonan 1992). Such cases provide additional evidence for considering the progressive RTR harmony described in Section 2.4 as exceptional. These facts could be accounted for by increasing the weight of regressive harmony and/or the trigger scaling of mid vowels. Notably, if we assume that triggering of regressive harmony is in fact the standard behaviour for mid vowels, and that the neutral mid vowel suffixes like the transitive/intransitive infinitive are exceptional, then the morpheme-specific indexation or scaling would be for the neutral cases, not the cases of mid vowels acting as triggers; instead, it is simply that the weight and/or scaling in general must be increased.

A final fact to consider is -V versus -CV suffixes and geminate versus non-geminate intervening clusters. Noonan (1992) notes examples such as [wèlò-ná] 'visitor-1SG', where the stem vowel [ò] is kept and progressive harmony does not occur, versus the same form with deletion by contraction of the stem vowel, [wèl-lə́] 'visitor-1SG', where cluster simplification creates a geminate out of the root-final and suffix consonants, and harmony does occur. Similarly, he notes [cwíɲ-wá] 'liver-1PL', with a non-geminate cluster and no harmony, versus [cwíɲ-ɲə́] 'liver-1SG', where the suffix consonant undergoes cluster simplification to create a geminate, and harmony does occur. The former fact could result from the presence of the stem vowel [-o] initiating a new harmony domain; Noonan (1992: 34) describes it as harmony not happening 'because there is a stem vowel', and stem vowels also fail to trigger harmony regressively onto the root, as in [wálô] 'to boil (intr.)' (compare the ventive infinitive and plural in the previous paragraph). In terms of the cluster versus geminate distinction, this could be analyzed as the relevant intervening unit being defined in terms of consonants rather than moras, where geminates count as a single consonant unit; progressive harmony can cross one consonant slot, but not two.

6.2 Dominance in ATR harmony

In the present analysis, ATR dominance is derived by having IDENT-IO[ATR] outweigh IDENT-IO[RTR]; ATR vowels will always remain faithful, while RTR ones can change in order to satisfy harmony. This view is consistent with the idea that the ‘problem’ solved by harmony is disagreeing sequences, and that the markedness constraint that drives harmony can penalize this problem but not dictate a particular solution. Instead, as is standard in OT, the reason the ATR value wins in solving the disagreement needs to be determined by faithfulness.

This version of accounting for ATR dominance raises some puzzling questions surrounding other properties of such systems. Specifically, as discussed earlier (see Section 3), ATR dominant harmony is highly correlated with the presence of a tongue root contrast among high vowels (e.g. Casali 2003; 2008), and with preferential triggering by high vowels (e.g. Rose 2018). Given these facts (at least about African ATR systems)[87], it is clear that ATR dominance is more complex than is suggested by the ranking IDENT-IO[ATR] >> IDENT-IO[RTR]. Specifically, in OT, the inventory of high vowels is derived by ranking RTR faithfulness above a constraint against RTR high vowels, IDENT-IO[RTR] >> *ɪʊ. There is no obvious reason why these two rankings would be related: nothing should prevent a language from setting the ranking IDENT-IO[RTR] >> IDENT-IO[ATR], *ɪʊ, creating a 2IU language with RTR-dominant harmony, yet such systems are extremely rare, if not unattested, according to typological surveys.[88] It seems that somehow, the ranking that permits high RTR vowels to surface must also enforce higher faithfulness to ATR than to RTR.

Beyond that, whether a language has ATR dominance also seems to affect which vowels behave as the best triggers and targets of the harmony system. In ATR dominant systems, it is evident that high vowels are the best triggers; they consistently trigger harmony in these cases, even if (contrastive) vowels of other height do not, like in some of the sub-patterns of Lango. High vowels are arguably also the best targets of these systems, in that languages often have sub-

---

[87] It is unclear what is the case for non-African ATR systems like in Tungusic (Li 1996), since Casali focuses only on African systems.

[88] Again, however, the surveys only focus on African languages.

patterns in which high vowels are also the only targets, as also occurs in Lango. Languages with non-ATR-dominant tongue root harmony, on the other hand, clearly prefer mid vowels as both triggers and targets; allophonic high vowels created through harmony are rare (though they do occur in some cases of RTR harmony, such as certain dialects of Yoruba; see Rose 2018), and harmony often occurs among mid vowels to the exclusion of other vowels.

What this distinction means is that there is yet another place in the grammar where ATR dominance is encoded: in the scaling factors. The fact of having ATR dominance means that high vowels become the best triggers and targets, where otherwise that status belongs to the mid vowels. Again, this fact is somewhat problematic: there is no obvious reason that the scaling factors should be able to 'see' the ranking of the faithfulness constraints in order to establish whether a system is ATR dominant. Like with ATR dominance and the presence of contrastive RTR high vowels, there is no clear reason in the grammar for these factors to be at all related.

These relationships between multiple parts of the analysis are clearly non-ideal; an ideal account would more accurately reflect the attested typology by allowing for some sort of relationship between the inventory, the ATR dominance, and the trigger and target preferences. However, it is worth noting several relevant facts before dismissing an analysis that does not accomplish this.

First, a great deal remains unknown about the typology of tongue root harmony, and of vowel harmony more generally. It is well-known that it is difficult to distinguish RTR high vowels from ATR mid vowels (e.g. Casali 2003; Hyman 1999; Rose 2018); as such, recent phonetic technology can result in reanalysis of the vowel inventories that have previously been described, including the number of contrasts. As such, it is not entirely clear to what extent we can rely on the existing typological generalizations. This is especially true with respect to claims about certain non-contrastive vowels being unaffected by harmony (as opposed to being allophonically affected), and so typological claims and arguments about that type of property are particularly vulnerable. Indeed, instances of supposedly unaffected vowels have been found under acoustic analysis to be allophonically subject to harmony, as in Kinande (Gick et al. 2006).

Moreover, there are obvious typological gaps across types of harmony, suggesting that we cannot provide any sort of consistent analysis of all harmony systems that does not make unusual typological predictions for at least some types of harmony. Notably, the type of dominant/recessive system that is common for tongue root harmony is apparently rare for other types of vowel harmony like backness or rounding (e.g. Casali 2008). For instance, there do not appear to be languages in which rounding is ‘dominant’ in the same way that ATR can be, in the sense where both root-to-suffix and suffix-to-root harmony occurs, depending on whether the ‘dominant’ feature value is in the root or the affix. This typological gap exists despite the fact that many Turkic languages with rounding harmony have affixes that are consistently round (e.g. the second vowel in the Turkish suffix [-edur]~[-adur]; Clements & Sezer 1982), which we might expect to trigger harmony onto roots in the same way that roots can trigger rounding onto suffixes. Thus, even allowing ATR dominant harmony to exist, as we need to, predicts systems that are unattested for vowel harmony in all other features. There is clearly more work to be done on understanding how harmony systems work and why before we can actually use the attested typology to determine whether an analysis makes sense.

Additionally, capturing the relationship between a contrast in high vowels and ATR dominance would create difficulties in any OT analysis. In OT, having a contrast is based on a ranking in which faithfulness to the given feature dominates a relevant markedness constraint. Specifically here, we need to ensure that RTR is maintained on high vowels. ATR dominance, on the other hand, needs to be encoded into which value of the feature gets preserved in harmony contexts, here ATR; harmony constraints are markedness constraints that cannot specify a solution. Maintaining RTR on high vowels in general and maintaining ATR on any vowel in disharmonic contexts cannot be connected in any obvious way, and no existing analysis is able to enforce this generalization formally.

6.3 Targets in ATR harmony (revisited)

As noted in Section 3.2, some literature (e.g. Rose 2018) claims that mid vowels are the best targets of ATR harmony, including ATR dominant harmony. This idea has been justified both phonetically and typologically. In terms of the former, the tongue root contrast in mid vowels has been argued to be acoustically and perceptually more distinct, compared to the contrast in high and low vowels (e.g. Fulop et al. 1998; Starwalt 2008; Becker-Kristal 2010). Given arguments about why harmony is phonologized, ‘good’ targets of harmony should be those that have a strong perceptual contrast for the relevant feature. Thus, phonetically, we expect mid vowels to be good targets. In terms of the typological evidence, most languages with a tongue root contrast in high but not mid vowels create allophonic mid vowels through harmony, while the reverse is quite rare (e.g. Casali 2003; 2008; Rose 2018). This fact suggests a greater drive for mid vowels to undergo harmony than for high vowels.

There are a number of questions that are raised by this idea. While mid vowels are quite clearly ideal participators in non-ATR dominant harmony (i.e. RTR dominant or root-controlled), as is the case in Yoruba, the same is difficult to establish in ATR dominant systems. While it is true that most such systems harmonize mid vowels even when ATR mid vowels do not exist contrastively in the inventory, Casali’s survey does not note systems that target mid vowels to the exclusion of vowels of other heights. In contrast, in a number of systems, there are at least some contexts in which high vowels are targeted to the exclusion of non-high vowels; such is the case in Lango for regressive harmony across multiple consonants triggered by [u], and in Kinande progressive harmony (e.g. Archangeli & Pulleyblank 2002). In the typology of ATR dominant harmony in general, there appears to be a height effect in the targets in any given context: they can be only high, only non-low, or all vowels, not any other combination. This effect is unexpected if the point of harmony is to extend the ATR feature to where it is most perceptible, and if it is most perceptible on mid vowels.

There are a few possibilities to consider. First, as mentioned before, we may simply not know enough about tongue root harmony; there may be undocumented or poorly documented cases in which mid vowels are in fact the only targets. Similarly, there may be systematic gaps in the typology, in the way that dominant/recessive systems seem to only exist for tongue root harmony. Here, it would be that there is a systematic gap in ATR dominant harmony targeting only the 'best' targets of tongue root harmony. In this particular situation, such a claim seems non-ideal, because it is in conflict with the phonetic motivations claimed for harmony systems. An alternative would be to claim that all tongue root harmony systems prefer targets that share features with the triggers. In this case, since high vowels (which are [+high, -low]) are the best triggers of ATR dominant harmony, the height effect could be reduced to the number of height features shared between target and trigger: [+high, -low], [-low], or none. Harmony preferentially occurs with a target that shares more features with the trigger. Such an approach, however, relies on either a binary feature system for height or reference to absence of features, which are not universally accepted.

An interesting effect of such an analysis is that it could be extended to understanding why non-ATR dominant systems like Yoruba typically preferentially involve mid vowels, both as triggers and targets. In these cases, what is strange on the surface is that mid vowels are also the preferential triggers, despite the contrast being apparently easily perceptible in these vowels. Here, we could argue that harmony is based on the target: mid vowels are the best targets of tongue root harmony, and the trigger should agree with it in height. The idea that harmony could be driven by particular vowels needing to be targets of harmony is a bit unusual, but worth exploring in understanding this difference in triggers and targets of distinct types of tongue root harmony. It is also worth noting that while almost all languages with an ATR/RTR contrast in high vowels have harmony, a much smaller percentage of languages without this contrast have harmony (Rose 2018). This fact is suggestive of there being less of a drive for harmony in the latter type of language, which could be expected if harmony in such cases is driven by the target rather than the trigger.

Regardless of the specific analysis, these distinctions make it difficult to capture all of the typological generalizations in any theoretical analysis. What we need to allow for ATR dominant harmony is almost directly in opposition from what we need to allow for other types of tongue root harmony, and no matter what, there will be unwanted typological predictions.

7 Conclusion

In conclusion, the basic facts of Lango ATR harmony that other analyses have accounted for become quite straightforward in an approach in which the strength of the harmony constraint, and therefore the penalty for disharmony, depends on the interaction of triggers, targets, and distance. Previous work on other languages has argued for the necessity of such factors incorporated into the harmony constraint, with phonetic properties used to determine the relative strength of particular vowels as triggers and targets. Applied to Lango, this approach reduces the complex pattern to an interaction of harmony and faithfulness. Critically, it is able to maintain both restrictiveness and generality missing from many of the previous accounts, while not relying on complex representational notions like a non-standard definition of head, which is difficult, if not impossible, to incorporate formally into a theoretical account.

Chapter 9: Neutral paired vowels in Mayak

1 Introduction

In many languages with tongue root harmony, the lack of an ATR low vowel in the inventory has allowed explanations of the neutrality of [a] to invoke general markedness against its harmonic counterpart (e.g. Lo/TR in Pulleyblank et al. 1995; *[+lo,+ATR] in Baković 2000). However, this analysis is not available for all cases of neutral [a] in tongue root harmony; in some languages, vowels are idiosyncratically neutral to harmony, despite having harmonic counterparts that are allowed to occur in the given context. This chapter considers such a case in Mayak (Andersen 1999), a Western Nilotic language from the Burun subfamily, spoken in Sudan. Crucially, in Mayak, there is a contrastive ATR low vowel [ʌ] that is permitted in the contexts in which the RTR low vowel [a] fails to undergo ATR harmony; [ʌ] also sometimes fails to trigger ATR harmony. This pattern is in direct opposition to what happens with Mayak mid vowels, which participate in (regressive) harmony (in the root, not in suffixes) despite not contrasting for ATR. Mayak therefore raises critical questions about how we account for neutrality and participation in ATR harmony, and the factors that establish which vowels participate. In this chapter, I adopt the same type of analysis used in Chapter 8 for Lango, but argue that instead of distinguishing high from non-high vowels, Mayak draws the distinction between non-low and low vowels. This approach is able to capture the Mayak pattern, where any account based on the relationship between contrast and neutrality fails.

2 Mayak vowel harmony and neutrality

2.1 Inventory

The vowel inventory of Mayak, from Andersen (1999) is given in Table 9.1. The facts relevant to harmony are symmetric across short and long vowels, so for notational ease, I mention only the short vowels in the prose about the facts (but will use examples with both short and long vowels). The mid ATR vowels [e] and [o] are in parentheses because Andersen (1999) argues that these vowels are not underlying, but rather occur only as the result of ATR harmony targeting /ɛ/ and /ɔ/. Crucially, the surface vowel inventory of Mayak is symmetrical in the tongue root dimension; each vowel has a counterpart that differs from it only in ATR/RTR. Nonetheless, the contrastive inventory is asymmetric, with high and low vowels paired for tongue root position, but no contrast in mid vowels.

	ATR			RTR		
	Front	Central	Back	Front	Central	Back
High	i, i:		u, u:	ɪ, ɪ:		ʊ, ʊ:
Mid	(e, e:)		(o, o:)	ɛ, ɛ:		ɔ, ɔ:
Low		ʌ, ʌ:			a, a:	

Table 9.1: Vowel inventory of Mayak (Andersen 1999)

2.2 Regressive harmony

Mayak has regressive ATR harmony triggered by the high vowels [i] and [u] and targeting the non-low RTR vowels, so that [ɪ, ʊ, ɛ, ɔ] never occur before [i, u] in the language. Examples are shown in (1), where the past tense suffix [-u] triggers an alternation in the root vowel. All examples in this paper come from Andersen (1999).

(1) Alternation of non-low RTR root vowels before a suffix [u]

	Underlying root vowel	Present tense	Past tense	Gloss
(a)	ɪ	ʔɪt̪	ʔið-u	‘shape with an axe’
(b)	ɛ	lɛp	lew-u	‘open’
(c)	ɔ	kɔc	koj-u	‘take’
(d)	ʊ	gʊt̪	guð-u	‘untie’

In contrast, the low RTR vowel [a] is unaffected by this process, surfacing as RTR even in the past tense, before the ATR [-u]. Examples are given in (2).

(2) Non-alternation of low RTR root vowels before a suffix [u]

	Underlying root vowel	Present tense	Past tense	Gloss
(a)	a	ʔam	ʔam-u	‘eat’
(b)	aa	maat̪	maað-u	‘drink’

The underlying ATR vowels are all allowed to surface unchanged before [i, u], as shown in (3). Given later data in (10) that will establish that [ʌ] is indeed the ATR counterpart of [a], the example of [ʌ] in (3b) is crucial, because it establishes that the ATR counterpart of [a] is permitted in the same context in which [a] fails to undergo harmony in (2); there is no reason to suspect that consonantal context in any way influences the tongue root quality of the vowel.

(3) Non-alternation of ATR root vowels before a suffix [u]

	Underlying root vowel	Present tense	Past tense	Gloss
(a)	i	ʈiŋ	ʈiŋ-u	‘hear’
(b)	ʌ	nʌk	nʌɣ-u	‘beat’
(c)	u	t̪uc	t̪uj-u	‘send’

As noted, the mid ATR vowels [e] and [o] are not contrastive, in that they occur only as the result of harmony, and so there cannot be examples in which they trigger harmony. In contrast, the low ATR vowel [ʌ] does occur underlyingly, yet cannot trigger harmony. RTR vowels are also unable to trigger harmony, as harmony is ATR dominant. The verbal paradigm examples in (4) illustrate the contrast between the 2SG suffix with /i/, which does trigger harmony, and the 1SG and 3SG suffixes with /ʌ/ and /ɛ/ respectively, which do not. Crucially, [ʌ] fails to trigger harmony in exactly the same context in which [i] triggers it, even though [ʌ] is ATR; the result is surface disharmony in the 1SG.

(4) Root vowel alternations triggered by /i/ but not /ʌ/ or /ɛ/

	Underlying root vowel	1SG	2SG	3SG	Gloss
(a)	ɪ	ʔɪd̪-ʌr	ʔid̪-ir	ʔɪd̪-ɛr	‘shape with axe’
(b)	ɔ	nɔn-ɖ-ʌr	non-ɖ-ir	nɔn-ɖ-ɛr	‘fold’
(c)	aa	caab-ʌr	caab-ir	caab-ɛr	‘cook’

2.3 Progressive harmony

Mayak also has a more limited system of progressive ATR harmony, primarily illustrated in suffixes that alternate between [i] and [ɪ] depending on the ATR value of the stem vowel. As for regressive harmony, only [i] and [u] are triggers; [ʌ] is not. Examples are shown in (5). Note how with RTR root vowels, in (a-c), the suffix surfaces as the RTR [ɪ]; with high ATR root vowels in (d-e), the suffix is ATR [i]. Critically, with the low ATR root vowel [ʌ], the suffix is RTR [ɪ] in (f), resulting in disharmony. Recall that there are no mid ATR root vowels (except as the result of harmony).

(5) Alternation of suffix [i]~[ɪ] triggered by high ATR root vowels

	Underlying root vowel	Non-possessed	1SG	Gloss
(a)	a	pal	pal-ɪ	‘navel’
(b)	ɔ	wɔŋ	wɔŋ-ɪ	‘eye’
(c)	ʊ	t̪ʊk	t̪ʊɣ-ɪ	‘outer mouth’
(d)	i	ʔic	ʔiɗ-i	‘ear’
(e)	u	ʔuŋ	ʔuŋ-i	‘knee’
(f)	ʌ	ʔʌm	ʔʌm-ɪ	‘thigh’

Andersen (1999) notes that the behaviour of /ʊ/ in suffixes has not been systematically studied, and that he was unable to establish whether there exist such suffixes, but he hypothesizes that they would behave like suffixes with /ɪ/. Nonetheless, he does discuss a suffix that alternates [-ʊk]~[-uk]; he does not count this as underlying /ʊ/, because it harmonizes with all preceding vowels, including the low [ʌ]. The data in (6) show the alternation.

(6) Alternation of suffix [ʊ]

	Underlying root vowel	Plural	Gloss
(a)	ɪ	mɪɣ-ʊk	‘spider-PL’
(b)	ʊ	gʊj-ʊk	‘bowl-PL’
(c)	i	cim-uk	‘knife-PL’
(d)	ʌ	jʌŋ-uk	‘crocodile-PL’
(e)	u	bul-uk	‘stomach-PL’

Note that this is the only suffix with [ʊ] (or alternating [u]) discussed by Andersen (1999), meaning that it is not clear whether the fact that it harmonizes with [ʌ] is an idiosyncratic/exceptional aspect of this particular morpheme, or whether it holds true of /ʊ/ in general. Andersen (1999) suggests treating this suffix morpheme as exceptional, and uses underspecification (lack of underlying RTR, unlike in the /ɪ/ suffixes in (5)) to encode this. I will take up the issue of progressive harmony from [ʌ] later on.

In contrast to these high vowel suffixes, suffixes with [ɛ] and [ɔ] do not alternate, nor does the plural suffix [-ak]. Examples are shown in (7-9). Note in particular how the suffix is RTR after ATR root vowels in (6d-f), (8c), and (9c-d), resulting in disharmonic forms.

(7) Non-alternation of suffix [ɛ]

	Underlying root vowel	Non-possessed	3SG	Gloss
(a)	a	pal	pal-ɛ	‘navel’
(b)	ɔ	wɔŋ	wɔŋ-ɛ	‘eye’

(c)	ʊ	t̪ʊk	t̪ʊɣ-ɛ	‘outer mouth’
(d)	i	ʔic	ʔiɗ-ɛ	‘ear’
(e)	u	ʔuŋ	ʔuŋ-ɛ	‘knee’
(f)	ʌ	ʔʌm	ʔʌm-ɛ	‘thigh’

(8) Non-alternation of suffix [ɔ]

	Underlying root vowel	Unmarked	1PEX	Gloss
(a)	ɛ	gɛp	gɛw-ɔnɔn	‘beat’
(b)	ɔ	nɔn	nɔn-ɔnɔn	‘fold’
(c)	i	ʔip	ʔiw-ɔnɔn	‘shoot’

(9) Non-alternation of suffix [a]

	Underlying root vowel	Singular	Plural	Gloss
(a)	ɪ	bɪl	bɪl-ak	‘iron’
(b)	ɛ	t̪ɛl	t̪ɛl-ak	‘lower leg’
(c)	i	kic	kij-ak	‘bee’
(d)	u	kuʈ	kuɗ-ak	‘nest’

2.4 Exceptional suffix

Critically for the present purposes, there is one low vowel suffix that does alternate depending on the ATR value of the stem vowel, with [a] after RTR roots and [ʌ] after ATR roots. Examples are given in (10). In (a-c), the root vowel is RTR, and the singular suffix surfaces with [a]; in (d-f), the root vowel is ATR, and the singular suffix surfaces with [ʌ]. Notably, this is true even in

(10e), where the root vowel is [ʌ], which typically does not trigger harmony, even when other ATR vowels do (cf. the 1SG column in (4) and example (5f)). Note that there are two suffixes that contain [a], the non-alternating one in (9) and the alternating one in (10); the suffix that I call exceptional, following Andersen's (1999) description of it as underspecified, is one of these two. With respect to [ʌ] triggering harmony, note that it never triggers regressive harmony onto any RTR vowels (cf. (4) above), and it does not trigger progressive harmony onto suffixal [ɪ, ɛ, ɔ], nor onto the other suffix [a], as seen in (5) and (7–9). The only places where the ATR low vowel does trigger harmony are onto the suffix in (10) and the one in (6).

(10) Suffix alternation [a]~[ʌ] based on the stem ATR value

	Underlying root vowel	Singular	Plural	Gloss
(a)	ɪ	rɪm-at̪	rɪm	'blood'
(b)	a	ɗaal-at̪	ɗaal	'flower'
(c)	ʊ	kʊm-at̪	kʊm	'egg'
(d)	i	ʔin-ʌt̪	ʔin	'intestine'
(e)	ʌ	ʔʌʌw-ʌt̪	ʔʌʌp	'bone'
(f)	u	ruuj-ʌt̪	ruuc	'worm'

The behaviour of this particular suffix is crucial because it illustrates that [ʌ] and [a] do appear to constitute an ATR/RTR pair for the purposes of the Mayak harmony system, even though [ʌ] typically does not trigger ATR harmony and [a] typically does not undergo it.[89] Andersen (1999) assumes that the suffix in (10) is underspecified for ATR and therefore not subject to ATR

---

[89] Of course, it is possible that this is just allomorphy, but given the phonetic description of the vowels as low and differing only in ATR/RTR, as well as the fact that ATR in the stem conditions the choice of allomorph, it is a better explanation to claim that this is exceptional ATR harmony, with the two low vowels constituting a pair.

faithfulness, allowing it to alternate even though low vowels are generally immune to harmony. Since there is a three-way suffix contrast (non-alternating [a], as in [kij-ak] ‘bee-PL’; non-alternating [ʌ], as in [nɔn-ɖ-ʌr] ‘fold-1SG’; and alternating [a]~[ʌ], as in [rɪm-at̪] ‘blood-SG’ versus [ʔin-ʌt̪] ‘intestine-SG’), an analysis in which this suffix is not in some way exceptional, either through underspecification or lexical indexation, fails to explain why the other suffix with [a] fails to alternate (see (9) above), as well as why [ɛ, ɔ] suffixes fail to alternate, and why alternations involving suffix /ɪ/ behave differently, in that [ʌ] does not trigger harmony to /ɪ/.

2.5 Summary

To summarize, the Mayak low vowels [ʌ] and [a] are demonstrably paired in the feature ATR, yet with the exception of one suffix (as a target, /-At̪/; two in terms of triggers, since /ʌ/ can cause both the /-At̪/ suffix and the /-Uk/ suffix to surface as ATR), they behave neutrally. [ʌ] does not trigger ATR harmony in many contexts where the other underlying ATR vowels [i] and [u] do (see (4) and (5)), and [a] generally does not undergo ATR harmony in the contexts where all other RTR vowels do (see (1) versus (2)). This pattern is most evident in suffix-to-root harmony; suffix [ʌ] is never a trigger, while suffix [i] and [u] are, and root [a] is never a target, while non-low RTR vowels are targets in roots. Root-to-suffix harmony is more complex. Mid RTR vowels are never targets in suffixes, and high RTR vowels always are. Andersen (1999) provides one suffix [-ak] that is not a harmony target, and another suffix [-at̪] that is. Further, root [ʌ] triggers harmony to /-at̪/ and to /-ʊk/, but not to any of three suffixes with /ɪ/ that are targets of harmony from high ATR triggers in the root. As such, there might be a generalization that root [ʌ] can trigger harmony only to phonemic back vowels. One reason for labelling [-at̪] and [-ʊk] as exceptions, instead of aspects of a general pattern, is that regressive harmony in

Mayak operates in such a way that [a] is a non-target and [ʌ] is a non-trigger, even to back vowels.

Indeed, even with the ‘exceptional’ suffixes taken into account, there are strong arguments for the need for some formal mechanism for differentiating between low and non-low vowels as potential targets and potential triggers. [ʌ] is more neutral than [i,u] as a potential trigger because it never triggers regressive harmony, and only triggers regressive harmony onto [-at̪] (one of two suffixes with [a]) and [-ʊk] (the only suffix with [ʊ]), but not onto any of the three suffixes with /ɪ/. In contrast, [i] and [u] do trigger harmony onto all suffixes with /ɪ/, so regardless of whether [-at̪] and [-ʊk] are exceptional, [ʌ] is a weaker trigger than [i,u]. In terms of targets, [a] is more neutral than [ɪ,ʊ,ɛ,ɔ] as a potential target, because /a/ never undergoes regressive harmony, whereas /ɪ,ʊ,ɛ,ɔ/ all do. With respect to progressive harmony, one of two /a/ suffixes never undergoes harmony from any ATR vowel; this corresponds to the behaviour of suffix /ɛ,ɔ/, whereas suffix /ɪ,ʊ/ do undergo harmony, such that there is a high/non-high distinction in target behaviour. The other /a/ suffix patterns like the one suffix with /ʊ/, in undergoing harmony from all ATR vowels. Notably, while this means that in progressive harmony, /a/ could in theory be considered intermediate in undergoing harmony between /ɪ,ʊ/ and /ɛ,ɔ/, it is worth noting that /ɛ,ɔ/ do not have contrastive ATR counterparts, while /a/ does.

The point about differential neutrality in targets is most crucial, because while trigger-specific conditions are commonly invoked in analyses of harmony (e.g. trigger scaling in Kimper 2011; alignment constraints specified to trigger features in Kaun 1995), failure to undergo harmony is most often ascribed to the absence of a harmonic counterpart (e.g. Ringen & Vago 1998; Baković 2000; Kiparsky & Pajusalu 2003; etc.). This view is adopted not only in front/back harmony systems, from which they are best known, but also in many analyses of why low vowels are neutral to tongue root harmony in other languages. However, in Mayak, the harmonic counterpart of [a], namely [ʌ], is present in the inventory and permissible in the same contexts in which [a] fails to undergo harmony.

One way to look at harmony in Mayak is as being parasitic. Progressive ATR harmony in Mayak generally involves only high vowels as triggers and targets, so it could be said to be generally parasitic on the feature [+high], with the exception of the [at̪]~[ʌt̪] alternation. Regressive harmony, on the other hand, targets all non-low vowels, and so could be said to be parasitic on the feature [-low]. These are simply different ways of stating the trigger and target conditions, such as the fact that for regressive harmony, every RTR vowel except for [a] is required to be a target. The failure of [a] to undergo harmony is specifically be about the suitability of [a], compared to non-low vowels, as a harmony target. In particular, this fact is not about contrast. Indeed, the fact that mid vowels do undergo regressive harmony, even though ATR mid vowels (unlike ATR low vowels) do not occur independently in the language, is particularly strong evidence that the pattern here is not about contrast.

While Mayak is uncommon in that its neutral [a] is paired, neutral (unpaired) low vowels are in fact quite frequent among ATR harmony systems, as we saw in Alur (Chapter 7). Indeed, [a] appears to be a poor target for ATR harmony; it does have an allophonic counterpart in some languages like Akan, and re-pairs in others like Alur, but it is also often neutral (transparent or opaque), as in Ngiti (Casali 2008); these patterns are all attested in 9-vowel ATR systems, so the pattern cannot be determined by the inventory. In contrast, mid vowels typically participate allophonically in ATR harmony, even when unpaired (Casali 2003; 2008; Rose 2018). As such, any analysis should predict languages like Mayak, where a paired low vowel is neutral to ATR harmony, but not languages in which paired low vowels participate but paired mid vowels are neutral.

3 Problems for a view where neutrality is equivalent to unpairedness

A traditional view of neutral low vowels in ATR harmony relies on the markedness of ATR low vowels. For example, [a] would be neutral due to a constraint against ATR low vowels (e.g. Lo/TR in Pulleyblank et al. 1995; *[+lo,+ATR] in Baković 2000). In Mayak, however, the ATR low vowel [ʌ] is permitted in the contexts in which [a] fails to undergo harmony, indicating that markedness alone cannot account for the pattern.

One possibility is that the Mayak pattern is a derived environment effect, in that [ʌ] can only surface if it is underlying and therefore cannot appear as the result of harmony. This idea is similar to that of Baković (2000), who uses conjunctions of (context-free) markedness constraints with faithfulness to [ATR]; this approach generates this sort of derived environment effect, in that the conjunction is only violated by a marked vowel if that vowel is unfaithful (derived), as by harmony, but not if it is preserving an underlying specification. Note, though, that Baković used this type of machinery primarily to avoid a 'majority rules' pathology, not for this type of effect.

However, this option is undesirable in Mayak; the examples in (9) show that [ʌ] is permitted as the result of harmony. Indeed, assuming that it is underlying in the [ʌt̪]~[at̪] alternation would substantially complicate the analysis, given non-alternating suffixes with [ʌ] as in (4). If [ʌ] were permitted only when underlying, which requires a ranking of IDENT-IO[ATR] >> *ʌ >> ATR-HARMONY, then we would predict underlying /ʌ/ to surface faithfully, and any underlyingly unspecified low vowels to surface as [a]. The tableaux in Tables 9.2-9.3 illustrate.

/ʌ/	IDENT-IO[ATR]	*ʌ	ATR-HARMONY
☞ʌ		*	
a	*!		

Table 9.2: Underlying /ʌ/ surfaces faithfully

/A/	IDENT-IO[ATR]	*ʌ	ATR-HARMONY
ʌ		*!	
☞a			

Table 9.3: Underlying unspecified /A/ surfaces as [a], regardless of harmony

Both options are inconsistent with the alternation in (9). Even following Andersen's (1999) assumption that the alternating low vowel suffix is underspecified, both [ʌ] derived from underlying /a/, which does not occur, and [ʌ] derived from underlying /A/, which is allowed, are equally derived; a derived environment analysis would have difficulty capturing that distinction in behaviour. Moreover, while we could potentially say that this exceptional suffix is simply lexically specified to undergo harmony (as in, indexed to be subject to a higher-ranked harmony constraint), and therefore allow it to have an underlyingly specified vowel, the choice most consistent with the rest of the data would be /a/. Indeed, harmony is consistently ATR dominant in Mayak, triggered only by the ATR vowels [i, u]; requiring an instance of RTR dominance simply to assume underlying /ʌ/ for a single exceptional affix would add unnecessary complexity to the analysis. Thus, it is unlikely that the behaviour of low vowels is a straightforward derived environment effect, as doing so would result in many complications for the analysis of the exceptional suffix.

As such, an account for Mayak in which neutrality arises only from general markedness is impossible, or at the very least, strongly dispreferred. In fact, I argue that markedness is relevant to Mayak only in the sense that the factors that cause ATR low vowels to be cross-linguistically dispreferred are the same as those that make them poor targets of harmony.

One way of considering the Mayak pattern in traditional terms is as a case of parasitic harmony, in that harmony operates only on vowels that are both [-low] (for regressive harmony). However, the actual label of 'parasitic' is not particularly meaningful in itself; what is critical is about the analytical tools and whether something that looks like parasitic harmony can be derived. The present analysis, while not based on agreement, succeeds in deriving this pattern. In terms of analyses that are actually based on agreement, a major one in vowel harmony is for rounding harmony (Kaun 1995; 2004). In that case, Kaun (1995; 2004) has argued that the reason for the parasitic pattern is uniformity of the phonetic (articulatory) realization of the feature; specifically, the gesture of lip rounding differs on high and non-high vowels. Kaun (1995; 2004) suggests that for languages in which rounding harmony occurs only among vowels of the same height, it is because harmony endeavours to extend a single type of lip rounding gesture. In contrast, as far as the literature suggests, there is no such categorical phonetic distinction between low and non-low vowels in the realization of ATR. At the very least, then, this distinction would not be like the parasitic harmony in cases of rounding harmony. Consonant harmony is also often parasitic on other features, such as sibilant harmony parasitic on manner or voicing; the Agreement by Correspondence framework (ABC; Rose & Walker 2004) was designed to deal with such cases. However, again we might expect, for instance, that realization of sibilant distinctions is categorically different depending on manner. Evidently, deriving parasitic patterns by separating trigger and target asymmetries can capture more diverse patterns than an ABC-style analysis requiring trigger and target to agree (e.g. cases where only trigger or only target is relevant). Further argumentation about this analysis versus one in which harmony operates only on [-low] vowels in an ABC way is beyond the scope of the present paper.

Additionally, other languages with ATR dominant harmony draw a distinction between participating and neutral vowels at a different point in the high-mid-low hierarchy. For instance, another Western Nilotic language Lango (Chapter 8) differentiates high versus non-high targets, but does not distinguish low from non-low. In other words, mid vowels can either pattern with high vowels or with low vowels. As such, if we were to argue that there is parasitic harmony due to a requirement for uniformity of feature realization, we would either need to claim that these related languages differ in the realization of ATR on mid vowels, or else that there is a gradient scale of how ATR is realized, and languages can choose to make the distinction in different places. The latter approach is simply an unnecessarily complex version of the idea argued here, that there is a gradient scale of target quality, and languages can choose different 'cut-off' locations for which vowels actually behave as targets.

4 Analysis of Mayak

4.1 Basics of analysis

Unlike in other types of analyses, a lack of highly weighted/ranked general markedness constraints does not preclude neutrality in the target-oriented approach advocated for in this dissertation. Specifically, a language can permit ATR low vowels, yet weight faithfulness and scale the harmony constraint such that non-low vowels are required to harmonize, but low vowels are allowed to be disharmonic. This type of analysis applied to Mayak can successfully capture the fact that low vowels are neutral to ATR harmony despite being paired, in a way that incorporates the cross-linguistic tendency for low vowels to be poor participators in ATR harmony.

In addition to the neutrality of the paired low vowels, an analysis of Mayak must account for the fact that all mid vowels and high vowels are targets of regressive ATR harmony, but only high vowels are targets of progressive ATR harmony. There are two ways in which this difference could be captured: the harmony-driving constraints for regressive versus progressive harmony could be weighted differently, or the scaling factors could differ. I adopt the former approach here, with the assumption that a language can set scaling factors only once per type of harmony constraint, and these will be consistent for both directionalities.[90] With regressive harmony weighted higher than progressive harmony, an asymmetry can emerge in terms of which vowels are targets of each directionality. Given that cross-linguistic evidence suggests a regressive bias in vowel harmony (e.g. Hyman 2002), there is independent reason to believe that languages might prefer to weight regressive harmony higher than progressive harmony. In contrast, beyond this regressive bias, there is no principled reason for vowels to differ in their suitability as a target depending on their position relative to the trigger.[91] Reducing directionality-based target asymmetries to the weighting of regressive versus progressive harmony is another advantage of this analysis; such asymmetries can often require more complex theoretical machinery to capture, for instance by having languages set target conditions separately for each directionality, as is done for example in the grounded conditions in

---

[90] Note, however, that a language could set scaling factors differently for different types of harmony (e.g. ATR versus rounding), or for harmony of the same feature in which the trigger/target (here, context/focus) roles have been switched. For example, the constraints adopted here, defined in (12), have ATR defined as the context, and therefore trigger, as in *[RTR]∞[ATR], where the underline indicates focus, or 'target' for the purposes of scaling. The corresponding constraint *[RTR]∞[ATR], where ATR is the focus, would not necessarily need to be scaled the same way; an ATR vowel would not necessarily be scaled the same way for disharmony when it is a 'target' as when it is a 'trigger'. This is potentially important given that ATR-dominant harmony, which I presume to have constraints in which RTR is specified as the focus, patterns differently in terms of triggers and targets compared to other types of tongue root harmony.

[91] Of course, there can be differences in suitability depending on morpheme affiliation; in Mayak, for instance, the regressive/progressive differences align with root/suffix distinctions.

Archangeli & Pulleyblank (1994).

4.2 Mid vowels created only through harmony

In Mayak, like in a number of other languages with a tongue root contrast in high vowels (2IU languages; Casali 2008), there is no contrast in mid vowels (1EO), yet ATR mid vowels are created allophonically through harmony. In fact, based on surveys of many 2IU, 1EO languages (Casali 2003; 2008; Rose 2018), these systems typically have allophonic ATR mid vowels as a product of harmony. Moreover, there is a very strong correlation between 2IU in general and ATR dominance, including in 2IU, 1EO languages. The result is an apparent conceptual and analytical paradox: why would a language preserve ATR in determining the result of harmony, but not in determining its contrastive inventory? In Mayak, for instance, the ATR of the high vowel [u] is preserved in the form [lew-u] 'open-PST', creating the ATR mid vowel [e] from the RTR [ɛ] in the unsuffixed root [lɛp] 'open.PRES'. However, [e] and [ɛ] are not contrastive, and [e] would not be preserved if it occurred underlyingly. This situation is odd, as we might expect instead to see RTR harmony, with /lɛp + u/ surfacing as *[lɛwʊ], in order to avoid creating ATR mid vowels.

To see where this problem arises in a basic constraint-based conceptualization, consider the rankings required to obtain each aspect of such a pattern: ATR dominance, lack of contrast in the mid vowels, and harmony. For ATR to be dominant, it must be more important to preserve ATR than RTR, which could be achieved through a ranking like IDENT-IO[ATR] >> IDENT-IO[RTR] (see (10) for constraint definitions). To avoid a contrast in mid vowels, the constraint against ATR mid vowels must outrank the faithfulness that preserves them, so *eo >> IDENT-IO[ATR]. Thus far, then, we need *eo >> IDENT-IO[ATR] >> IDENT-IO[RTR]. However, with this ranking, no matter where the harmony-driving constraint appears, ATR mid vowels will never be created as the result of harmony. Even if it is the highest-ranked constraint, HARMONY >> *eo >> IDENT-IO[ATR] >> IDENT-IO[RTR], we predict that a combination like /ɛ…i/ will harmonize to [ɛ…ɪ], not [e…i], because avoiding [e] is more important than preserving ATR. This result is shown in Table 9.4.

/ɛ…i/	HARMONY	*eo	IDENT-IO[ATR]	IDENT-IO[RTR]
☞ɛ…ɪ			*	
☹ e…i		*!		*

Table 9.4: Failure to predict ATR dominance

Specifically, we would incorrectly predict that ATR is dominant except in cases in which that would create an ATR mid vowel; in the presence of a mid vowel, RTR will win to avoid such vowels. Clearly, this result is not how Mayak works, nor does it occur in the many other 2IU, 1EO languages considered in existing surveys. However, with these basic constraints and the conceptualization behind them, there is no way to avoid this prediction, because changing the ranking of *eo and IDENT-IO[ATR] would incorrectly predict a tongue root contrast in mid vowels. This is like a reverse derived environment effect, in that derived [e,o] are allowed, while underived ones are banned.[92]

It is worth noting that this prediction of ATR dominance except with mid vowels is in

---

[92] This is the type of pattern that Stratal OT or Lexical Phonology could deal with straightforwardly, by having the phonemic inventory defined at the Stem stratum, where *eo is high ranked, and harmony existing only at the Word stratum, where IDENT-IO[ATR] would outrank *eo. This multi-level approach would mean that the 'input' to harmony at the Word stratum would not contain /e,o/, since they have been eliminated from outputs at the Stem stratum. As I adopt a non-derivational approach here, this type of multi-level analysis is not possible in the present framework. It is also possible to deal with this type of pattern in McCarthy's Comparative Markedness (2003), with 'old markedness' >> faithfulness >> 'new markedness'; this approach was built for cases similar to this one, but again is a different framework from what I adopt here.

fact doubly wrong for the 2IU, 1EO language (Ethiopian) Komo (Otero 2015). In Komo, there are seven contrastive vowels, where the only contrastive ATR vowels are high; however, there are ten surface vowels, with ATR mid and low vowels created as the result of harmony. Specifically, in combinations of ATR high vowels and RTR mid or low vowels, ATR is dominant, and allophonic ATR mid or low vowels are created. For example, /hám-úk/ surfaces as [hə́m-úk] 'yawn-PFV', and similarly /kɛ́ʃ-úk/ as [kéʃ-úk] 'thresh-PFV'. However, in combinations of ATR and RTR vowels where both are high, RTR wins. For example, the same perfective suffix surfaces as [-ʊ́k] in examples like [yìl-ʊ́k] 'see-PFV' and [ɓʊ́g-ʊ́k] 'wait-PFV'. This pattern is precisely the opposite of what we would expect from the discussion above: ATR is dominant exactly when it creates vowels that exist only as the result of harmony, and RTR is dominant otherwise. Note that I do not have an analysis of Komo to propose here; I simply raise it as a particularly compelling example of how dominance and contrast are not related in any straightforward way, as we might predict from simple faithfulness. Table 9.4 also applies to this Komo case, while Table 9.5 shows an incorrect prediction in Komo high vowels.

/ɪ…i/	Harmony	*eo	Ident-IO[ATR]	Ident-IO[RTR]
☹ ɪ…ɪ			*!	
☞ i…i				*

Table 9.5: Failure to predict dominance reversal in Komo

A critical insight for solving the issue in standard 2IU, 1EO languages (not including Komo) comes from previous work such as Archangeli & Pulleyblank (1994), who note that high vowels are better at being ATR, in that ATR high vowels are phonetically and typologically preferable to both RTR high vowels and ATR non-high vowels. This fact is motivated by the articulatory relationship between tongue root advancement and tongue body raising, which makes it easier to produce an ATR high vowel. We might therefore expect that preserving ATR should be more important with high vowels than with non-high vowels. Indeed, while there are a large number of languages in which the only ATR vowels are high and/or the only high vowels are ATR, I am unaware of any languages with a tongue root contrast in which there are no ATR high vowels.

As such, if we are separating Ident-IO[ATR] from Ident-IO[RTR] in order to account for different degrees of faithfulness to ATR versus RTR vowels, we might also distinguish these constraints by height; Ident-IO[ATR]$_{\text{HI}}$ is different from the general Ident-IO[ATR], because high vowels are preferentially ATR, and ATR preferably associates with high vowels. This type of analysis has the effect of allowing RTR high vowels, but not ones derived by harmony; underlying ATR high vowels must remain ATR. Note that something distinct would need to be proposed for Komo. Moreover, note that regardless of the ranking of the high-specific and general ATR faithfulness constraints, the prediction is that only high vowels can behave specially; these constraints are in a stringency relationship.

It is worth noting that these sorts of ‘phonetic grounding’ considerations are typically applied to markedness constraints, rather than faithfulness, as is the case for Archangeli & Pulleyblank’s (1994) Hi/ATR grounded condition. The present idea is a version of suggesting that faithfulness constraints should prefer to preserve the better grounded, or less marked, configuration. However, the idea of specifying faithfulness constraints in this type of way is not novel; faithfulness hierarchies have been proposed and developed in work like de Lacy (2002) and Howe & Pulleyblank (2004), among others. As Howe & Pulleyblank (2004) note, markedness hierarchies can alternatively be re-established as faithfulness hierarchies; as such, the fact that a markedness constraint against ATR high vowels would be at the lowest end of a markedness hierarchy on ATR in vowels of different heights could be translated into it being a high-ranked faithfulness constraint. Howe & Pulleyblank (2004) also discuss preservation of perceptibility, which could be relevant here, if ATR is a more perceptible feature on high vowels. It has been argued that the tongue root contrast is less distinguishable on high than mid vowels, in terms of both acoustics (Starwalt 2008; Becker-Kristal 2010), and perception (Fulop et al. 1998; see also Rose 2018) though it is not clear how that translates into looking into the feature ATR on its own, separately from RTR. Moreover, if we consider articulation rather than perception, then the combination of [+high] and ATR is articulatorily good, and so it would be sensible to preserve it from that perspective.

Given that for other languages, we independently require a markedness constraint against high RTR vowels, this analysis sounds like it could be reformulated as a classic gang effect: mid ATR vowels are banned, except when harmony and the combination of Ident-IO[ATR] and *ɪʊ forces their occurrence in order to prevent changing ATR high vowels into RTR. However, such an approach is in fact impossible. In order to derive the general absence of ATR mid vowels in Mayak, *eo must outweigh Ident-IO[ATR] on its own. Similarly, in order to derive the general presence of RTR high vowels, which are contrastive, Ident-IO[RTR] must outweigh *ɪʊ. Thus, *eo + Ident-IO[RTR] will outweigh Ident-IO[ATR] + *ɪʊ. As can be seen by the violation profiles of the two harmonic candidates given a disharmonic input /ɛ…i/, the result is necessarily a preference for ɛ… ɪ over e…i, exactly the opposite of the attested pattern. We can spell out this argument more explicitly if we give the labels a, b, c, and d to the weights of *eo, Ident-IO[RTR], *ɪʊ, and Ident-IO[ATR] respectively. We know that a>d and b>c, therefore (a+b)>(c+d).

/ɛ…i/	*eo	IDENT-IO[RTR]	*ɪʊ	IDENT-IO[ATR]
☞ ɛ… ɪ			*	*
☹ e…i	*	*		

Table 9.6: Violations of the relevant constraints by harmonic candidates

Instead, then, the facts of many 2IU, 1EO languages, including Mayak – that mid ATR vowels are only permitted when derived by harmony, while RTR high vowels are permitted except if derived by harmony – must be due to a more complex mechanism. Here, I propose a set of constraints in a stringency relationship: IDENT-IO[ATR]$_{\text{HI}}$ and IDENT-IO[ATR], justified by the previously mentioned phonetic and typological facts about the relationship between ATR and height. When the latter outranks or outweighs the former, all vowels will behave the same with respect to ATR harmony[93], but in the reverse situation, high vowels may preserve ATR in situations in which non-high vowels do not. In 2IU, 1EO languages, so long as the relevant harmony constraint outranks *eo, the ranking IDENT-IO[ATR]$_{\text{HI}}$ >> *eo >> IDENT-IO[ATR] will ensure that ATR mid vowels occur as the result of harmony, in order to avoid changing ATR high vowels to RTR, but will not occur elsewhere.

---

[93] However, it is worth noting that if there were some need to change ATR in a combination of high ATR and non-high ATR vowels, then everything else being equal, even this ranking would predict that the non-high ones should be unfaithful. Whether such situations exist and how they behave is unclear. Since the present faithfulness constraints refer only to changing ATR to RTR, not the reverse, harmony is not a context in which this issue would arise.

4.3 Constraints, weights, and scaling

The relevant faithfulness constraints for the Mayak analysis are defined in (11). In order to derive ATR dominance, I separate faithfulness preserving ATR from that preserving RTR; weighting the former higher will ensure that harmony does not change a vowel from ATR to RTR, but instead makes both relevant vowels RTR. Moreover, as discussed above, an additional ATR faithfulness constraint specific to high vowels is needed, in order to capture the generalization that ATR is more crucial to preserve on high than non-high vowels. Note that according to these definitions, an underspecified input vowel would violate none of the faithfulness constraints, regardless of whether it were to surface as ATR or RTR.

(11)

IDENT-IO[ATR]: Assign a violation for every segment that is [ATR] in the input but [RTR] in the output.

IDENT-IO[RTR]: Assign a violation for every segment that is [RTR] in the input but [ATR] in the output.

IDENT-IO$[ATR]_{HI}$: Assign a violation for every segment that is [ATR, +high] in the input but [RTR] in the output.

Given that ATR mid vowels occur only as the result of harmony, we require a markedness constraint against them, which needs to be weighted lower than ATR faithfulness for high vowels, but higher than general ATR faithfulness. The former weighting will ensure that ATR mid vowels do get derived through harmony from a high trigger, in order to preserve ATR on a high vowel, while the latter ranking will ensure that ATR mid vowels do not occur elsewhere.

These weightings will be established through the tableaux in the next subsection. The relevant constraint is defined in (12).

(12) *eo: Assign a violation for every ATR mid vowel in the output.

Of course, constraints forcing harmony are also necessary; I adopt ones that penalize combinations of ATR and RTR, the same as those used in Chapter 8. Since Mayak has harmony to some degree in both directions, the constraints must penalize both ATR followed by RTR and the reverse. The fact that these directions pattern differently in terms of their targets will be derived through the different weights of the two harmony constraints; target scaling factors will be the same in both directions. The harmony constraints are defined in (13); these are adapted from the no-disagreement constraints argued for by Pulleyblank (2004), and developed further in a Harmonic Grammar framework by Chapters 3, 4, and 8 on Hungarian, Votic, and Lango.

(13)
*[RTR]∞[ATR]: Assign a violation for every RTR vowel followed at any distance by an ATR vowel.
*[ATR]∞[RTR]: Assign a violation for every RTR vowel preceded at any distance by an ATR vowel.

Both of these constraints are oriented towards RTR, with ATR as the context; the orientation is necessary for determining how to count the violations of the constraint, which in this case is for every RTR vowel in an ATR context. Moreover, I use these notions to determine the assignment of scaling factors; the context is scaled with the trigger scaling factor, while the focus is scaled with the target scaling factor. Note that this means that, in theory, the actual trigger and target of the harmony that occurs may not be the same as the ‘trigger’ and ‘target’ as defined for the scaling. Indeed, it is the ranking of faithfulness constraints that determines which vowel is actually the trigger, and so ‘trigger’ and ‘target’ for scaling are somewhat misleading terms. For example, an input /ɔ…i/ violates the constraint *[RTR]∞[ATR], and the penalty for the faithful candidate will be assigned the target scaling of [ɔ] and the trigger scaling of [i]. This will be true regardless of how harmony actually operates. For instance, if IDENT-IO[ATR]>>IDENT-IO[RTR], as in Mayak, then this form will harmonize to [o…i], in which case /i/ is indeed the trigger and /ɔ/ is indeed the target. However, this approach does not exclude the possibility that IDENT-IO[RTR]>>IDENT-IO[ATR], in which case the harmonic form [ɔ…ɪ] would be preferred. In this case, /ɔ/ was the actual trigger, and /i/ was the actual target, even though for scaling purposes, the disharmonic candidate would still be scaled with /ɔ/ as the ‘target’ and /i/ as the ‘trigger’. The tableaux in Tables 9.7-9.8 illustrate, using the scaling factors for Mayak. Even though the outcome differs depending on the relative weighting of faithfulness (here represented as ranking for simplicity), the ‘trigger’ and ‘target’ scaling remains the same.

/ɔ…i/	*[RTR]∞[ATR]	IDENT-IO[ATR]	IDENT-IO[RTR]
ɔ… ɪ		-1	
☞ o…i			-1
ɔ…i	-1*4(target ɔ)*4(trigger i)		

Table 9.7: Target ɔ and trigger i with ATR harmony

/ɔ…i/	*[RTR]∞[ATR]	IDENT-IO[RTR]	IDENT-IO[ATR]
☞ɔ… ɪ			-1
o…i		-1	
ɔ…i	-1*4(target ɔ)*4(trigger i)		

Table 9.8: Target ɔ and trigger i with RTR harmony

We know that in Mayak, harmony depends on both trigger and target quality. Only high ATR vowels are triggers of harmony, while low ATR vowels are not. In regressive harmony, high and mid vowels are targets, while low vowels are not; in progressive harmony, only high vowels and one low vowel suffix are targets. Thus, we require the two scaling factors defined in (14) and (15), with trigger scaling modified from Kimper (2011) and both factors adopted previously in Chapters 3, 4, and 8. It is worth noting that it is necessarily the case here that trigger strength always applies to ATR vowels, while target quality always applies to RTR vowels. This is simply due to the formulation of the harmony constraints, specifically the fact that RTR is always the focus, and the focus receives the target scaling by definition. However, as discussed, this does not at all imply that the constraints themselves are forcing ATR dominance; they could just as easily be satisfied through RTR harmony.

(14) Scaling factor: trigger strength

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that x ≥ 1) for each degree i to which α is perceptually impoverished with respect to ±F.

(15) Scaling factor: target quality

For a trigger α, a target β, and a feature F, multiply the penalty earned by a constant x (such that

$x \geq 1$) for each degree i to which β is a quality target (defined by phonetic factors discussed in Section 3) with respect to ±F.

The weights I assume for the constraints are given in (16), while the trigger and target scaling factors, which apply to both of the harmony constraints, are given in (17) and (18) respectively. The specific numbers are not crucial; the numbers simply need to satisfy certain criteria relative to each other. These conditions will be stated and justified in the next subsection, in which I illustrate the analysis with tableaux. As noted above, Mayak makes a distinction between low and non-low vowels in terms of their trigger and target properties; this is the only distinction encoded into these scaling factors. The difference between high and mid vowels in being targets of progressive harmony will be derived using the constraint *eo.

(16) Constraint weights

IDENT-IO[ATR]$_{\text{HI}}$: weight 100

*eo: weight 32

IDENT-IO[ATR]: weight 14

IDENT-IO[RTR]: weight 13

*[RTR]∞[ATR]: weight 3

*[ATR]∞[RTR]: weight 1

(17) Trigger scaling factors for tongue root harmony constraints

4 for non-low vowels

1 for low vowels

(18) Target scaling factors for tongue root harmony constraints

4 for non-low vowels

1 for low vowels

4.4 Regressive harmony

Tableaux illustrating Mayak regressive harmony are given in Tables 9.9 through 9.14. There are several major aspects of the pattern that need to be illustrated: harmony must be triggered by high but not low vowels; harmony must target high and mid but not low vowels; the failure to target the low vowel is not due to an independent ban on ATR low vowels; and ATR mid vowels exist only as the result of harmony from high vowels.

Table 9.9 shows regressive harmony from an ATR high vowel to an RTR mid vowel. Candidate (a), with disharmony, is ruled out by the high scaled weight of the harmony constraint. Given a mid vowel target and a high vowel trigger, the harmony constraint is scaled by 4*4=16; since *[RTR]∞[ATR] has weight 3, the total penalty this candidate incurs is 3*16=48. In contrast, the combined violations of *eo and IDENT-IO[RTR] in candidate (b) incur a total penalty of 32+13=45 for this candidate; it is preferred, since -45>-48. Candidate (c), which changes an ATR high vowel to RTR, is strongly dispreferred by the large weight of IDENT-IO[ATR]$_{HI}$. Thus, ATR mid vowels occur as the result of regressive harmony triggered by high vowels, so long as weight(*eo) + weight(IDENT-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(non-low trigger); this guarantees that candidate (b) will be preferred over candidate (a). Additionally, weight(*eo) + weight(IDENT-IO[RTR]) <

weight(Ident-IO[ATR]$_{HI}$) + weight(Ident-IO[ATR]) is required to ensure that candidate (b) wins over candidate (c).

/ɔ…i/	Ident-IO[ATR]$_{HI}$ 100	*eo 32	Ident-IO[ATR] 14	Ident-IO[RTR] 13	*[RTR]∞[ATR] 3	H
a. ɔ…i					-1*4(ɔ)*4(i)	-48
☞ b. o…i		-1		-1		-45
c. ɔ…ɪ	-1		-1			-114

Table 9.9: ATR harmony among non-low vowels

These conditions will also enforce regressive ATR harmony with a high target (and a high trigger), since the harmonic candidate in such a case violates a subset of the constraints violated by candidate (b) in Table 9.9 (i.e. Ident-IO[RTR] but not *eo). Since weight(*eo) + weight(Ident-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(non-low trigger), it is also the case that weight(Ident-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(non-low trigger), so that candidate (b) is preferred over candidate (a) in the relevant tableau in Table 9.10.

/ɪ…i/	IDENT-IO[ATR]$_{\text{HI}}$ 100	IDENT-IO[ATR] 14	IDENT-IO[RTR] 13	*[RTR]∞[ATR] 3	H
a. ɪ…i				-1*4(ɔ)*4(i)	-48
☞ b. i…i			-1		-13
c. ɪ…ɪ	-1	-1			-114

Table 9.10: ATR harmony among high vowels

While mid vowels do undergo regressive harmony in Mayak, the low vowel [a] does not, even though unlike the mid vowels, it has a contrastive ATR counterpart. As such, it is critical here that regressive harmony that targets [a] is less important than regressive harmony targeting mid vowels. In Table 9.11, candidate (a) violates the harmony constraint *[RTR]∞[ATR], scaled to a low target and high trigger, while candidate (b) violates only IDENT-IO[RTR]. As such, weight(*[RTR]∞[ATR]) * scaling(low target) * scaling(non-low trigger) < weight(IDENT-IO[RTR]) is required to choose the disharmonic candidate (a) over its harmonic competitor candidate (b). Recall from above that weight(*eo) + weight(IDENT-IO[RTR]) < weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(non-low trigger). The only way that both of these requirements can hold is if scaling(low target) < scaling(mid/non-low target), proving the need for the target scaling approach advocated for here. Notably, while I have assumed for illustrative purposes that a constraint *ʌ is so low-weighted in Mayak as to not be relevant to the grammar, this assumption is not necessary. Indeed, in order to derive the contrastive presence of [ʌ] in Mayak, *ʌ would need to be outweighed by IDENT-IO[ATR]; moreover, since [e,o] are not contrastive in Mayak, IDENT-IO[ATR] must be outweighed by *eo. Thus, weight(*ʌ) < weight(*eo), and so even if we included *ʌ, violated here by candidate (b), it would still be necessary to have scaling(low target) < scaling(mid target) to derive both the weighting requirements of this tableau and those of the previous one.

/a…i/	IDENT-IO[ATR]$_{HI}$ 100	IDENT-IO[ATR] 14	IDENT-IO[RTR] 13	*[RTR]∞[ATR] 3	H
☞ a. a…i				-1*1(a)*4(i)	-12
b. ʌ…i			-1		-13
c. a… ɪ	-1	-1			-114

Table 9.11: Neutrality of low vowels

Critically, this analysis allows for the ATR low vowel [ʌ] to surface in the same environment in which [a] fails to undergo harmony, as shown in Table 9.12. Specifically, a combination of [ʌ] with an ATR high vowel does not violate any of the constraints illustrated here, while changing /ʌ…i/ to [a…i], as in candidate (b), violates both IDENT-IO[ATR] and *[RTR]∞[ATR]; it will therefore be ruled out. Candidate (a) would still win even if we included the markedness constraint *ʌ, since as discussed above, it would need to have weight lower than that of IDENT-IO[ATR].

/ʌ …i/	Ident-IO[ATR] 14	*[RTR]∞[ATR] 3	H
☞ a. ʌ…i			0
b. a…i	-1	1*1(a)*4(i)	-26

Table 9.12: No independent ban on the ATR low vowel

The fact that ATR mid vowels will not occur in isolation is derived straightforwardly by the fact that *eo outweighs Ident-IO[ATR]. However, what is more crucial is the fact that ATR mid vowels do not occur even in combination with ATR low vowels; in such cases, there is additional pressure to maintain hypothetical underlying ATR mid vowels from the fact that an RTR mid – ATR low output violates the harmony constraint. To fully satisfy richness of the base, we therefore need to ensure that a hypothetical underlying ATR mid vowel would not surface even before an ATR low vowel. This is shown in Table 9.13. While candidate (a) violates both Ident-IO[ATR] and the harmony constraint, the total penalty it incurs is only -26, since ʌ is a poor trigger of ATR harmony, resulting in a relatively low weighting of the harmony constraint. Thus, this candidate is preferable to (b), which violates *eo, since -26>-32. This result requires weight(*eo) > weight(Ident-IO[ATR]) + weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(low trigger). Changing both relevant vowels to RTR, as in candidate (c), is ruled out by the requirement that weight(Ident-IO[ATR]) > weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(low trigger).

/o…ʌ/	*eo 32	IDENT-IO[ATR] 14	*[RTR]∞[ATR] 3	H
☞ a. ɔ…ʌ		-1	-1*4(ɔ)*1(ʌ)	-26
b. o…ʌ	-1			-32
c. ɔ…a		-2		-28

Table 9.13: No ATR mid vowels even in the context of an ATR low vowel

As a final note for this directionality, regressive harmony from a low potential trigger will not occur even if the potential target is high (and therefore not subject to the constraint *eo). Again, this result is derived from the fact that low vowels are poor triggers. This is shown in Table 9.14, and requires the weights of each of IDENT-IO[ATR] and IDENT-IO[RTR] to be greater than weight(*[RTR]∞[ATR]) * scaling(non-low target) * scaling(low trigger).

/ɪ…ʌ/	IDENT-IO[ATR] 14	IDENT-IO[RTR] 13	*[RTR]∞[ATR] 3	H
☞ a. ɪ…ʌ			-1*4(ɪ)*1(ʌ)	-12
b. i…ʌ		-1		-13
c. ɪ…a	-1			-14

Table 9.14: No ATR harmony from low vowels

4.5 Progressive harmony

Turning now to progressive harmony, the main distinction is that in this case, harmony must target only high vowels, not mid vowels. Since I assume scaling factors must be consistent between directions, this result comes from the lower weight of the progressive harmony constraint compared to the regressive constraint. Indeed, while the regressive constraint scaled for a mid vowel target has sufficient weight to force unfaithfulness plus a marked ATR mid vowel in the output, the progressive one does not. In the present analysis, this result is derived through the constraint *eo, since high and mid vowels have the same target scaling factors. There exists a markedness constraint against ATR mid vowels, but not against ATR high vowels, which is sufficient for creating a weighting schema that will result in harmony occurring with high but not mid targets. Specifically, in ATR high – RTR high combinations, ATR harmony will violate just IDENT-IO[RTR], while in ATR high – RTR mid combinations, harmony also violates *eo. If progressive harmony scaled to a non-low target outweighs IDENT-IO[RTR], but does not outweigh *eo (or the combination of IDENT-IO[RTR] plus *eo), then harmony will occur with a high target, but not with a mid target. As such, scaling high and mid vowels differently as targets is not necessary here; the difference in this case can be attributed to contrast, and there is no reason for learners to do otherwise.

However, it is worth emphasizing that the difference between mid and low vowels still necessitates target scaling, as discussed above. Indeed, *eo needs to outweigh IDENT-IO[ATR], which needs to outweigh *ʌ, since mid ATR vowels do not occur contrastively while low ones do. By transitivity, *eo outweighs *ʌ. As such, without scaling, if regressive harmony is sufficient to force mid ATR vowels to occur as the result of harmony, it must also be sufficient to create low ATR vowels. Thus, despite the fact that there does not need to be a distinction between high and mid vowels for target scaling in Mayak, and the progressive harmony target asymmetries can instead be derived from the constraint against ATR mid vowels, we do still need target scaling in order to derive the difference between mid and low vowels in regressive harmony.

Turning to the tableaux, Table 9.15 shows that progressive harmony occurs between high vowels; this result requires weight(IDENT-IO[RTR]) < weight(*[ATR]∞[RTR]) * scaling(non-low trigger) * scaling(non-low target). The high penalty incurred by disharmony between high vowels rules out the disharmonic candidate (a).

/i…ɪ/	IDENT-IO[RTR] 13	*[ATR]∞[RTR] 1	H
a. i…ɪ		-1*4(i)*4(ɪ)	-16
☞ b. i…i	-1		-13

Table 9.15: Progressive harmony among high vowels

Table 9.16 shows that the same effect does not occur with mid vowels, because *eo is relevant here. While weight(*[ATR]∞[RTR]) * scaling(non-low trigger) * scaling(non-low target) is sufficient to outweigh IDENT-IO[RTR], it is not sufficient to outweigh *eo, which is violated by the harmonic candidate (b). As such, the disharmonic candidate (a) surfaces. The condition weight(*[ATR]∞[RTR]) * scaling(non-low trigger) * scaling(non-low target) < weight(IDENT-IO[RTR]) + weight(*eo) is necessary here.[94]

---

[94] Note that I am using *eo here, instead of treating mid and low vowels in suffixes in the same way. Since I treat the alternating low vowel suffix as exceptional, all non-high suffix vowels behave the same way, but the analysis does not unify them, since it uses *eo to eliminate mid vowel alternations and target scaling to eliminate low vowel ones. This is for two reasons. First, a high weight of the *eo constraint is independently necessary in Mayak to prevent ATR mid vowels from occurring except as the result of regressive harmony; this constraint already does the work of preventing mid vowel suffix alternations, so there is no need to assume different target scaling for mid vowels. This analysis is not possible for low vowels, which are contrastive. Second, I assume that scaling must be the same for both directionalities of a given type of harmony, such that the fact that mid vowels participate in regressive harmony, while low vowels do not, necessitates that they have different scaling for regressive harmony and therefore also for progressive harmony. Notably, since *eo must be weighted higher than *ʌ, as ʌ occurs contrastively, the only way in which mid vowels but not low vowels can participate in regressive harmony is to give mid vowels higher target scaling. The result is that the exclusion of mid and low vowels from progressive harmony is due to different reasons, but those separate reasons are consistent with the independent behaviours of mid versus low vowels elsewhere in the language (inventory and regressive harmony).

/i…ɛ/	*eo	IDENT-IO[RTR]	*[ATR]∞[RTR]	H
	32	13	1	
☞ a. i…ɛ			-1*4(i)*4(ɛ)	-16
b. i…e	-1	-1		-45

Table 9.16: No progressive harmony to mid targets

Since ATR mid vowels occur only as the result of regressive harmony from high vowels, it is also necessary to ensure that a hypothetical underlying ATR mid vowel would not occur even following an ATR high vowel. Table 9.17 illustrates, with candidate (b) ruled out by the high weight of *eo. Here, we require weight(IDENT-IO[ATR]) + weight(*[ATR]∞[RTR]) * scaling(non-low trigger) * scaling(non-low target) < weight(*eo).

/i…e/	*eo	IDENT-IO[ATR]	*[ATR]∞[RTR]	H
	32	14	1	
☞ a. i…ɛ		-1	-1*4(i)*4(ɛ)	-30
b. i…e	-1			-32

Table 9.17: No ATR mid vowels after high vowels, even if present in input

Progressive harmony also does not occur with a low target; this is already apparent because regressive harmony does not target low vowels and the progressive harmony constraint is weighted lower than the regressive one, but is illustrated in Table 9.18. For identical reasons, the ATR low vowel cannot act as a trigger of progressive harmony, as shown in Table 9.19.

/i…a/	IDENT-IO[RTR] 13	*[ATR]∞[RTR] 1	H
☞ a. i…a		-1*4(i)*1(a)	-4
b. i…ʌ	-1		-13

Table 9.18: No progressive harmony to low targets (except for the one alternating suffix, to be discussed in the next subsection)

/ʌ…ɪ/	IDENT-IO[RTR] 13	*[ATR]∞[RTR] 1	H
☞ a. ʌ…ɪ		-1*1(ʌ)*4(ɪ)	-4
b. ʌ…i	-1		-13

Table 9.19: No progressive harmony to high front vowels from low triggers[95]

A final case worth noting is that an ATR mid vowel will not occur even if it could act as a trigger of regressive harmony onto an RTR high vowel. This result essentially occurs because the constraint IDENT-IO[ATR]$_{\text{HI}}$ is not relevant in this situation, and thus the best option is to make the mid vowel RTR, in order to satisfy *eo. This result could be considered an instance of RTR harmony, although we have no evidence that there are any underlying situations like this in the language. Table 9.20 illustrates.

---

[95] The harmony from [ʌ] to the suffix /-ʊk/ needs to be considered in further research; it is possibly due to harmony being parasitic on agreement in backness, if low vowels are [+back] in Mayak. See Section 4.6 for more discussion on this suffix.

/ɪ…e/	*eo 32	IDENT-IO[ATR] 14	IDENT-IO[RTR] 13	*[RTR]∞[ATR] 3	*[ATR]∞[RTR] 1	H
☞ a. ɪ…ɛ		-1				-14
b. i…e	-1		-1			-45
c. i…ɛ		-1	-1		-1*4(i)*4(ɛ)	-43
d. ɪ…e	-1			-1*4(ɪ)*4(ɛ)		-80

Table 9.20: No ATR mid vowels, even if they could act as a trigger

Sections 4.3 and 4.4 have shown that all of the (non-exceptional) Mayak data can be derived using these constraints, with the conditions on weighting as discussed above. Worth noting is that the distinction between low and non-low triggers and targets is critical here. With triggers, a mechanism is required to ensure that ATR high vowels act as triggers while ATR low vowels do not. With targets, the situation is even more interesting: the low RTR vowel is not a target, despite having a contrastive ATR counterpart, while the mid RTR vowels do act as targets for regressive harmony, even though ATR mid vowels are not permitted in any other context. This discrepancy, in which ATR mid vowels are permitted only as a result of harmony, while ATR low vowels are permitted in general but never result from harmony, necessitates a distinction between low and mid (or non-low) targets to be built into the harmony-driving constraint.

4.6 Exceptional suffix(es)

Before concluding the analysis, I briefly discuss some options for the exceptional suffix that alternates between [a] and [ʌ], as well as the suffix that alternates between [ʊ] and [u], which Andersen (1999) also considers exceptional. Andersen (1999) assumes that these suffixes are underspecified underlyingly; this assumption would be consistent with the present approach. Indeed, progressive harmony will target low vowels if they are underlyingly unspecified for [ATR]; the candidates will tie on the faithfulness constraints, which means that the decision falls to the harmony constraints. It is therefore preferable to harmonize, so that an unspecified suffix would surface as [ʌ] following ATR roots and [a] following RTR roots. Table 9.21 illustrates; compare Table 9.18, with a suffix /a/.

/i…A/	IDENT-IO[ATR]$_{\text{HI}}$ 100	*eo 32	IDENT-IO[ATR] 14	IDENT-IO[RTR] 13	*[RTR]∞[ATR] 3	*[ATR]∞[RTR] 1	H
☞ a. i…ʌ							0
b. i…a						-1*4(i)*1(a)	-4

Table 9.21: Progressive harmony to underspecified vowel

Note that whether feature-filling of the sort in these candidates violates IDENT-IO[ATR] and IDENT-IO[RTR] is a matter of debate; the definitions in (11) mean that it would not be, and the same assumption is made by Inkelas (1995) for Turkish voicing. This is irrelevant here, as either way, the candidates tie on faithfulness, and the decision is left to the harmony constraint.

Nonetheless, this pattern could instead be derived as morpheme specificity rather than phonological representation, through an analysis of exceptionality that adopts indexation. There are two possible ways of implementing this: one follows Chapter 3 on Hungarian, in which the vowel is exceptionally indexed to behave like a target from another category; the other follows more standard exceptionality accounts, as discussed for instance by Jurgec & Bjorkman (2018), and indexes the harmonizing low vowel suffix to a higher-weighted version of the progressive harmony constraint. The former would result in /a/ behaving like non-low vowels, and harmony would occur exactly as in Table 9.15. For the latter approach, we would requiring marking both the suffix /a/ and a higher-weighted version of the constraint *[ATR]∞<u>[RTR]</u> as exceptional, notated here as a subscript 'e'. As long as the exceptional harmony constraint is weighted highly enough, doing so correctly derives harmony of the exceptional suffix after ATR vowels, including after root [ʌ], as illustrated in Table 9.22.

/ʌ…a$_e$/	*[ATR]∞<u>[RTR]</u>$_e$ 15	IDENT-IO[RTR] 13	*[ATR]∞<u>[RTR]</u> 1	H
☞ a. ʌ…ʌ		-1		-13
b. ʌ…a	-1*1(ʌ)*1(a)		-1*1(ʌ)*1(a)	-16

Table 9.22: Progressive harmony to an exceptional low vowel

The same indexation approach would obviously work for the suffix /-ʊk/, to ensure that it harmonizes after root [ʌ]; an underspecification account would also work in this case, since the only relevant constraint would be the harmony constraint. Either of these analyses would amount to considering this suffix as exceptional. However, there remains the possibility that the harmony of this suffix after [ʌ] is a genuinely phonological fact, if there is a (higher-weighted) drive for ATR-RTR harmony when the trigger and target agree in backness. This possibility remains for future research.

5 Discussion and conclusion

5.1 Predictions

This way of analyzing the pattern makes an important prediction regarding the behaviour of low vowels. If the only low vowel in a 2IU, 1EO language is the RTR [a], then it can behave like the mid vowels or neutrally, as expected. Critically, however, are cases in which there is a tongue root contrast in low but not mid vowels, like in Mayak. In such cases, the present analysis predicts one of two options, neither of which is ATR dominance. The low ATR vowel could be a non-trigger of harmony, like in Mayak; this option is accomplished by scaling the harmony constraint such that with a low trigger, the result is insufficient to force unfaithfulness. Alternatively, combinations of ATR low vowels with RTR mid vowels could harmonize to RTR; the low vowel is not subject to the high-specific ATR faithfulness, and so we would obtain the result discussed in relation to Table 9.4, where RTR wins to avoid ATR mid vowels. I am not aware of any such cases, but also not aware of any languages in which low but not mid vowels contrast that do not behave like Mayak. A language in which ATR low vowels and RTR mid vowels harmonize to ATR, even though ATR mid vowels exist only as a result of harmony in the language, would force a re-analysis, since it is incompatible with the present view.

Relatedly, an odd fact about this analysis is that it in some sense puts the fact that high vowels are better triggers into two places: in the faithfulness constraint, where ATR gets preserved on high vowels to the exclusion of other vowels, and in the scaling of the harmony constraint, where disharmony is penalized more severely for high triggers than non-high ones. Whether or not this fact constitutes a problem or loss of generalization depends in part on our view of why high vowels are good triggers. In a footnote, Potts et al. (2010), referencing personal communication with McCarthy, note that ATR high vowels might be good triggers of harmony because the tongue root is particularly advanced with high vowels. This reasoning sounds related to why ATR is good on high vowels: tongue root advancement and tongue body raising are correlated. In that case, we might want there to be a tighter connection between wanting to preferentially preserve ATR on high vowels and having high vowels be the best triggers. However, this explanation is not the only one proposed for why high vowels trigger harmony: it could be due to a weak perceptual tongue root contrast in high vowels, as argued for instance by Rose (2018). If that is the case, the fact that high vowels are good triggers of ATR harmony is essentially independent of the fact that ATR is good on high vowels; the former is about the properties of the contrast, while the latter is about the properties of a specific vowel.

In either case, the present analysis requires both trigger scaling and preferential preservation of ATR in high vowels; without the former, harmony would be expected with the low ATR vowel as a trigger, and without the latter, we encounter a paradox regarding the behaviour of mid vowels. Whether these components of the analysis could be brought together remains to be explored in future work.

5.2 Conclusions

Mayak is a particularly interesting case of ATR harmony because the low vowel [a] is paired for ATR, but generally does not alternate with its ATR counterpart. The fact that [a] and [ʌ] alternate harmonically in one affix is sufficient to show that they are indeed paired for ATR, yet the fact that [a] does not otherwise undergo harmony in Mayak is mysterious under the often-held perspective that unpairedness is the reason that [a] is often neutral in tongue root harmony systems. Moreover, it is particularly strange in Mayak that the mid vowels, which do not have contrastive ATR counterparts, undergo harmony allophonically, in the same contexts in which the contrastively paired [a] fails to harmonize.

Mayak illustrates that the cross-linguistic dispreference for low vowels to harmonize in tongue root harmony is unrelated to whether they have harmonic counterparts. Instead, some independent factor about low vowels makes them poor triggers and targets of ATR harmony, resulting in a pattern in which [ʌ] fails to trigger harmony in contexts in which non-low vowels do trigger, and in which [a] is neutral despite having a counterpart that is perfectly permitted to occur in the contexts in which it fails to harmonize. It is not sufficient to claim that this pattern results from a dispreference for [ʌ]; that would fail to explain why [ʌ] does not act as a trigger. Moreover, in general, Mayak has a dispreference for ATR mid vowels, given that [e] and [o] do not occur except as a result of harmony; however, that dispreference is overridden by harmony. It would therefore be unusual to claim that the failure of [a] to harmonize is due to a dispreference for [ʌ], given that [ʌ] occurs contrastively in the language.

In the present analysis, these facts are accounted for by scaling the harmony constraint by trigger strength and target quality; low vowels are worse triggers and targets than non-low vowels, meaning the pressure to harmonize is less, and in particular is insufficient to force unfaithfulness. The fact that mid vowels do participate in some contexts is due to the greater pressure for them to be targets, which is sufficient to force not only unfaithfulness but also the presence of vowels that are otherwise absent from the language. Treating mid and low vowels as having distinct pressures to harmonize is consistent with their cross-linguistic patterning; mid vowels almost always participate in ATR dominant harmony, even if only allophonically (due to the absence of contrastive ATR mid vowels), while low vowels are frequently neutral, even with contrastive counterparts like in Mayak.

The Mayak pattern also shows a distinction between high and mid vowels in targets, in that high but not mid vowels are subject to progressive harmony. It would be possible to account for this fact using an additional differentiation between the quality of high versus mid vowels as targets, but doing so is not necessary. Instead, the pattern can be derived through stronger regressive than progressive harmony, combined with the fact that mid vowels do not contrast for ATR, while high vowels do. Regressive harmony with a mid target is sufficient to outweigh the pressure against ATR mid vowels, while progressive harmony is not. With high targets, there is no such additional relevant markedness constraint in conflict with harmony, and so the fact that progressive harmony with a high target outweighs faithfulness is sufficient to ensure harmony will occur. Thus, the neutrality of mid vowels as targets to progressive harmony in this analysis is due to the same constraints that causes the absence of their contrastive counterparts in the inventory; Mayak makes a distinction in targets between low and non-low, but does not need any further distinctions.

It is worth noting that, cross-linguistically, there is in fact a three-way target distinction in ATR dominant harmony: high versus mid versus low. As noted, Mayak requires a low versus non-low distinction, but other languages like Lango require high versus non-high (see Chapter 8). In Lango, like in Mayak, there are contexts of harmony where the only targets are high. However, unlike Mayak, Lango has contrastive ATR mid vowels, meaning that it is not possible to employ the solution proposed here for Mayak, of using general markedness against ATR mid vowels to explain harmony targeting only high vowels. Instead, Lango does appear to require a high versus non-high target distinction, while Mayak requires separating low versus non-low targets. Future research should examine whether any languages critically require all three levels of target quality inherent in these two patterns, or whether languages consistently choose to make a binary distinction, with either the feature [+high] or [-low] determining participation, but not both.

A crucial point raised here is that there are multiple factors that can contribute to neutrality. Despite the existence of such simplifying claims in the literature (e.g. van der Hulst

2016), neutrality is not equivalent to the lack of a harmonic counterpart. Mayak is an excellent illustration of the many facets in the relationship (or lack thereof) between neutrality and contrast. The low vowels contrast for ATR, yet are neutral, both in terms of the ATR low vowel not triggering harmony and the RTR low vowel not acting as a target. This fact is the most problematic for the view in which neutrality corresponds to lack of contrast; even in a view with a more nuanced view of what contrast means, like the Contrastive Hierarchy approach (Dresher 2009), this type of pattern with neutrality of contrastive vowels should not be possible. Indeed, the Contrastive Hierarchy assigns features to the segment inventory by successive divisions with features, and so the ordering of the features determines in part the specifications that a segment has. In such an approach, a segment X that is not contrastive for a given feature [F] may nonetheless be specified for [F], if [F] is ordered before X has been fully differentiated from the rest of the inventory. However, there is no way in this approach in which a segment that has a counterpart for [F] can be unspecified for it, since X will not be differentiated from its counterpart until the feature [F] has been specified. Thus, even in this more nuanced view of contrast, Mayak low vowels must be specified for ATR, and so if neutrality is just about contrastive specifications for ATR, they should not be neutral.

Nonetheless, the Mayak pattern is even more complicated in the types of relationships it illustrates. Of course, high vowels contrast for ATR and act as both triggers and targets, as expected. With mid vowels, the situation is more interesting, because here is the one height where there is a lack of harmonic counterpart in the underlying inventory, yet mid vowels participate in one direction and not the other. This result, like that of low vowels, is unexpected by any view of contrast in which contrast determines participation. Even in a view like the Contrastive Hierarchy, in which segments without harmonic counterparts can be contrastive for the harmonic feature depending on the ordering of features, we would expect mid vowels to either consistently participate in harmony, or else consistently be neutral. Instead, in Mayak, their behaviour varies depending on directionality. We therefore cannot say that contrast has a critical role in determining which vowels participate in Mayak harmony, for low vowels or for mid vowels. However, as noted, it was the lack of contrast in mid vowels that allowed for an account of Mayak progressive harmony without three levels of target quality scaling; the existence of the *eo constraint meant that there was no need to distinguish mid from high vowels within the harmony constraint itself.

An additional consideration worth exploring is whether the Mayak pattern is really about targets, or whether it is preferable to consider it a case of parasitic harmony. It is unclear whether there is any clear conceptual or analytical advantage to either of these options. Ultimately, harmony in Mayak is essentially parasitic on [-low], with [+high] also potentially relevant for progressive harmony depending on the analysis. However, if we consider this a basic case of parasitic harmony, without further considering the nature of the targets, we would want to know why harmony would be parasitic on [-low]. The answer to that seems to be that low vowels are simply not good targets for ATR harmony, in that harmonizing [a] is bad in Mayak even in contexts in which its ATR counterpart can occur, and even in contexts in which harmonizing RTR mid vowels to their non-contrastive ATR counterparts is good. We simply do not know enough about the articulation and perception of tongue root contrasts among vowels of different heights to know whether there is any additional motivation for harmony among non-low but not low vowels. Specifically, we do not know enough to determine whether there is more motivation for this being a case of trigger-target agreement or of independent trigger and target properties that happen to result in a pattern that looks like trigger and target need to agree in [-low].

Understanding the broader typology of the behaviour of vowels of different heights in ATR harmony systems could help to answer these questions. Typologically, there does seem to be a scale of participation, in which high vowels participate more than mid vowels, which participate more than low vowels, at least for ATR dominant harmony. This same scale appears to be true for triggers. For both cases, it is not entirely clear why, and therefore it is difficult to speculate whether a pattern like Mayak is due to agreement or due to independent properties of triggers and targets that happen to overlap in this instance. It is worth noting that, unlike in some other types of parasitic harmony (e.g. cases of rounding harmony; see Kaun 1995), it is not the case in Mayak that there is harmony as long as the vowels agree in [low]; there is simply no harmony if either of the relevant vowels is [+low].

Overall, the Mayak pattern is particularly intriguing, given the lack of relationship between contrast and neutrality. It can be analyzed fairly straightforwardly, as long as we assume that the harmony constraint is sensitive to the height of the interacting vowels, with a greater pressure to harmonize among non-low vowels than when one (or both) of the interacting vowels is low.

Chapter 10: Conclusion

1 Summary

This dissertation has dealt with the complexities in which vowels participate in a number of front/back and ATR harmony systems. Evidently, participation versus neutrality in vowel harmony is more complicated than simply looking at which segments have harmonic counterparts in the phonemic inventory of the language. Indeed, there are differences both within and across languages in whether unpaired vowels are neutral and whether paired vowels participate. Despite this lack of direct correlation, there are contrast-independent patterns in which vowels participate in harmony, but these appear to be based on vowel quality rather than contrast.

It is worth noting that, in a large number of languages, neutral vowels are the ones that are unpaired, and unpaired vowels are neutral; the generalization made in previous literature does come from an actual tendency, just one that does not hold universally. One question raised is where that cross-linguistic tendency comes from. In the analysis presented through this dissertation, one answer is that the mechanisms behind lack of contrast are in fact still one of the ways in which neutrality to harmony can arise within the phonological grammar; this was the analysis I adopted for Alur. Beyond that, though, it is unclear whether or how the connection between vowels that tend to be unpaired and vowels that tend to be neutral is encoded in the grammar, other than that the same factors happen to contribute to both properties.

Both typologically and within individual languages, there are height effects for participation in both front/back and ATR harmony. In front/back harmony, lower vowels are better targets, while in ATR-dominant harmony, higher vowels are the better targets. These generalizations are independent of the harmonic pairing of the vowels, and as such, cannot be incorporated into the theory without a mechanism for determining participation that is independent from the structure of the inventory.

I have proposed a formal analysis in which the penalty for disharmony depends on which segments are interacting, as well as how far apart they are. Building on previous work that incorporates scaling of harmony based on trigger and distance (Kimper 2011), I have argued for the necessity of scaling for target and added this mechanism to the account. I implemented this approach for two cases of front/back harmony, in Hungarian and Votic, and for two of ATR harmony, in Lango and Mayak. The same principles and formal mechanisms could in theory be used to account for Alur, but are not in fact necessary. Critically, the point here is not that differentiation of specific vowels as harmony targets is the only mechanism by which vowels may be neutral (i.e. it can also be due to the same mechanisms that cause lack of contrast), but that it must be a possibility afforded by the phonological grammar.

Additionally, I examined whether these types of cross-linguistic generalizations about segment participation also hold gradiently in lexical corpora of individual languages with front/back harmony. While this aspect of the dissertation was only a proof-of-concept, it is important because it asks this question about how participation is reflected in gradient co-occurrence relationships, and proposes a quantitative metric that can be used to examine this question in future research.

2 Future directions

There are many opportunities for further work based on this dissertation. Evidently, the number of languages analyzed in detail here is small, and there are a large number of additional cases of harmony that could be explored using the theoretical principles introduced here. There are many languages that have been described as having complicated systems of front/back or ATR-dominant harmony that could be studied using this analysis. This includes looking into not only categorical patterns of participation versus neutrality, but also which vowels tend to be exceptional. For example, low vowels seem to often be exceptionally disharmonic in ATR-dominant harmony systems, even in cases where they generally participate (e.g. Baka, Parker 1985: [a] is paired and usually harmonizes, but is neutral in some stems; Lulubo, Andersen 1987: ATR-[a] sequences are forbidden, while [a]-ATR is permitted, even though generally RTR-ATR is not; Bari, Steinberger & Vago 1987: non-target exceptionality involves the contrastively paired vowel [a]; Kalenjin, Lodge 1995: three of four opaque non-harmonizing morphemes involve [a], which is contrastively paired and usually harmonizes).

Additionally, future research needs to look at other types of tongue root harmony, beyond ATR-dominant systems. As noted in Chapter 5, trigger and target asymmetries appear to pattern differently between these harmony types, and there is a clear correlation between inventory contrasts in high vowels and the type of tongue root harmony a language exhibits. At present, this connection is incorporated into multiple places in the analysis: contrasts in high vowels are encoded by the ranking of *[RTR,+hi] with respect to IDENT-IO[RTR], ATR dominance is encoded by the ranking of ATR faithfulness above RTR faithfulness, and trigger and target asymmetries are encoded through the scaling factors. Finding a way to express these correlations formally is critical in further work. In particular, if the same tongue root feature is involved in all types of tongue root harmony, then presumably any articulatory and perceptual factors influencing triggers and targets are the same; understanding why triggers and targets nonetheless pattern differently in these systems is therefore important. Given the many different phonetic factors involved in the realization of tongue root contrasts (e.g. tongue root movement, phonation type, etc.), as discussed in Chapter 5, it is not entirely clear that the same tongue root feature is involved in all tongue root harmony systems; such potential phonetic differences should be further explored as a way to understand the trigger/target differences. In terms of front/back harmony, it may also be worth examining umlaut patterns (e.g. Walker 2011; van der Hulst 2018), which often involve front/back, to determine whether they behave the same way as front/back harmony and whether similar trigger and target asymmetries exist there.

Further, corpora of more languages should be analyzed using the segment-specific participation metric developed in Chapter 4. In particular, it should be tested on languages with other types of harmony. For example, it could be tested on languages with tongue root harmony, to see whether the height-related asymmetries in vowel behaviour seen there are also reflected in the corpus statistics. Moreover, looking at more refined corpora, in which loans, derived forms, and so on are removed, would also be helpful, in order to determine the extent to which the generalizations are still apparent in native roots. It could also be worth incorporating the corpus generalizations into the formal analysis, by extending the theoretical model to predict fine-grained segment-specific participation properties. Whether this is necessary depends on one's view of what the phonological grammar can see, in other words, on whether these gradient patterns are part of the phonological grammar or simply provide the basis for patterns that can get phonologized by subsequent generations of learners. This question requires more research to answer; it would be useful to test whether speakers' phonological representations are sensitive to these gradient differences in segment-specific participation. If speakers do not have any sort of knowledge of these patterns, then it would be reasonable to claim that they are not part of the phonology and do not require formal analysis as an aspect of the synchronic grammar. In contrast, if speakers are sensitive to the distinctions, then we could incorporate them, for instance, into a type of phonological analysis that models specific likelihoods of particular forms (e.g. MaxEnt; Hayes & Wilson 2008).

A large amount of experimental work could also be done to examine the target asymmetries discussed throughout this dissertation. Articulatory and acoustic work could help confirm the reasons behind the types of asymmetries we see in the typology, such as whether contrasts are produced differently at different heights. For example, if it is more difficult to produce front/back disharmonic forms involving lower vowels, that could provide an articulatory reason for front/back harmony being preferred among lower vowels. In terms of perception, we might want to look at whether speakers are better at hearing contrasts in certain vowels in a way that reflects their likelihood of participating in harmony. This could be extended to look at whether height effects are apparent in the perception of disharmony. For example, given that cross-linguistically ATR disharmony involving low vowels appears to be less problematic than disharmony involving high vowels, we could ask whether speakers of a language that harmonizes across-the-board perceive these types of disharmony differently.

References

Abondolo, Daniel. 1998. Introduction. *The Uralic languages*, ed. Daniel Abondolo. New York: Routledge.

Agresti, Alan. 2013. *Categorical data analysis*. John Wiley & Sons.

Ahlqvist, August. 1856. *Wotisk grammatik: Jemte språkprof och ordförteckning*. Helsinki: Friis.

Allen, Blake, Pulleyblank, Douglas & Ajíbóyè, Ọládiípọ̀. (2013). Articulatory mapping of Yoruba vowels: An ultrasound study. *Phonology* 30, 183-210.

Andersen, Torben. 1986. Tone splitting and vowel quality: Evidence from Lugbara. *Studies in African Linguistics*, *17*(1), 55-68.

Andersen, Torben. 1987. An outline of Lulubo phonology. *Studies in African Linguistics*, *18*(1), 39-66.

Andersen, Torben. 1989. The Päri vowel system with an internal reconstruction of its historical development. J*ournal of African Languages and Linguistics* 11: 1-20.

Andersen, Torben. 1990. Vowel length in Western Nilotic languages. *Acta Linguistica Hafniensia* 22: 5-26.

Andersen, Torben. 1999. Vowel harmony and vowel alternation in Mayak (Western Nilotic). *Studies in African Linguistics* 28(1). 1–30. https://journals.linguisticsociety.org/elanguage/sal/article/download/1337/1337-2131-1-PB.pdf

Andersen, Torben. 2004. Jumjum phonology. *Studies in African linguistics*, 33(2), 133-162.

Andersen, Torben. 2007. Kurmuk phonology. *Studies in African Linguistics*, 36(1), 30-90.

Anderson, Lloyd B. 1980. Using asymmetrical and gradient data in the study of vowel harmony. In Robert Vago (ed.), *Issues in vowel harmony*, 271–340. Amsterdam: John Benjamins.

Anttila, Arto & Paul Boersma. 2001. Variation in Finnish phonology and morphology. *Glot*

International, *5*(1), 31-40.

Archangeli, Diana & Douglas Pulleyblank. 1994. *Grounded phonology*. Cambridge, MA: MIT Press.

Archangeli, Diana & Douglas Pulleyblank. 2002. Kinande vowel harmony: domains, grounded conditions and one-sided alignment. *Phonology* 19(2). 139–188. DOI: 10.1017/S095267570200430

Archangeli, Diana & Douglas Pulleyblank. 2007. Vowel harmony. In Paul de Lacy (ed.), *The Cambridge handbook of phonology*, 353–378. Cambridge: Cambridge University Press.

Ariste, Paul. 1968. *A grammar of the Votic language*. Bloomington: Indiana University.

Arsenault, Paul. 2012. Retroflex Consonant Harmony in South Asia. PhD dissertation, University of Toronto.

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

Baković, Eric & Colin Wilson 2000. Transparency, strict locality, and targeted constraints. In Roger Billerey & Brook Danielle Lillehaugen (eds.), *Proceedings of the 19th West Coast Conference on Formal Linguistics*, 43–56. Somerville: Cascadilla Press.

Becker-Kristal, Roy. 2010. *Acoustic Typology of Vowel Inventories and Dispersion Theory: Insights from a Large Cross-Linguistic Corpus.* PhD Dissertation, UCLA.

Beckman, Jill 1998. *Positional faithfulness*. Amherst: University of Massachusetts, Amherst dissertation.

Beddor, Patrice, Rena A. Krakow & Stephanie Lindemann. 2001. Patterns of perceptual compensation and their phonological consequences. In Elizabeth Hume & Keith Johnson (eds.), *The role of perceptual phenomena in phonology*, 55–78. San Diego: Academic

Press.

Beddor, Patrice & Handan Yavuz. 1995. The relationship between vowel-to-vowel coarticulation and vowel harmony in Turkish. In Kjell Elenius & Peter Branderud (eds.), *Proceedings of the 13th International Congress of Phonetic Sciences*, vol. 2, 44–51. Stockholm, Sweden.

Benkő, Ágnes, Ágnes Lukács, Katalin Mády, Fanni Patay, Péter Rebrus & Miklós Törkenczy. 2017. Testing variability effects in Hungarian vowel harmony. *The 13th International Conference on the Structure of Hungarian* (ICSH13), MTA Nyelvtudományi Intézet, Budapest. http://budling.nytud.hu/~tork/owncikk/ICSH13.pdf.

Benus, Stefan. 2005. *Dynamics and transparency in vowel harmony.* New York: New York University dissertation.

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

Berta, Árpád. 2015. Tatar and Bashkir. *The Turkic Languages,* ed. Lars Johanson & Éva Á. Csató. New York : Routledge.

Blaho, Sylvia & Dániel Szeredi. 2013. Hungarian neutral vowels: A microcomparison. *Nordlyd*, *40*(1), 20-40.

Blumenfeld, Lev & Ida Toivonen. 2016. A featural paradox in Votic harmony. *Natural Language & Linguistic Theory* 34:1167–1180.

Boyd, Virginia L. 2015. *The phonological systems of the Mbam languages of Cameroon with a focus on vowels and vowel harmony*. LOT, Utrecht.

Brown, Jason Camy. 2008. *Theoretical aspects of Gitksan phonology*. PhD dissertation, University of British Columbia.

Bowman, Samuel R. 2013. Two arguments for vowel harmony by trigger competition. *Proceedings of the 49th Annual Meeting of the Chicago Linguistic Society*. http://www.nyu.edu/projects/bowman/CLS_Seto.pdf.

Calabrese, Andrea, ed. 2005. *Markedness and economy in a derivational model of phonology*. Berlin: Mouton de Gruyter.

Casali, Roderic 2003. [ATR] value asymmetries and underlying vowel inventory structure in Niger-Congo and Nilo- Saharan. *Linguistic Typology* 7.307–82.

Casali, Roderic. 2008. ATR harmony in African languages. *Language and Linguistics Compass* 2/3: 496–549.

Casali, Roderic. 2016. Some inventory-related asymmetries in the patterning of tongue root harmony systems. *Studies in African Linguistics* 45:95-140.

Chumbow, Beban Sammy. 1982. Ogori vowel harmony: An autosegmental perspective. *Linguistic Analysis.* 10(1): 61-93.

Clark, Larry. 2015. Chuvash. *The Turkic Languages,* ed. Lars Johanson & Éva Á. Csató. New York : Routledge.

Clements, George N. & Engin Sezer. 1982. Vowel and consonant disharmony in Turkish. *The structure of phonological representations*, *2*, 213-255.

Coates, Heather. 1985. Otuho phonology and orthography. *Occasional papers in the study of Sudanese languages*, *4*, 86-118.

Coetzee, Andries W. & Shigeto Kawahara. 2013. Frequency biases in phonological variation. *Natural Language & Linguistic Theory* 31(1). 47–89. DOI: 10.1007/s11049-012-9179-z

Coetzee, Andries W. & Joe Pater. 2008. Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. *Natural Language & Linguistic Theory*, 26(2), 289-337.

Cole, Jennifer & Charles Kisseberth. 1994. An optimal domains theory of harmony. *Studies in the Linguistic Sciences*, *24*(2), 1-13.

De Lacy, Paul. 2004. Markedness conflation in optimality theory. *Phonology* 21(2). 145–199. DOI: 10.1017/S095267570400019

Dimmendaal, Gerrit J. 2002. Constraining disharmony in Nilotic: What does an optimal system look like?. *Journal of African Languages and Linguistics*, *23*(2), 153-182.

Dimmendaal, Gerrit J. 2012. Metrical structures: A neglected property of Nilotic (and other African language families). *Studies in Nilotic Linguistics*, *5*, 1-26.

Dresher, B. Elan. 2009. *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.

Edmondson, Jerold A. & John H. Esling. 2006. The valves of the throat and their functioning in tone, vocal register and stress: laryngoscopic case studies. *Phonology*, *23*(2), 157-191.

Farkas, Donka & Patrice Beddor. 1987. Privative and equipollent backness in Hungarian. In *23rd annual regional meeting of the Chicago Linguistics Society*, vol. 2, 90–105.

Finley, Sara. 2015. Learning nonadjacent dependencies in phonology: Transparent vowels in vowel harmony. *Language* 91(1). 48–72. DOI:10.1353/lan.2015.0010

Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe. 2004. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22: 179–228.

Fulop, Sean, Ethelbert Kari & Peter Ladefoged. 1998. An acoustic study of tongue root contrast in Degema vowels. *Phonetica* 55:80-98.1988.

Gafos, Adamantios & Amanda Dye. 2011. Vowel harmony: transparent and opaque vowels. In Marc van Oostendorp (ed.), *The Blackwell companion to phonology*, vol. 4, 2164–2189. Hoboken, NJ: John Wiley & Sons.

Gick, Byran, Douglas Pulleyblank, Fiona Campbell & Ngessimo Mutaka. 2006. Low vowels and transparency in Kinande vowel harmony. *Phonology*, *23*(1), 1-20.

Goldsmith, John. 1985. Vowel harmony in Khalkha Mongolian, Yaka, Finnish and Hungarian. *Phonological Yearbook* 2. 253–275. http://www.jstor.org/stable/4419959

Gordon, Matthew. 1999. The "neutral" vowels of Finnish: How neutral are they. *Linguistica Uralica*, *35*(1), 17-21.

Gósy, Mária & Péter Siptár 2015. Abstractness or complexity? The case of Hungarian /a:/. In Katalin É. Kiss, Balázs Surányi & Éva Dékány (eds.), *Approaches to Hungarian 14: Papers from the 2013 Piliscsaba Conference*, 147–165. Amsterdam: John Benjamins,

Graff, Peter & T. Florian Jaeger. 2009. Locality and feature specificity in OCP effects: Evidence from Aymara, Dutch, and Javanese. In *Proceedings from the annual meeting of the Chicago linguistic society* (Vol. 45, No. 1, pp. 127-141). Chicago Linguistic Society.

Hahn, Reinhard F. 1991. *Spoken Uyghur*. Seattle: University of Washington Press.

Hall, Beatrice L. & R. M. R. Hall. 1980. Nez Perce vowel harmony: an Africanist explanation and some theoretical questions. In Robert Vago (ed.), *Issues in Vowel Harmony.* Amsterdam: John Benjamins, 201-36.

Hall, Beatrice L. & Eluzai M. Yokwe. 1981. Bari Vowel Harmony: The evolution of a cross-height vowel harmony system. *Occasional Papers in the Study of Sudanese Languages*, *1*, 55-63.

Hall, Daniel C. 2018. Locality, neutrality, and contrast: A new resolution to the Votic paradox. In *Proceedings of NELS 48*, University of Iceland.

Hall, Kathleen C., Blake Allen, Michael Fry, Scott Mackie & Michael McAuliffe. 2015. Phonological CorpusTools, Version 1.1.0. [Computer program]. Available from https://sourceforge.net/projects/phonologicalcorpustools/.

Halle, Morris. 1964. On the bases of phonology. In *The structure of language*, eds. Jerry A. Fodor and Jerrold J. Katz, 324–333. New York: Prentice Hall. [Reprinted in Makkai, Valerie B. (ed.) 1972. Phonological theory: evolution and current practice, pp. 393–400. New York: Holt, Rinehart and Winston].

Halle, Morris, Bert Vaux & Andrew Wolfe. 2000. On feature spreading and the representation of place of articulation. *Linguistic Inquiry* 31: 387-444.

Hansson, Gunnar Ó. 2001. *Theoretical and typological issues in consonant harmony*. Berkeley: University of California, Berkeley dissertation.

Hansson, Gunnar Ó. 2008. Effects of contrast recoverability on the typology of harmony systems. *Contrast in phonology: Theory, perception, acquisition*, 115-41.

Hansson, Gunnar Ó. 2010. *Consonant Harmony: Long-Distance Interaction in Phonology*. (University of California Publications in Linguistics, 145.) Berkeley, CA: University of California Press.

Hansson, Gunnar Ó. & Patrick Moore. 2011. The phonetics of transparency in Kaska vowel harmony. In S. Tuttle & O. Lovick (eds.), *Working Papers in Athabaskan Languages 2010 [Proceedings of the 2010 Athabascan/Dene Languages Conference, Eugene, Oregon, June 25–27, 2010]*, Fairbanks, AK: Alaska Native Language Center, 122–138.

Harrison, David, Emily Thomforde & Michael O'Keefe. 2004. The vowel harmony calculator. Available from http://www.swarthmore.edu/SocSci/harmony/public_html/.

Harrison, Kevin David. 2000. *Topics in the Phonology and Morphology of Tuvan*. Doctoral dissertation, Yale University.

Haspelmath, Martin. 1995. *A grammar of Lezgian*. The Hague: Mouton de Gruyter.

Hayes, Bruce & Zsuzsa Cziráky Londe 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology* 23(1). 59–104. DOI: 10.1017/S095267570600076

Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic inquiry*, *39*(3), 379-440.

Hayes, Bruce, Kie Zuraw, Péter Siptár & Zsuzsa Londe 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85(4). 821–862. https://www.jstor.org/stable/40492955

Howe, Darin & Douglas Pulleyblank. 2004. Harmonic scales as faithfulness. *Canadian Journal of Linguistics/Revue canadienne de linguistique*, *49*(1), 1-49.

Hsu, Brian & Jesney, Karen. 2017. Loanword Adaptation in Québec French: Evidence for Weighted Scalar Constraints. In Aaron Kaplan, Abby Kaplan, Miranda K. McCarvel & Edward J. Rubin (eds.), *Proceedings of the 34th West Coast Conference on Formal Linguistics*, 249–258.

Hyman, Larry. 1985. *A Theory of Phonological Weight*. Foris: Dordrecht.

Hyman, Larry. 1999. The historical interpretation of vowel harmony in Bantu, in Jean-Marie Hombert & Larry Hyman (eds), *Bantu historical linguistics*. Stanford: CSLI Publications. 235-295.

Hyman, Larry M. 2002. Is there a right-to-left bias in vowel harmony? Ms, University of California, Berkeley.

Inkelas, Sharon, Cemil O. Orgun & Cheryl Zoll. 1997. The implications of lexical exceptions for the nature of grammar. In Iggy Roca (ed.), *Derivations and constraints in phonology*, 393–418. Oxford: Clarendon Press.

Inkelas, Sharon & Cemil O. Orgun. 2003. Turkish stress: a review. *Phonology*, *20*(1), 139-161.

Inkelas, Sharon & Cheryl Zoll. 2007. Is grammar dependence real? A comparison between cophonological and indexed constraint approaches to morphologically conditioned phonology. *Linguistics* 45(1). 133–172. DOI: 10.1515/LING.2007.004

Itô, Junko & Armin Mester. 1995. The core-periphery structure of the lexicon and constraints on reranking. In Jill N. Beckman, Laura Walsh Dickey & Suzanne Urbanczyk (eds.), *Papers in Optimality Theory*, University of Massachusetts Occasional Papers in Linguistics 18, 181–209. Amherst: University of Massachusetts.

Jacobson, Leon. 1980. Voice-quality harmony in Western Nilotic languages. In Robert Vago (ed.), *Issues in Vowel Harmony* 6. Amsterdam: John Benjamins, pp.183– 200.

Johanson, Lars. 2015. The structure of Turkic. *The Turkic Languages,* ed. Lars Johanson & Éva Á. Csató. New York : Routledge.

Jurgec, Peter. 2010. Disjunctive lexical stratification. *Linguistic Inquiry* 41(1). 149–161. DOI: 10.1162/ling.2010.41.1.149

Jurgec, Peter. 2011. *Feature spreading 2.0: A unified theory of assimilation*. Tromsø: University of Tromsø dissertation.

Jurgec, Peter & Bronwyn M. Bjorkman, B. 2018. Indexation to stems and words. *Phonology*, *35*(4), 577-615.

Kabak, Bariş. 2011. Turkish vowel harmony. *The Blackwell companion to phonology*, 1-24.

Kabak, Bariş & Irene Vogel. 2001. The phonological word and stress assignment in Turkish. *Phonology*, *18*(3), 315-360.

Kangasmaa-Minn, Eeva. 1998. Mari. *The Uralic languages*, ed. Daniel Abondolo. New York: Routledge.

Kaplan, Aaron F. 2008. Noniterativity is an emergent property of grammar. PhD dissertation, University of California, Santa Cruz.

Kaun, Abigail R. 1995. *The Typology of rounding harmony: an Optimality Theoretic approach*. Los Angeles, CA: UCLA dissertation.

Kaun, Abigail. 2004. The typology of rounding harmony. In Bruce Hayes, Robert Kirchner & Donca Steriade (eds.) *Phonetically based phonology,* 87–116. Cambridge: Cambridge University Press.

Kilpatrick, Eileen. 1985. Bongo Phonology. *Occasional Papers in the Study of Sudanese Languages 4*, 1-62.

Kimper, Wendell A. 2011. *Competing triggers: Transparency and opacity in vowel harmony*. Amherst: University of Massachusetts Amherst dissertation.

Kiparsky, Paul. 1973. Phonological representations. In Osamu Fujimura (ed.), *Three dimensions of linguistic theory*, 1–136. Tokyo: TEC.

Kiparsky, Paul & Karl Pajusalu. 2003. Toward a typology of disharmony. *The Linguistic Review*, 20. 217–241. DOI: 0167–6318/03/020-0217

Köhler, Oswin. 1948. Die nilotischen Sprachen. Darstellung ihres Lautsystems, nebst einer Einleitung über die Geschichte ihrer Erforschung, ihre Verbreitung und Gliederung. Ph. D. dissertation, Berlin.

Korn, David. 1969. Types of labial vowel harmony in the Turkic languages. *Anthropological Linguistics* 11 (3). 98–106. www.jstor.org/stable/30029212

Krämer, Martin. 2003. *Vowel harmony and correspondence theory*. Berlin: Walter de Gruyter.

Krueger, John R. 1961. *Chuvash manual*. Indiana University.

Kurisu, Kazutaka. 2001. *The phonology of morpheme realization*. University of California, Santa Cruz.

Kutsch Lojenga, Constance. 1991. Vowel harmony in Alur: On the crossroads of two systems. In Franz Rottland and Lucia N. Omondi (eds.), *Proceedings of the Third Nilo-Saharan Linguistics Colloquium*, 131–141. Hamburg: Dietrich Reimer.

Legendre, Géraldine, Yoshiro Miyata & Paul Smolensky. 1990a. Harmonic Grammar: a formal multi-level connectionist theory of linguistic well-formedness: theoretical foundations. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, 388–395. Hillsdale: Erlbaum, 388–395. DOI:10.1.1.20.4453

Legendre, Géraldine, Yoshiro Miyata & Paul Smolensky. 1990b. Harmonic Grammar: a formal multi-level connectionist theory of linguistic well-formedness: an application. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, 884–891. Hillsdale: Erlbaum. DOI:10.1.1.20.4266

Leitch, Myles (1996). *Vowel harmonies of the Congo basin: An optimality theory analysis of variation in the Bantu zone C*. Doctoral dissertation, University of British Columbia.

Li, Bing. 1996. *Tungusic Vowel Harmony: Description and Analysis*. Holland Institute of Generative Linguistics.

Lightner, Theodore. 1965. On the description of vowel and consonant harmony. *Word* 21: 244-250.

Lindau, Mona. 1987. Tongue mechanisms in Akan and Luo. *UCLA Working Papers in Phonetics* 68: 46–57.

Lindau, Mona E., Leon Jacobson & Peter Ladefoged. 1972. The feature advanced tongue toot. *UCLA Working Papers in Phonetics* 22: 76–94.

Lindau, Mona E. & Peter Ladefoged. 1986. Variability of feature specifications. In Joseph S. Perkell & Dennis H. Klatt (eds.), *Invariance and Variability in Speech Processes*, 464–

479. Hillsdale, N.J.: Erlbaum.

Lindblad, Vern M. 1990. *Neutralization in Uyghur*. Master's thesis, University of Washington.

Lodge, Ken. 1995. Kalenjin phonology and morphology: A further exemplification of underspecification and non-destructive phonology. *Lingua*, 96(1), 29-43.

Luutonen, Jorma, et al. (ed.). 2007. *Electronic Word Lists: Mari, Mordvin and Udmurt. With SFOu WordListTool 1.3*. Lexica Societatis Fenno-Ugricae XXXI:1. ISBN 978-952-5150-98-8. Helsinki.

Luutonen, Jorma, et al. (ed.). 2016. *Electronic Word Lists: Komi, Chuvash and Tatar. With SFOu WordListTool 1.4*. Lexica Societatis Fenno-Ugricae XXXI:2. ISBN 978-952-5667-79-0. Helsinki.

Mahanta, Shakuntala. 2008. *Directionality and locality in vowel harmony: with special reference to vowel harmony in Assamese*, PhD thesis, Utrecht University.

Majors, Tivoli. 1998. *Stress-dependent harmony.* Austin: University of Texas, Austin dissertation.

Martin, Andrew T. 2007. *The evolving lexicon*. PhD dissertation, University of California, Los Angeles.

McCarthy, John J. 2003. Comparative markedness. *Theoretical linguistics*, *29*(1-2), 1-51.

McCarthy, John. 2004. Headed spans and autosegmental spreading. ROA.

McCarthy, John J. & Alan Prince. 1995. Faithfulness and reduplicative identity. *Papers in Optimality Theory* 10. University of Massachusetts, Amhurst. https://scholarworks.umass.edu/cgi/viewcontent.cgi?article=1009&context=linguist_faculty_pubs

McCarthy, John J. & Alan Prince. 1999. Faithfulness and identity in prosodic morphology. *The prosody-morphology interface* 79. 218–309.

McCollum, Adam. 2019. *Gradience and locality in phonology: Case studies from Turkic vowel harmony*. San Diego: University of California San Diego dissertation.

Muratori, Carlo 1938. *Grammatica Lotuxo*. Verona: Missioni Africane.

Nevins, Andrew. 2010. *Locality in Vowel Harmony.* Cambridge, MA: MIT Press.

Nevins, Andrew & Bert Vaux. 2004. Consonant harmony in Karaim. *MIT Working Papers in Linguistics*, *46*, 175-194.

Ní Chiosáin, Máire & Jaye Padgett. 2001. Markedness, segment realization, and locality in spreading. In Linda Lombardi (ed.), *Constraints and representations: Segmental phonology in Optimality Theory*, 118–156. Cambridge: Cambridge University Press.

Noonan, Michael. 1992. *A grammar of Lango*. Berlin & New York: Mouton de Gruyter.

Noske, Manuela. 1990. Vowel harmony in Turkana. *Studies in the Linguistic Sciences*, *20*(1), 123-134.

Noske, Manuela. 1996. [ATR] harmony in Turkana. *Studies in African Linguistics*, *25*(1), 61.

Ohala, John J. 1994. Towards a universal, phonetically-based, theory of vowel harmony. In *Third International Conference on Spoken Language Processing*, 491–494. Yokohama, Japan,

Oyelaran, Olasope O. 1973. Yoruba vowel co-occurrence restrictions. *Studies in African Linguistics* 4: 155–182.

Otero, Manuel A. 2015. [+ ATR] Dominant Vowel Harmony Except When It's Not? Evidence from Ethiopian Komo. In *44th Annual Conference on African Linguistics* (pp. 212-220). Cascadilla Proceedings Project.

Ozburn, Avery. 2019a. A target-oriented approach to neutrality in vowel harmony: Evidence from Hungarian. *Glossa: a journal of general linguistics* 4(1): 47. 1–36. DOI: https://doi.org/10.5334/gjgl.681

Ozburn, Avery. 2019b. A segment-specific metric for quantifying participation in harmony. In Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose & Matthew Zaslansky (eds.), *Supplemental Proceedings of the 2018 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America.

Ozburn, Avery. To appear. An analysis of ATR harmony in Alur. In *Proceedings of the 2019 Canadian Linguistics Association Conference.*

Ozburn, Avery, Samuel Akinbo, Alexander Ansongna, Murray Schellenberg, and Douglas Pulleyblank. 2018. Dagaare [a] is not neutral to ATR harmony. *Canadian Acoustics* 46(4): 36-37.

Ozburn, Avery, and Alexei Kochetov. 2018. Ejective harmony in Lezgian. *Phonology* 35(3): 407-440.

Parker, Kirk. 1985. Baka phonology. *Occasional Papers in the Study of Sudanese Languages*, *4*, 63-85.

Pater, Joe. 2007. The locus of exceptionality: Morpheme−specific phonology as constraint indexation. In Leah Bateman, Michael O'Keefe, Ehren Reilly & Adam Werle (eds.), *University of Massachusetts occasional papers in linguistics 32: Papers in Optimality Theory III*, 259-296. Amherst: GLSA, University of Massachusetts. Available on Rutgers Optimality Archive, ROA 866, http://roa.rutgers.edu.

Pater, Joe. 2009. Weighted constraints in generative linguistics. *Cognitive Science* 33(6). 999–1035. DOI: 10.1111/j.1551-6709.2009.01047.x

Pearce, Mary. 2006. The interaction between metrical structure and tone in Kera. *Phonology* 23: 259–286.

Pierrehumbert, Janet B. 1993. Dissimilarity in the Arabic verbal roots. *Proceedings of NELS 23*: 367–381.

Polgárdi, Krisztina. 1998. *Vowel Harmony: An Account in Terms of Government and Optimality*. The Hague: Holland Academic Graphics.

Poppe, Nikolaus. 1960. *Vergleichende Grammatik der altaischen Sprachen*. Wiesbaden: Otto Harrassowitz.

Potts, Christopher, Joe Pater, Karen Jesney, Rajesh Bhatt & Michael Becker. 2010. Harmonic Grammar with linear programming: from linear systems to linguistic typology. *Phonology* 27(1). 77–117. DOI:10.1017/S0952675710000047

Prince, Alan & Paul Smolensky. 1993. *Optimality Theory: Constraint interaction in Generative Grammar.* Linguistic Inquiry Monograph. MIT Press.

Przezdziecki, Marek A. 2005. *Vowel harmony and coarticulation in three dialects of Yoruba: Phonetics determining Phonology*. Ph.D. dissertation, Cornell University, Ithaca, N.Y.

Pulleyblank, Douglas. 1996. Neutral vowels in Optimality Theory: A comparison of Yoruba and Wolof. *Canadian Journal of Linguistics* 41(4). 295–347. DOI: 10.1017/S0008413100016601

Pulleyblank, Douglas. 2001. Nothing follows: harmony & disharmony. Northwest Linguistics Conference, University of British Columbia.

Pulleyblank, Douglas 2004. Harmony drivers: No disagreement allowed. In Julie Larson & Mary Paster (eds.), *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, 249–267. Berkeley, CA: Berkeley Linguistics Society.. DOI: 10.3765/bls.v28i1.3841

Pulleyblank, Douglas, Ping Jiang-King, Myles Leitch & Nike Ola. 1995. Typological variation through constraint rankings: low vowels in tongue root harmony. *Proceedings of South Western Optimality Theory Workshop*. The University of Arizona Coyote Papers, 184–208. https://arizona.openrepository.com/bitstream/handle/10150/227278/apc-v-184-208.pdf;jsessionid=FBE165F7213AAD0DE6B7AF9575DD179B?sequence=1

Quinn-Wriedt, Lindsey. 2013. *Vowel harmony in Maasai*. PhD Dissertation, University of Iowa.

Rebrus, Péter, Péter Szigetvári & Miklós Törkenczy 2012. Dark secrets of Hungarian vowel harmony. In Eugeniusz Cyran, Bogdan Szymanek & Henryk Kardela (eds.), *Sound, structure and sense: Studies in memory of Edmund Gussmann,* 491–508. Lublin Wydawnictwo KUL,

Rebrus, Péter, Péter Szigetvári & Miklós Törkenczy. 2017. Asymmetric variation. In Jeff Lindsey & Andrew Nevins (eds.), *Sonic signatures*, 163–187. Amsterdam: John Benjamins. http://budling.nytud.hu/~tork/owncikk/Yodharmony_prepub_version.pdf.

Rebrus, Péter & Miklós Törkenczy. 2015. Monotonicity and the typology of front/back harmony. *Theoretical Linguistics* 41(1-2). 1–61. DOI: 10.1515/tl-2015-0001

Rebrus, Péter & Miklós Törkenczy. 2016a. Types and degrees of vowel neutrality. *Linguistica* 56. 239–251. DOI: 10.4312/linguistica.56.1.239-252

Rebrus, Péter & Miklós Törkenczy. 2016b. A Non-cumulative pattern in vowel harmony: A frequency-based account. In Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin & Douglas Pulleyblank (eds.), *Proceedings of the 2015 Annual Meeting on Phonology*, Linguistic Society of America, Washington, DC. http://budling.nytud.hu/~tork/owncikk/Vancouver_pap_3692-4832-1-PB.pdf.

Reh, Mechthild. 1996. *Anywa language. Description and internal reconstruction*. Cologne:

Rüdiger Köppe.

Rialland, Annie and Redouane Djamouri. 1984. Harmonie vocalique, consonantique et structures de dépendance dans le mot en mongol khalkha. *Bulletin de la Société de Linguistique de Paris* 89:333-383.

Ringen, Catherine. 1975. Vowel harmony: theoretical implications. Bloomington: Indiana University dissertation. Published 1988, New York: Garland.

Ringen, Catherine. 1978. Another view of the theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry* 9(1). 105–115. www.jstor.org/stable/4178039

Ringen, Catherine. 1980. A concrete analysis of Hungarian vowel harmony. In Robert Vago (ed.), *Issues in vowel harmony*, 135–154. Amsterdam: John Benjamins.

Ringen, Catherine. 1988. Transparency in Hungarian vowel harmony. *Phonology 5*(2). 327–342. http://www.jstor.org/stable/4419984

Ringen, Catherine & Orvokki Heinämäki. 1999. Variation in Finnish vowel harmony: An OT account. *Natural Language & Linguistic Theory* 17(2). 303–337. DOI: 10.1023/A:1006158818498

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

Ringen, Catherine & Robert M. Vago. 1998. Hungarian vowel harmony in Optimality Theory. *Phonology* 15(3). 393–416. https://www.cambridge.org/core/services/aop-cambridge-core/content/view/A6FC6BF883046CFB910EB6A6B084EBF1/S0952675799003632a.pdf/div-class-title-hungarian-vowel-harmony-in-optimality-theory-div.pdf

Ritchart, Amanda & Sharon Rose. 2017. Moro vowel harmony: implications for transparency and representations. *Phonology*, *34*(1), 163-200.

Rose, Sharon. 2018. ATR Vowel Harmony: new patterns and diagnostics. In Gillian Gallagher, Maria Gouskova, Sora Heng Yin (eds.), *Proceedings of the 2017 Annual Meeting on Phonology*. Washington, DC: Linguistic Society of America.

Rose, Sharon & Lisa King. 2007. Speech error elicitation and co-occurrence restrictions in two Ethiopian Semitic languages. *Language and Speech* 50(4): 451-504.

Rose, Sharon, Michael Obiri-Yeboah & Sarah Creel. 2019. Perception of ATR vowel contrasts by Akan speakers. Talk presented at the 50th Annual Conference on African Linguistics, University of British Columbia.

Rose, Sharon, and Rachel Walker. 2004. A typology of consonant agreement as correspondence. *Language* 80(3): 475–531.

Rottland, Franz. 1982. *Die Südnilotischen Sprachen. Beschreibung, Vergleichung und Rekonstruktion*. (Kölner Beiträge zur Afrikanistik, 7.) Berlin: Dietrich Reimer.

Schlindwein, Deborah. 1987. P-bearing units: a study of Kinande vowel harmony. In *Proceedings of NELS* (Vol. 17, pp. 551-567).

Siptár, Péter. 2015. Neutrality and variation: What are they? *Theoretical Linguistics* 41(1-2). 97–103. DOI: 10.1515/tl-2015-0006.

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

Smolensky, Paul. 2006. Optimality in phonology II: harmonic completeness, local constraint conjunction, and feature domain markedness. In Smolensky & Legendre (vol. 2). 27–160.

Smolensky, Paul & Matthew Goldrick. 2016. Gradient symbolic representations in grammar: The case of French liaison. *Ms. Johns Hopkins University and Northwestern University. Available as ROA*, *1286*.

Smolensky, Paul & Géraldine Legendre. 2006. *The harmonic mind: From neural computation to optimality-theoretic grammar* vols. 1-2. Cambridge, MA: MIT Press.

Spagnolo, Lorenzo M. 1933. *Bari Grammar*. Verona: Missioni Africane.

Starwalt, Coleen 2008. *The Acoustic correlates of ATR harmony in seven- and nine-vowel African languages: a phonetic inquiry into phonological structure*. PhD Dissertation, University of Texas, Arlington.

Steinberger, Kathy E. & Robert Vago. 1987. A multileveled autosegmental analysis of vowel harmony in Bari. *Current approaches to African linguistics*, 4, 357-368.

Stewart, John M. 1967. Tongue root position in Akan vowel harmony. *Phonetica* 16: l85–204.

Suomi, Kari. 1983. Palatal vowel harmony: a perceptually motivated phenomenon? *Nordic Journal of Linguistics* 6, 1-35.

Svantesson, Jan-Olof. 1985. Vowel harmony shift in Mongolian. *Lingua* 67.4:283-329.

Svantesson, Jan-Olof, Anna Tsendina, Anastasia Karlsson & Vivan Franzén. 2005. *The phonology of Mongolian*. Oxford University Press.

Tiede, Mark K. 1996. An MRI-based study of pharyngeal volume contrasts in Akan and English. *Journal of Phonetics* 24: 399–421.

Törkenczy, Miklós. 1997. *Hungarian verbs and essentials of grammar*. Lincolnshire, Illinois: Passport Books.

Törkenczy, Miklós. 2010. Hungarian vowel harmony. In Marc van Oostendorp, Colin Ewen, Elizabeth Hume & Keren Rice (eds.), *The Blackwell Companion to Phonology*, 2963–2989. Oxford: Wiley-Blackwell.

Törkenczy, Miklós. 2013. The close-ups can get rough: exceptional behaviour, transparency and variation in Hungarian vowel harmony. *The Twenty-First Manchester Phonology*

Meeting, 23–25. Manchester, UK.

Törkenczy, Miklós, Péter Szigetvári & Péter Rebrus. 2013. Harmony that cannot be represented. In Johan Brandtler, Valéria Molnár & Christer Platzack (eds.), *Approaches to Hungarian 13: Papers from the 2011 Lund Conference,* 229–252. Amsterdam: John Benjamins,.

Tupper, Paul, and John Alderete. In prep. Risk ratios for linguists.

Ultan, Russell. 1973. Some reflections on vowel harmony. *Working Papers on Language Universals*, *12*, 37-67.

Vago, Robert. 1973. Abstract vowel harmony systems in Uralic and Altaic languages. *Language* 49(3). 579–605. DOI: 10.2307/412352

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

Vago, Robert. 1978. Some controversial questions concerning the description of vowel harmony. *Linguistic Inquiry* 9(1). 116–125. https://www.jstor.org/stable/4178040

Vago, Robert. 1980a. A critique of suprasegmental theories of vowel harmony. In Robert Vago (ed.), *Issues in vowel harmony*, 155–183. Amsterdam: John Benjamins.

Vago, Robert. 1980b. *The sound pattern of Hungarian*. Washington: Georgetown University Press.

van der Hulst, Harry. 1985. Vowel harmony in Hungarian: A comparison of segmental and autosegmental analysis. In Harry van der Hulst & Norval Smith (eds.), *Advances in nonlinear phonology*, 267-303. Dordrecht: Foris.

van der Hulst, Harry. 2016. Vowel harmony. In Mark Aronoff (ed.), *Oxford Research Encyclopedia of Linguistics*. DOI:10.1093/acrefore/9780199384655.013.38

van der Hulst, Harry. 2018. *Asymmetries in vowel harmony: A representational account*. Oxford

University Press.

van der Hulst, Harry & Jeroen van de Weijer. 1995. Vowel harmony. In John Goldsmith (ed.), *The Handbook of Phonological Theory*, 495–534. Cambridge, MA: Blackwell.

Vaux, Bert. 2000. Disharmony and derived transparency in Uyghur Vowel Harmony. In NELS 30: *Proceedings of the North East Linguistic Society* 30, edited by Masako Hirotani, Andries Coetzee, Nancy Hall & Ji-yung Kim, volume 2, 671–698. Amherst, MA: GLSA.

Vaux, Bert. 2009. [atr] and [back] harmony in the Altaic languages. In Investigations into Formal Altaic Linguistics: *Proceedings of WAFL3*, edited by Sergei Tatevosov, pp. 50-67. Moscow: MAKS Press.

Vaysman, Olga. 2009. *Segmental alternations and metrical theory.* PhD dissertation, MIT.

Viitso, Tiit-Rein. 1998. Fennic. *The Uralic languages*, ed. Daniel Abondolo. New York: Routledge.

Vossen, Rainer. 1982. *The Eastern Nilotes: Linguistic and Historical Reconstruction*. (Kölner Beiträge zur Afrikanistik, 9.) Berlin: Dietrich Reimer.

Vroomen, Jean, Jyrki Tuomainen & Beatrice de Gelder. 1998. The roles of word stress and vowel harmony in speech segmentation. *Journal of Memory and Language*, *38*(2), 133-149.

Walker, Rachel. 1998. Nasalization, neutral segments, and opacity effects. Santa Cruz: University of California Santa Cruz dissertation.

Walker, Rachel. 2011. *Vowel patterns in language* (Vol. 130). Cambridge University Press.

Walker, Rachel. 2012. Vowel harmony in optimality theory. *Language and Linguistics Compass* 6(9). 575–592. DOI: 10.1002/lnc3.340.

Wilson, Colin & Marieke Obdeyn. 2009. Simplifying subsidiary theory: statistical evidence from

Arabic, Muna, Shona, and Wargamay. Ms, Johns Hopkins University.

Woock, Edith Bavin & Michael Noonan. 1979. Vowel harmony in Lango. *CLS* 15. 20–29.