Long-Distance Phonotactics as Tier-Based Strictly 2-Local Languages[*]

Kevin McMullin and Gunnar Ólafur Hansson

University of British Columbia

1 Introduction

This paper shows that the properties of locality observed for patterns of long-distance consonant agreement and disagreement belong to a well-defined and relatively simple class of subregular formal languages (stringsets) called the TIER-BASED STRICTLY 2-LOCAL (TSL$_2$) languages (Heinz et al. 2011), and argues that analyzing them as such has desirable theoretical implications. Specifically, treating the two elements of a long-distance dependency as adjacent segments on the computationally defined notion of a tier allows for a unified account of locality that necessarily extends to the cross-linguistically variable behavior of neutral segments (transparency and blocking). This result is significant in light of the long-standing and persistent problems that long-distance dependencies have raised for phonological theory, with current approaches still predicting several pathological patterns that have little or no empirical support.

The primary challenge posed by long-distance phonotactic dependencies is that in some languages they appear to hold with no upper limit on the number of segments that may intervene between the two interacting elements. Example (1) illustrates such a pattern, long-distance consonant agreement in Aari, an Omotic language of southern Ethiopia, which does not allow subsequences of [αant]…[−αant] sibilants within a word. Evidence for this is seen when the perfective suffix /-s/ surfaces as [-ʃ] when preceded by a [−ant] sibilant, as in (1b-e). For clarity, relevant segments are presented in boldface throughout this paper.

(1)   Unbounded sibilant harmony in Aari (Hayward 1990)
      a.   /ba**ʔ-s**-e/   ba**ʔse**   'he brought'
      b.   /**ʔuʃ-s**-it/   **ʔuʃʃit**   'I cooked'
      c.   /**tʃʼ**a̠ːq-**s**-it/   **tʃʼ**a̠ːq**ʃit**   'I swore'
      d.   /**ʃ**ed-er-**s**-it/   **ʃ**eder**ʃit**   'I was seen'
      e.   /**ʒ**a̠ːg-er-**s**-e/   **ʒ**a̠ːger**ʃe**   'it was sewn'

Though long-distance dependencies of this kind are seemingly more complex than string-adjacent co-occurrence restrictions (e.g., *CC or *VV), they are also known to be constrained in certain ways with respect to locality. For example, there is a robust dichotomy present in the typology of consonant harmony between UNBOUNDED patterns that hold at any distance (like that of Aari shown above) and TRANSVOCALIC harmony, which holds across at most a vowel, i.e. where no other consonant intervenes (Hansson 2010a). A transvocalic variant of sibilant harmony is illustrated in (2) with data from Koyra, another Omotic language related to Aari, in which the dependency holds only in transvocalic contexts.

(2)   Transvocalic sibilant harmony in Koyra (Koorete; Hayward 1982)
      a.   /tim-d-o**s**ːo/   tindo**s**ːo   'he got wet'
      b.   /pa**tʃ**-d-o**s**ːo/   pa**tʃ**ːo**ʃ**ːo   'it became less'
      c.   /**dʒ**a**ʃ**-us-e**s**ːe/   **dʒ**a**ʃ**u**ʃ**e**ʃ**ːe   'let him/them frighten!'
      d.   /**ʃ**od-d-o**s**ːo/   **ʃ**od:o**s**ːo   'he uprooted'   *ʃodːoʃːo
      e.   /**ʔatʃ**-ut-d-o**s**ːo/ **ʔatʃ**ut:o**s**ːo   'he (polite) reaped'   *ʔatʃutːoʃːo

Transvocalic consonant harmony can be thought of as a long-distance dependency that is *bounded* in terms of distance; the pattern still holds when two elements of the dependency are in a non-adjacent, but still relatively local …CVC… context. Beyond the transvocalic window, in …CV…C…VC… contexts, the restriction does not apply since the intervening material exceeds the transvocalic threshold for locality in that it contains one or more consonants. Of particular interest to this paper is the fact that there are no attested cases of a long-distance dependency that is bounded by any other measure (e.g., sibilant harmony across at most one consonant, two vowels, five segments, or any other metric of distance). Likewise, there are no examples of a language with a dependency that holds across exactly one intervening consonant (not more or less), or across at least one consonant (in *beyond-transvocalic* contexts). These typological generalizations, true of both long-distance consonant agreement (harmony) and disagreement (dissimilation), are summarized in the table below.

(3)   Typological properties of locality for long-distance consonant (dis)agreement

Locality	Status	CVCV	CVCVCV	CVCVCVCV
Unbounded	Attested	+	+	+
Transvocalic	Attested	+	−	−
At most 1 consonant	Unattested	+	+	−
Exactly 1 consonant	Unattested	−	+	−
At least 1 consonant	Unattested	−	+	+

In this paper, we first show (§2) how these cross-linguistic generalizations can be viewed as reflecting the range of locality relations that are possible within the Tier-based Strictly 2-Local class of formal languages. Furthermore, we argue (§3) that this formalization offers a better fit with the attested patterns than does the factorial typology of optimality-theoretic constraints employed in the Agreement by Correspondence model of long-distance consonant assimilation and dissimilation (Rose & Walker 2004, Hansson 2001/2010a, Bennett 2013). The latter generates a host of pathological patterns that are not only empirically unattested but also exceed the limits of computational complexity that otherwise seem to constrain natural-language phonotactics.

2   Non-Adjacent Phonotactics in Formal Language Theory

2.1   *Phonotactic Patterns as Subregular Stringsets*   Languages can be thought of as sets of grammatical words, whose members include only the sequences of sounds that are well-formed in the language. A phonotactic pattern is thus manifested as a restriction on the strings of segments that are permitted in the set. Strings that do not adhere to the pattern will be ungrammatical, and are therefore not members of the stringset. As the scope of this paper is limited to phonotactic complexity in particular, the terms *pattern* and *language* are used interchangeably in reference to the stringset (i.e., a formal language) that reflects the phonotactic grammar of a language.

A long-established property of phonological mappings (e.g., input strings to output strings) is that any pattern that can be described as an ordered set of rewrite rules belongs to the class of regular relations (Johnson 1972, Kaplan & Kay 1994). Consequently, as Rabin & Scott (1959) show, all stringsets generated by these relations (the surface phonotactics) are members of the regular region of the Chomsky hierarchy (Chomsky 1956). While certain syntactic processes result in relatively complex stringsets (of words rather than segments; Culy 1985, Shieber 1985, Kobele 2006), it turns out that nearly all phonological patterns are indeed regular, including long-distance consonant agreement and disagreement (Heinz 2010, Heinz et al. 2011, Payne 2014).

While attested phonotactic patterns are overwhelmingly regular, not every pattern that can be described as a regular stringset is attested in natural language (e.g., a dependency that holds between the first and last segments of a word; Lai 2012). However, the regular region can be further broken down into a hierarchy of well-studied subregular languages (McNaughton & Papert 1971, Rogers & Pullum 2011). Below, figure (4) illustrates this hierarchy with the subset classes presented on the right. Each language class is thus a proper subset of any class that is situated to the left and connected by a line.

(4)		The Subregular Hierarchy (figure adapted from Heinz et al. 2011)
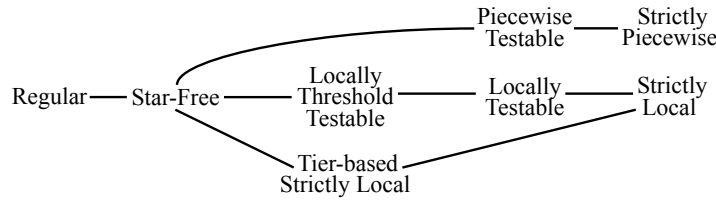
For example, restrictions on certain pairs of string-adjacent segments are members of Strictly 2-Local languages (SL_2; found at the bottom of the above hierarchy), as there is a prohibition against certain 2-factors (bigrams) of segments. Likewise, the transvocalic sibilant harmony pattern of Koyra in (2) is SL_3 since it can be described in terms of banned CVC substrings (*sVʃ, *ʃVs, and so on). By contrast, an unbounded consonant harmony pattern like that of Aari in (1) is not SL_k for any value of k because the dependency holds at all distances, including $k+1$. However, by referring not to string-adjacency but instead to precedence relations, which are by definition blind to distance and intervening material, unbounded harmony can be described as a Strictly 2-Piecewise (SP_2) pattern that disallows certain subsequences of length 2 (e.g., *s…ʃ, *ʃ…s).

	Heinz (2010) demonstrates that nearly all cases of unbounded consonant harmony from two typological surveys (Hansson 2001, Rose & Walker 2004) can be described as members of the SP_2 class of formal languages. However, the potential extension of this approach to all types of long-distance dependencies is impeded by a number of attested phonotactic patterns that are not SP_2. For example, certain cases of consonant dissimilation are known to exhibit blocking effects (Odden 1994, Heinz et al. 2011, Bennett 2013). In the case of Georgian liquid dissimilation, illustrated in (5), the ethnonymic suffix /-uri/ surfaces as [-uli] when preceded by [r] (5c-d), except if [l] intervenes (5e-f).

(5)		Georgian liquid dissimilation (Odden 1994)
	a.	dan-u**r**i			'Danish'
	b.	p'o**l**on-u**r**i		'Polish'
	c.	ung**r**-u**l**i			'Hungarian'
	d.	ap**r**ik'-u**l**i		'African'
	e.	avst'**ral**-u**r**i		'Australian'
	f.	ka**rtl**-u**r**i		'Kartvelian'

Note that the above pattern, akin to the well-known case of Latin liquid dissimilation (Jensen 1974, Odden 1994), cannot be classified as SP_2 with a ban against subsequences of *l…l as well as *r…r. Such an analysis would account for the basic generalization seen in (5a-d), but would fail to permit cases like (5e-f), which exhibit a blocking effect, as these still contain a supposedly-banned *r…r subsequence.

	The typological evidence thus suggests that a description of long-distance phonotactics in terms of precedence relations is too restrictive, accounting for only a subset of attested patterns. The remainder of this section pursues a different strategy, demonstrating that the Tier-based Strictly 2-Local class of formal languages (Heinz et al. 2011) not only encompasses each of the attested parameters of locality and blocking, but also excludes the unattested patterns presented in (3) above.

**2.2	*Tier-Based Strictly 2-Local Languages***	As a means of ignoring irrelevant intervening material, alternative to precedence relations, the pattern in (5) can instead be thought of as a restriction against contiguous [αlateral][αlateral] segment pairs (*ll, *rr), where adjacency is crucially assessed only among liquid consonants. Patterns that can be similarly described are members of the Tier-based Strictly 2-Local class of formal languages (TSL_2; Heinz et al. 2011). This characterization of long-distance dependencies, while still relatively simple, can account for several typological properties of locality.

	In more formal terms, TSL_2 languages can be defined as follows. For some alphabet Σ (a segment inventory), the grammar G of a Tier-based Strictly 2-Local language is a two-tuple $G=(T, \bar{S})$, where the tier T is some subset of Σ over which adjacency is assessed, and $\bar{S}$ is the set of bigrams prohibited on that tier (see Heinz et al. 2011 for an exhaustive formal definition of TSL languages and proofs for several

computational properties of the TSL class). For the Georgian liquid dissimilation pattern presented above in (5), the components of the TSL_2 grammar encoding the dependency are $T=\{l, r\}$ and $\bar{S}=\{*ll, *rr\}$.

Patterns of long-distance consonant agreement can equally be described in TSL_2 terms. The unbounded sibilant harmony pattern of Aari in (1) above is a restriction on sequences of [–anterior][+anterior] sibilants on a tier that includes all and only [+strident] segments, which can be generated by a TSL_2 grammar with the tier $T=\{s, ʃ, tʃˈ, ʒ, …\}$ and $\bar{S}=\{*ʃs, *tʃˈs, *ʒs, …\}$. Heinz (2010) hesitates to describe patterns of long-distance consonant agreement in terms of tiers, citing a lack of known systems that exhibit blocking. We argue however, that the widely attested locality type of transvocalic consonant harmony, such as the Koyra sibilant harmony shown in (2), can be recast as long-distance dependencies that are blocked *by any intervening consonant*. From the perspective of TSL_2 languages, the phonotactic grammar of Koyra bans sequences of [–ant][+ant] sibilants (exactly as in the unbounded case of Aari), but in Koyra violations are assessed on a tier that includes all of the consonants, rather than only sibilants. Forms like (2d) [ʃodːosːo] 'he uprooted' are grammatical precisely because another consonant, in this case [dː], remains present when the segment string is reduced to the consonants. The [dː], intervening between a pair of sibilants that constitutes a member of $\bar{S}$, can therefore be construed as a 'blocker', since it interrupts a sequence of sibilants that would otherwise violate the phonotactic grammar.

Further motivation for including long-distance consonant agreement within the scope of this approach is that additional cases have come to light that involve more obvious instances of consonant harmony with blocking. Relevant cases include some Berber dialects (Elmedlaoui 1995, Hansson 2010b), Kinyarwanda (Walker & Mpiranya 2005, Walker et al. 2008) and Slovenian (Jurgec 2011). The data in (6) show that the regressive sibilant harmony pattern found in Slovenian, which bans *[+ant]…[–ant] subsequences in (6a-c), is blocked when a coronal obstruent such as [t] or [d] intervenes, as in (6d-e). Note that in (6b) the coronal sonorants [n] and [l] are transparent, just like non-coronals are.

(6)   Sibilant harmony with coronal obstruent blocking in Slovenian (Jurgec 2011: 329-331)
   a. **s**pi   'sleeps'   **ʃ**pi-**ʃ**   '(you) sleep'
   b. **z**a-klɔn   'shelter'   **ʒ**a-klɔn-i**ʃtʃ**e 'bomb shelter'
   c. **ts**epəts   'fool'   **tʃ**eptʃ-ək   'fool-DIM'
   d. **s**i**t**   'full'   na-**s**it-i**ʃ**   '(you) feed'
   e. **z**i**d**a   '(s/he) builds'   **z**i**d**a-**ʃ**   '(you) build'

The distinction between unbounded dependencies, transvocalic dependencies, and dependencies with blocking can thus be attributed to a difference in the particular subset of Σ that comprises the designated tier T, rather than a change in the membership of $\bar{S}$. The grammars presented in (7) show this for three hypothetical languages with sibilant harmony, which are representative of the range of attested patterns (e.g., Aari, Slovenian, and Koyra, respectively). In order to facilitate a direct comparison between each type of language, we assume CV syllables, the segment inventory is restricted to $\Sigma=\{s, ʃ, p, t, a\}$, and the set of prohibited bigrams is always $\bar{S}=\{*sʃ, *ʃs\}$.

(7) TSL_2 grammars for three types of sibilant harmony

Type of Pattern	Σ	T	$\bar{S}$
Unbounded	$\{s, ʃ, p, t, a\}$	$\{s, ʃ\}$	$\{*sʃ, *ʃs\}$
Blocking	$\{s, ʃ, p, t, a\}$	$\{s, ʃ, t\}$	$\{*sʃ, *ʃs\}$
Transvocalic	$\{s, ʃ, p, t, a\}$	$\{s, ʃ, p, t\}$	$\{*sʃ, *ʃs\}$

With Σ and $\bar{S}$ held constant, the table in (7) shows that variation in (what appears to be) locality is merely a by-product of manipulating the contents of the relevant tier T over which violations are assessed. Unbounded dependencies are a result of a tier that includes only segments that are present in members of $\bar{S}$. If any additional segments are included in T, such as a coronal obstruent [t], these will block the pattern. The transvocalic locality type arises when all consonants are included in T, but the class of vowel segments is systematically absent from the tier. More generally, we note that the set of segments occurring in the members of $\bar{S}$ (such as $\{s, ʃ\}$ for the languages above) is a subset of T. Likewise, T is by definition a subset of Σ. Because of this property, there are overall just three possible types of segments. First, any segment that occurs in some member of $\bar{S}$ (and is therefore in both T and Σ) is a segment that participates in the

dependency and is a potential trigger or target for a repair strategy, such as harmony or dissimilation. Second, any other segment that is in T (and therefore also in Σ) will act as a neutral segment that blocks interaction between two segments on either side of it. Finally, a segment that is in Σ but not T (and therefore not present in any member of $\bar{S}$) will be neutral and transparent to the dependency.

We argued above that the attested types of locality for long-distance consonant (dis)agreement with and without blocking are easily captured as members of the Tier-based Strictly 2-Local class of formal languages. As a final note, we point out that none of the unattested types of locality shown in (3) are TSL_2. Dependencies that hold across either _at most_ or _exactly_ one consonant are TSL_3 since one must keep track of at least three consonants to prohibit words of the form *sVCVʃ, *ʃVCVs. The difference between the two is whether or not the grammar also prohibits words that include sVʃ and ʃVs substrings (i.e., sVʃVC, CVsVʃ, ʃVsVC, CVʃVs). A _beyond-transvocalic_ dependency, one that holds across _at least_ one consonant, is not TSL_k for any value of _k_, but instead is (Tier-based) Locally Testable at best (see §3.3 for details).

3 Locality Relations in Agreement by Correspondence

3.1 _Harmony as Agreement by Correspondence_ Long-distance phonotactics have long been a topic of debate in the literature of theoretical phonology (e.g., Jensen 1974, Steriade 1987a, 1987b, Odden 1994, Gafos 1999, Ní Chiosáin & Padgett 2001, Nevins 2010). While there exist many proposals for constraint-based analyses of consonant harmony as well as dissimilation, we focus here on the predictions of the Agreement by Correspondence framework (ABC; Rose & Walker 2004, Hansson 2001/2010a). This approach has seen relative success in accounting for the typology of consonant harmony, and has more recently been extended to analyses of long-distance consonant dissimilation (Bennett 2013, 2015; see §3.2 below). The ABC framework posits constraints that require pairs of segments to enter into a surface correspondence relation if they surpass some similarity threshold, due to sharing a certain set of feature values (e.g. CORR[+strid] requires all co-occurring sibilants in the output to be surface correspondents of one another). Other constraints (CC·Limiters; Bennett 2013) impose restrictions on surface-corresponding segments. Among these are CC-IDENT[F] constraints, which require corresponding consonants to agree with respect to some feature (e.g., CC-IDENT[ant] requires agreement in [±anterior]) and thus serve as potential triggers and targets of harmony for that feature. Other CC·Limiters that are relevant for the present argument include ones that have been proposed as an account for the locality dichotomy of unbounded vs. transvocalic consonant harmony systems as in (1)–(2). In particular, the constraint PROXIMITY (Rose & Walker 2004) penalizes correspondence for any pair of consonants that are not in adjacent syllables (later redefined as CC-SYLLADJ in Bennett 2013; see §3.2). Depending on the ranking of CORR and CC-IDENT constraints (and other CC·Limiters) relative to Faithfulness constraints (e.g., IO-IDENT[ant]), different variants of consonant harmony can be generated.

Tableau (8) shows a derivation of the Aari form in (1d) with an OT grammar for unbounded sibilant harmony. For expository reasons, in this and subsequent examples we abstract away from the thorny issue of directionality by restricting ourselves to cases involving [–ant]…[+ant] sequences like /ʃ…s/, and assume that an undominated IO-IDENT[–ant] (omitted from the tableaux) ensures that such sequences are always repaired by progressive assimilation, /ʃ…s/ → [ʃ…ʃ], rather than by regressive assimilation to [s…s]. In any case, since the focus in this paper is on the proper formal characterization of the phonotactics themselves (permitted vs. prohibited output sequences), the question of which repair strategy emerges as optimal, and how this gets determined in the grammar, is not directly relevant.

(8)   Aari: /ʃed-er-s-it/→[ʃed-er-ʃ-it] 'I was seen' (Hansson 2010a, Hayward 1990)

	/ʃed-er-s-it/	CORR[+strid]	CC-IDENT[ant]	IO-IDENT[+ant]	PROXIMITY
a.	ʃ$_x$e.der.s$_y$it	*!			
b.	ʃ$_x$e.der.s$_x$it		*!		*
c. ☞	ʃ$_x$e.der.ʃ$_x$it			*	*

Following standard practice, surface correspondence in output candidates is shown with matching subscript indices. Candidate (8a) loses as the two sibilants are not in correspondence, a violation of CORR[+strid]. Candidate (8b) is not optimal because its two sibilants, although in surface correspondence, have mismatched [±ant] specifications. The winner, then, is (8c), which incurs an IO-IDENT[+ant] violation

for the unfaithful mapping /s/ → [ʃ] in order to satisfy both the CORR and CC-IDENT constraints. With PROXIMITY ranked below CORR[+strid], the harmony will be enforced across any number of intervening segments or syllables. Note, however, that the constraint ranking in (8) would not be a suitable analysis of a transvocalic harmony pattern such as that exhibited by Koyra in (2), since in such cases harmony must be prevented from applying across an intervening syllable. We turn to transvocalic harmony in the next section, as it is one of several areas where the apparent formal simplicity and elegance of the ABC model turns out to be deceptive, and where various amendments and stipulations seem to be required in order to avoid undesirable pathological predictions.

3.2 *Pathological ABC Patterns*   In this section, we point out that the factorial typology generated by the ABC model, especially in its simplest and original form, includes several types of phonotactic patterns that are unattested in natural language. Our main focus here is on issues concerning trigger-target locality relations, but the typological predictions of ABC include other oddities as well (Hansson 2014).

   Consider, first, how transvocalic (or syllable-adjacent) harmony, such as the Koyra pattern in (2), is generated.[1] At first glance, this appears easily captured by ranking PROXIMITY above the CORR constraint responsible for placing trigger and target consonants into correspondence. Indeed, this is how Rose & Walker (2004) analyze the transvocalic nasal consonant harmony of Ndonga (in contrast to its unbounded counterpart in Kongo). In situations analogous to example (8) above, where the two sibilants are in non-adjacent syllables, correspondence as in (8b-c) would be prohibited by top-ranked PROXIMITY, and the faithful non-harmony (and non-correspondence) candidate (8a) [ʃₓe.der.sᵧit] would emerge as optimal.

   In words where the two consonants are in adjacent syllables, by contrast, CORR[+strid] can be satisfied without violating higher-ranked PROXIMITY; a correspondence relation is therefore established and agreement in [±ant] is enforced over that relation, as illustrated in (9).

(9)   Transvocalic harmony with one trigger, one target

	/paʃasa/	PROXIMITY	CORR[+strid]	CC-IDENT[ant]	IO-IDENT[+ant]
a.	pa.ʃₓa.sᵧa		*!		
b.	pa.ʃₓa.sₓa			*!	
c. ☞	pa.ʃₓa.ʃₓa				*

   However, since each collection of surface correspondents within an output form constitutes a set (an equivalence class; see Bennett 2013), in which every member is a correspondent of every other member, things quickly become complex once the number of correspondents goes above two. Tableau (10) shows how things go wrong for cases involving three consonants of the relevant class. Here each local pair of consonants is in adjacent syllables (straddling just a vowel), and we should therefore expect harmony to apply in a stepping-stone fashion; this is indeed what happens in real cases of transvocalic consonant harmony. Instead, the same constraint ranking that enforces harmony in a /…ʃVsV(CV)…/ sequence like in (9) will fail to do so in a /…ʃVsVsV…/ sequence.

(10)   Transvocalic harmony with one trigger, two potential targets: predicted non-harmony

	/paʃasasa/	PROXIMITY	CORR[+strid]	CC-IDENT[ant]	IO-IDENT[ant]
a.	pa.ʃₓa.sᵧa.s_za		***!		
b.	pa.ʃₓa.sₓa.sₓa	*!		**	
c. ☹	pa.ʃₓa.ʃₓa.ʃₓa	*!			**
d.	pa.ʃₓa.ʃₓa.sᵧa		**		*!
e. ☞	pa.ʃₓa.sᵧa.sᵧa		**		

---

[1] Although most ABC analyses define the transvocalic vs. unbounded dichotomy in terms of syllable-adjacency (following Odden 1994) rather than absence of intervening consonants (Hansson 2010a), we will continue to refer to the relevant short-range contexts as "transvocalic" even when discussing such analyses. The differences are subtle, and the predictions are indistinguishable as long as only …CV.CV.CV… strings are involved. The typological evidence is rather ambiguous on this point but, if anything, favours the transvocalic as opposed to syllable-adjacent characterization (see Hansson 2010a:175–178 for discussion).

The problem in situations like (10) is that in a chain of two or more transvocalic pairs of consonants of the relevant type, placing all of the consonants in correspondence will always result in a violation of PROXIMITY, since at least two of the correspondents will necessarily be separated by one or more intervening syllables, such as the first and third sibilants in (10b–c). Given the ranking PROXIMITY >> CORR (which as explained above is the defining property of transvocalic harmony), the optimal resolution is to leave one of the consonants out of correspondence. The choice of which consonant to leave out falls to lower-ranked considerations such as Faithfulness, as can be seen in (10d) vs. (10e). Similarly, an input like /paʃaʃasa/ will surface without harmony, because [pa.ʃ$_x$a.ʃ$_x$a.s$_y$a], with the first two sibilants in correspondence, will do better on Faithfulness than [pa.ʃ$_x$a.ʃ$_y$a.ʃ$_y$a] with correspondence (and hence harmony) between the second and third sibilants. The crucial factor is thus not the number of potential harmony targets but the overall number of sibilants in the sequence.

In fact, the nature of the pathology is even more bizarre, in that the key criterion is the *parity* of that number, where the predictions for the (non)application of harmony are different for odd-parity vs. even-parity cases. With PROXIMITY outranking CORR, an even number of potentially-harmonizing consonants is best partitioned into a series of adjacent-syllable correspondence pairs (…C$_x$V.C$_x$V.C$_y$V.C$_y$V…). Harmony is therefore predicted to be enforced in such words only between the 1st and 2nd consonant in the sequence, as well as in the 3rd and 4th consonant, etc., whereas there should be no requirement for harmony between the 2nd and 3rd, or the 4th and 5th (etc.) consonants. For an odd number of consonants, by contrast, the optimal correspondence configuration is for one consonant to stand outside of correspondence with any of the others, and for the (even-parity) sets of consonants on either side of that consonant to be partitioned into individual, harmonizing correspondence pairs as described above. Just as in (10), the determination of which consonant is the "odd man out" will fall to lower-ranked considerations such as Faithfulness. Needless to say, no natural language displays anything remotely resembling such a sound pattern.

In part to avoid this problem, Bennett (2013) proposes a redefined version of PROXIMITY, called CC-SYLLADJ, the evaluation of which is limited to those pairs of surface-corresponding consonants that are "local", in the sense of not being separated by a member of the same correspondence class. Thus, in a sequence …C$_x$…C$_x$…C$_x$… the C$_1$↔C$_2$ and C$_2$↔C$_3$ correspondent pairs are subject to CC-SYLLADJ but not the C$_1$↔C$_3$ pair. This avoids the pathologies described above, by eliminating the PROXIMITY violation of (10b–c) such that (10c) emerges as the optimal output. A similar modification to the way CC·Limiters are evaluated has been proposed by Hansson (2007) for the CC-IDENT[F] constraints that drive harmony. The motivation, again, is to avoid pathological predictions of factorial typology, including harmony systems where "majority rule" (Lombardi 1999; Baković 2000) determines the directionality of assimilation, or even the status of an intervening segment as opaque vs. transparent to the harmony, on a word-by-word basis. It is notable that in both cases, the proposed modifications are in the direction of treating collections of surface-correspondents as if they were *tiers* (subsequences of the output string), rather than as the unordered *sets* (equivalence classes) defined by the formal correspondence relation.

In relation to the dichotomy between transvocalic (syllable-adjacent) and unbounded dependencies, another area where ABC makes typological predictions of questionable validity is with respect to long-distance dissimilation (Bennett 2013). In the ABC model, dissimilation emerges as a strategy to avoid the need for two consonants to stand in correspondence. By becoming less similar, the pair of consonants evades the scope of a high-ranked CORR constraint that would otherwise require them to be correspondents. Such avoidance can be driven by one or more high-ranked CC·Limiter constraints on surface correspondence configurations, if these cannot be satisfied in any other way. This idea, that dissimilation happens only in situations where correspondence is penalized, whereas harmony can only take place where it is permitted (since correspondence is the vehicle for agreement), gives rise to Bennett's (2013) "mismatch prediction" regarding the typology of these two families of sound patterns. That is, contexts that favour harmony should, other things being equal, be ones that fail to trigger dissimilation, and vice versa. Given that one of the CC·Limiters in the theory is the aforementioned constraint CC-SYLLADJ (or PROXIMITY), which penalizes any pair of correspondents that are located too far apart, that constraint too should be able to either facilitate harmony, as we have seen, or trigger dissimilation. As a result, we expect a typological mismatch with respect to the locality patterns possible under consonant harmony and dissimilation, respectively.

In particular, the ABC model predicts the possibility of dissimilation patterns that are strictly "beyond-transvocalic" (or, more accurately, "beyond-syllable-adjacent"), applying in exactly the complement set of

environments to what we have seen for transvocalic consonant harmony. The tableaux in (11)–(12) illustrate this schematically. The hypothetical case in question involves surface correspondence between obstruents (CORR[-son]) and a demand that surface correspondents agree in voicing (CC-IDENT[voi]). The latter requirement can in principle be satisfied either directly, by voicing harmony under correspondence (/p…b/ → [p…p], violating IO-IDENT[voi]), or indirectly, by manner dissimilation out of correspondence (/p…b/ → [p…m], violating IO-IDENT[son]). However, as (11) shows, CC-SYLLADJ will, when ranked highly enough, trigger dissimilation on its own accord in beyond-transvocalic obstruent pairs, even if CC-IDENT[voi] is too low-ranked to play any active role. For transvocalic/syllable-adjacent pairs of consonants as in (12), on the other hand, correspondence is permitted and dissimilation is therefore not triggered. The result is dissimilation in beyond-transvocalic environments only. (As the comparison between 12b and 12c shows, such a dissimilation pattern can coexist either with harmony or faithful non-interaction in transvocalic contexts, depending on the ranking of the relevant CC-IDENT and IO-IDENT constraints.)

(11)   CC-SYLLADJ as dissimilation trigger: beyond-transvocalic consonant pairs

/palaba/	CORR[-son]	CC-SYLLADJ	IO-IDENT[son]	IO-IDENT[voi]	CC-IDENT[voi]
a. p_xa.la.b_ya	*!				
b. p_xa.la.b_xa		*!			*
c. p_xa.la.p_xa		*!		*	
d. ☞ p_xa.la.m_ya			*		

(12)   CC-SYLLADJ as dissimilation trigger: no dissimilation in transvocalic pairs

/lapaba/	CORR[-son]	CC-SYLLADJ	IO-IDENT[son]	IO-IDENT[voi]	CC-IDENT[voi]
a. la.p_xa.b_ya	*!				
b. ☞ la.p_xa.b_xa					*
c. ☞ la.p_xa.p_xa				*	
d. la.p_xa.m_ya			*!		

Unfortunately for the ABC model, this predicted mismatch between consonant harmony and dissimilation is a very poor fit for the attested typology. The only case that exhibits anything resembling a strictly beyond-transvocalic pattern, Sundanese rhotic dissimilation (Bennett 2015), is replete with other complications which make it far less persuasive as a test case (infixing morphology, co-existence with lateral harmony, sensitivity to stem-initial vs. non-stem-initial position, root vs. affix affiliation and onset vs. coda status). Furthermore, artificial language learning experiments show that when exposed to a strictly beyond-transvocalic dependency, English adults fail to learn such a pattern and tend to instead interpret it as unbounded (Hansson & McMullin 2014); importantly, this is true for harmony and dissimilation alike, even though the latter is predicted to be possible by the ABC model.

Sound patterns in which dissimilation is limited to transvocalic consonant pairs, by contrast, are both amply attested in the cross-linguistic typology (Bennett 2013) and easily learned in the laboratory (Hansson & McMullin 2014). This lack of mismatch between harmony and dissimilation is again a problem for the ABC model, as no permutation of the constraint types discussed thus far is capable of generating a dissimilation that is confined to transvocalic contexts. To deal with this problem, Bennett (2013) is forced to augment the model with special "CVC" versions of the CORR constraints. These only call for correspondence in transvocalic consonant pairs; for example, CORR-CVC[-son] would penalize (12a) but not (11a). Replacing CORR[-son] with this constraint, and moving IO-IDENT[son] to the bottom of the ranking, would produce a transvocalic-only dissimilation pattern in (11)–(12), with (11a) and (12d) as the respective winners. Such CORR-CVC[αF] constraints had previously been advocated by Hansson (2001/2010a), but as an alternative to PROXIMITY/CC-SYLLADJ rather than complementary to it. The inclusion of both constraint types in the model creates an undesirable duplication of effort as well as rampant ambiguity of analysis, as practically every case of transvocalic consonant harmony can be interpreted either as involving the ranking CC-SYLLADJ >> CORR[αF] (with CORR-CVC[αF] ranked too low to be relevant) or else the undominated status of CORR-CVC[αF] (with the other two constraints low-ranked).

Other pathological predictions of factorial typology arise from the constraint machinery of the ABC

framework. For example, the assumption that surface correspondence is an equivalence relation (Bennett 2013), partitioning the set of co-occurring segments within the output into equivalence classes, can give rise to bizarre "agreement by proxy" effects (Hansson 2014). Two co-occurring segments that are normally not required to correspond (nor, therefore, to interact) can be forced to do so when a third segment is present somewhere in the word, provided that this third segment is sufficiently similar to each of the other two to force them into (covert) correspondence with itself, and thereby also with each other. For example, we can imagine a hypothetical obstruent voicing harmony, in which only homorganic obstruent pairs interact, and where the harmony is moreover strictly regressive, harmonizing [–voice]…[+voice] sequences to [+voice]…[+voice] (while leaving [+voice]…[–voice] sequences intact). An input like /sada/ is thus changed to [zada], whereas a form like /saga/ surfaces faithfully as [saga]. A pattern like this is straight-forwardly captured with the ranking in (13)–(14), where CORR[–son, αPlace] requires that co-occurring homorganic obstruents stand in correspondence and $C_R C_L$-IDENT[+voi] penalizes [–voi]…[+voi] sequences of surface correspondents.[2] The more general constraint CORR[–son], which demands correspondence for heterorganic as well as homorganic obstruent pairs, is ranked too low to have any effect.

(13)　　　Obstruent voicing harmony parasitic on place: homorganic pairs

/sada/	CORR [-son, αPlace]	$C_R C_L$- IDENT[+voi]	IO- IDENT[voi]	CORR [-son]
a.　　$s_x ad_y a$	*!			*
b.　　$s_x ad_x a$		*!		
c. ☞　$z_x ad_x a$			*	

(14)　　　Obstruent voicing harmony parasitic on place: heterorganic pairs

/saga/	CORR [-son, αPlace]	$C_R C_L$- IDENT[+voi]	IO- IDENT[voi]	CORR [-son]
a. ☞　$s_x ag_y a$				*
b.　　$s_x ag_x a$		*!		
c.　　$z_x ag_x a$			*!	

Let us now imagine that the same language has another highly ranked CORR constraint, one which demands that segments that agree in both manner ([±continuant]) and voicing must also stand in surface correspondence to one another: CORR[αcont, βvoi]. In cases like (13)–(14), such a constraint has no bearing on the outcome, as all the relevant output candidates vacuously satisfy it. However, in a case like (15), where the same kind of /s…g/ sequence as in (14) co-occurs with a /x/ somewhere else in the word, we see how the mere presence of this /x/ causes regressive harmony to be triggered in the /s…g/ sequence.

(15)　　　Transitive correspondence relation causes "agreement by proxy"

/sagaxa/	CORR [-son, αPlace]	CORR [αcont, βvoi]	$C_R C_L$- IDENT[+voi]	IO- IDENT[voi]	CORR [-son]
a.　　$s_x ag_y ax_z a$	*!	*!			***
b.　　$s_x ag_y ax_y a$		*!			**
c.　　$s_x ag_y ax_x a$	*!				**
d.　　$s_x ag_x ax_x a$			*!		
e. ☞　$z_x ag_x ax_x a$				*	
f.　　$z_x ag_y ax_y a$				*	*!*

For a [s…g…x] sequence like in (15a–d), one CORR constraint requires [g…x] to be in correspondence (both are velar obstruents) while the other requires the same of [s…x] (both are voiceless fricatives). The

---

[2] For this approach to directionality effects, see Rose & Walker (2004). A constraint like $C_R C_L$-IDENT[+voi] states that for any pair of surface correspondents (C_L, C_R) where C_L linearly precedes C_R (hence the mnemonic "L" vs. "R" labels), if C_R is [+voice] then C_L must also be [+voice].

only way to satisfy both CORR constraints is to place all three segments into the same correspondence class (15c–d), but this means that [s…g] are also in correspondence with each other, unlike in cases like (14). Consequently, the generalization is that the regressive voicing harmony applies to heterorganic obstruent pairs if and only if the word also happens to contain a third obstruent that agrees in place with one but in manner with the other. Nothing resembling this kind of pattern has ever been attested, to our knowledge, and this is unlikely to be an accidental gap. Note, finally, that restricting the evaluation of CC-IDENT[F] to segment pairs that are adjacent in the "correspondence chain" (Hansson 2007; cf. discussion above and in §3.3) is of no help here; the [s…g] sequence is "chain-adjacent" in (15d) just as in (13b)/(14b).

3.3   *Typological Predictions and Formal Complexity*   The unusual varieties of long-distance dependency patterns that are predicted to be possible by the basic architecture of the ABC model are pathological not merely in terms of typological attestation but also from the standpoint of computational complexity and learnability. For example, majority-rule harmony patterns are known to fall outside the class of regular relations (Heinz & Lai 2013), and indeed fail to be learned in a laboratory setting (Finley & Badecker 2008). As for the bizarre and unattested parity-sensitive harmony pattern generated by high-ranked PROXIMITY, as described under (10) above, its status is unclear to us at present (we suspect it is regular but not star-free), but it in any case falls beyond the SL, SP or TSL subregular classes. Also outside those classes are the "agreement by proxy" effects discussed in the previous section. For example, the bizarre pattern in (13)–(15) where assimilation in /s…g/ is dependent on a nearby /x/ cannot be expressed in TSL_k terms for any k, even if the tier T is defined as {s, g, x}; this is because the pattern holds no matter how many additional instances of [g] (or [s]) intervene between the [s…g] pair and the "proxy" [x].

As described in §2.3, unbounded consonant harmony and dissimilation are both easily construed in TSL_2 terms, as are the blocking effects that are sometimes observed in such patterns (less easily generated in the ABC model, though see Hansson 2007, 2010b and Bennett 2013). Transvocalic dependencies are likewise TSL_2, requiring merely that the relevant tier T be defined as comprising all consonants. Indeed, the move to include CORR-CVC[αF] constraints in the ABC model (see §3.2) indirectly incorporates this notion of a "consonant tier", in that the presence of any intervening consonant nullifies the demand for correspondence (and thereby the motivation for harmony or dissimilation) between the segments of interest. Furthermore the abovementioned modification to the evaluation of CC·Limiter constraints like CC-SYLLADJ (Bennett 2013) or CC-IDENT[F] (Hansson 2007), such that these ignore pairs of correspondents that are separated by one or more intervening correspondents, is roughly equivalent to treating correspondence sets as tiers, and evaluating agreement only for tier-adjacent pairs (≈ 2-factors on tier T).

By contrast, the strictly beyond-transvocalic dissimilation pattern generated in the ABC model is not TSL_2, nor indeed TSL_k for any k, as the phonotactic legality of a word cannot be determined solely in terms of presence vs. absence of individual k-factors regardless of how the tier T is construed. In the illustrative example in (11)–(12), sequences like …pVbV… or …bVpV… are permitted, but ones like *…pVCVbV… or *…bVCVpV… are prohibited. Crucially, the latter cannot be ruled out simply by setting T as including all consonants and defining $\bar{S}$ as comprising the 3-factors {*pCb, *bCp}. This is because the dissimilation holds regardless of the number of intervening consonants; there is thus no upper bound on k that will suffice to rule out all illegal sequences. Instead, a phonotactic pattern of this type falls into (a tier-based instantiation of) the Locally Testable class (Rogers & Pullum 2011), which is defined in terms of Boolean operations over sets of k-factors. Here an illegal word is one that contains *both* a member of {#p, Cp} *and* a member of {#b, Cb} among its 2-factors on the consonantal tier.[3] In this respect, the computational status of this phonotactic pattern is somewhat analogous to the (unattested) "first-last harmony" pattern described by Lai (2012), where words of the structure #s…ʃ# or #ʃ…s# are banned. The evidence from artificial language learning experiments, showing a failure to learn first-last harmony (Lai 2012) as well as strictly beyond-transvocalic dissimilation and harmony alike (Hansson & McMullin 2014) strongly suggest that Locally Testable patterns that lie outside the TSL (and SP) regions of the subregular hierarchy are beyond

---

[3] To be precise, this beyond-transvocalic pattern lies outside even the Locally Testable class, given that identical-consonant pairs are presumably also subject to the dissimilation (i.e. *…pVCVpV…, *…bVCVbV…). This means that the grammar would need to be able to count the *number of instances* of certain k-factors; in this case, a word is illegal if it contains *two or more occurrences of* any member of {#p, Cp, #b, Cb} on the consonant tier. The relevant class is therefore (a tier-based instantiation of) the Locally Threshold Testable languages (Rogers & Pullum 2011).

the grasp of the human phonological learner.

A final concern regarding the ABC model is how it often requires surprisingly roundabout ways to capture phonotactic dependency patterns that are exceedingly simple from a formal perspective. This is especially true in the case of dissimilation that applies to identical consonants, such as where /l…l/ surfaces as $[l_x…r_y]$ in order to avoid the $[l_x…l_x]$ correspondence configuration demanded by a high-ranked constraint like CORR[+lat] (Bennett 2013). Since a total-identity pair like $[l_x…l_x]$ satisfies all CC-IDENT[F] constraints, the motivation to avoid correspondence must come from some other CC·Limiter constraint(s). But since each such constraint penalizes correspondence only in some specific configuration—such as when the two segments belong to different morphemes, different syllables, have different syllable constituency affiliations, etc.—it becomes surprisingly hard to capture a pattern whereby the dissimilation /l…l/ → [l…r] happens in *all* contexts. The only solution is to blame the exceptionless, across-the-board nature of the pattern on a fortuitous coincidence of CC·Limiter constraints that happen to state conflicting demands. For example, among the CC·Limiters proposed by Bennett (2013), CC-SROLE requires that corresponding consonants have identical syllable roles (e.g., are both onsets, or both codas), whereas CC-EDGE(σ) requires that they be tautosyllabic. Since no pair of consonants could ever satisfy both of these requirements simultaneously, a ranking in which these two constraints along with CORR[+lat] are all undominated will happen to produce (and is, as far as we can tell, the only way of producing) an across-the-board /l…l/ → [l…r] dissimilation pattern. Needless to say, such a pattern is exceedingly simple to characterize in TSL_2 terms, with $\bar{S}=\{*ll\}$ and the tier T defined as either {l, r} or just {l} (depending on whether intervening [r] acts as a blocker or not).

4   Summary and Conclusions

We have explored and assessed the typologies generated by two theoretical approaches to the properties of locality in long-distance consonant phonotactics. The hypothesis that all long-distance phonotactic patterns of this kind are members of the Tier-based Strictly 2-Local class of stringsets withstands the scrutiny of comparison with the set of patterns observed in natural language. The TSL_2 approach offers an account of both unbounded and transvocalic patterns, with the latter being a special case of a more general category of long-distance dependencies with blocking. By contrast, the seemingly simple constraints used in Agreement by Correspondence to account for the locality dichotomy generate many pathological patterns that are unattested cross-linguistically. These include cases of harmony that are sensitive to the *count* (even vs. odd) of potential correspondents, cases of *agreement by proxy*, and cases of dissimilation that applies only *outside* of a relatively local $\underline{C}V.\underline{C}$ window. Moreover, from the perspective of formal language theory, we have shown that these patterns are relatively complex, residing outside of the TSL_2 region of the subregular hierarchy. We conclude that the TSL_2 approach provides a more favourable set of typological predictions about locality in long-distance phonotactics than does the factorial typology of ABC constraints. We note, however, that certain more general properties of long-distance phonology, such as the role of trigger-target similarity, have yet to be explored and characterized within a TSL_2 approach. Nevertheless, we suggest that membership in the class of TSL_2 formal languages is a useful criterion for assessing the predictions of any approach to locality in segmental phonology.

References

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

Bennett, William G. 2013. *Dissimilation, consonant harmony, and surface correspondence*. New Brunswick, NJ: Rutgers University dissertation.

Bennett, William G. 2015. Assimilation, dissimilation, and surface correspondence in Sundanese. *Natural Language and Linguistic Theory* 33(2). 371–415.

Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions on Information Theory* 2. 113–124.

Culy, Christopher. 1985. The complexity of the vocabulary of Bambara. *Linguistics and Philosophy* 8(3). 345–351.

Elmedlaoui, Mohamed. 1995. *Aspects de representations phonologiques dans certains langues chamito-semitiques*. Rabat: Faculté des Lettres et des Sciences Humaines. [Université Mohammed V dissertation, 1992.]

Finley, Sara & William Badecker. 2008. Analytic biases for vowel harmony languages. *West Coast Conference on*

Formal Linguistics (WCCFL) 27. 168–176.

Gafos, Adamantios I. 1999. *The articulatory basis of locality in phonology*. New York: Garland. [Johns Hopkins University dissertation, 1996.]

Hansson, Gunnar Ólafur. 2001. *Theoretical and typological issues in consonant harmony*. Berkeley, CA: University of California dissertation.

Hansson, Gunnar Ólafur. 2007. Blocking effects in agreement by correspondence. *Linguistic Inquiry* 38(2): 395–409.

Hansson, Gunnar Ólafur. 2010a. *Consonant harmony: Long-distance interaction in phonology*. Berkeley, CA: University of California Press. [Revised version of Hansson (2001).]

Hansson, Gunnar Ólafur. 2010b. Long-distance voicing assimilation in Berber: Spreading and/or agreement? *Annual Conference of the Canadian Linguistic Association (CLA)* 2010. Online: http://homes.chass.utoronto.ca/~cla-acl/actes2010/CLA2010_Hansson.pdf

Hansson, Gunnar Ólafur. 2014. (Dis)agreement by (non)correspondence: inspecting the foundations. Paper presented at ABC↔Conference, Berkeley. [Slides published in *2014 UC Berkeley Phonology Lab Annual Report: ABC↔Conference Archive*, 3–62. Online: http://linguistics.berkeley.edu/phonlab/annual_report/documents/2014/annual_report_2014_ABCC.html]

Hansson, Gunnar Ólafur & Kevin McMullin. 2014. Biased learning of long-distance assimilation and dissimilation. Poster presented at LAGB Workshop on Learning Biases in Natural and Artificial Language Acquisition, Oxford.

Hayward, Richard J. 1982. Notes on the Koyra language. *Afrika und Übersee* 65. 211–268.

Hayward, Richard J. 1990. Notes on the Aari language. In Richard J. Hayward (ed.), *Omotic language studies*, 425–493. London: School of Oriental and African Studies, University of London.

Heinz, Jeffrey. 2010. Learning long-distance phonotactics. *Linguistic Inquiry* 41(4). 623–661.

Heinz, Jeffrey & Regine Lai. 2013. Vowel harmony and subsequentiality. *Mathematics of Language (MoL)* 13. 52–63.

Heinz, Jeffrey, Cheten Rawal & Herbert G. Tanner. 2011. Tier-based strictly local constraints for phonology. *Association for Computational Linguistics (ACL)* 49(2). 58–64.

Jensen, John. 1974. Variables in phonology. *Language* 50(4): 675–686.

Johnson, C. Douglas. 1972. *Formal aspects of phonological description*. The Hague: Mouton.

Jurgec, Peter. 2011. *Feature spreading 2.0: A unified theory of assimilation*. Tromsø: Univ. of Tromsø dissertation.

Kaplan, Ronald M. & Martin Kay. 1994. Regular models of phonological rule systems. *Computational Linguistics* 20(3). 331–378.

Kobele, Gregory M. 2006. *Generating copies: An investigation into structural identity in language and grammar*. Los Angeles, CA: University of California dissertation.

Lai, Regine. 2012. *Domain specificity in learning phonology*. Newark, DE: University of Delaware dissertation.

Lombardi, Linda. 1999. Positional faithfulness and voicing assimilation in optimality theory. *Natural Language and Linguistic Theory* 17(2). 267–302.

McNaughton, Robert & Seymour Papert. 1971. *Counter-free automata*. Cambridge, MA: MIT Press.

Nevins, Andrew I. 2010. *Locality in vowel harmony*. Cambridge, MA: MIT Press.

Ní Chiosáin, Máire & Jaye Padgett. 2001. Markedness, segment realization, and locality in spreading. In Linda Lombardi (ed.), *Segmental phonology in optimality theory*, 118–156. Cambridge: Cambridge University Press.

Odden, David. 1994. Adjacency parameters in phonology. *Language* 70(2). 289–330.

Payne, Amanda. 2014. Dissimilation as a subsequential process. *North East Linguistic Society (NELS)* 44(1). 79–90.

Rabin, Michael O. & Dana Scott. 1959. Finite automata and their decision problems. *IBM Journal of Research and Development* 3(2). 114–125.

Rogers, James & Geoffrey K. Pullum. 2011. Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information* 20(3). 329–342.

Rose, Sharon & Rachel Walker. 2004. A typology of consonant agreement as correspondence. *Language* 80(3). 475–531.

Shieber, Stuart. 1985. Evidence against the context-freeness of natural language. *Linguistics and Philosophy* 8: 333–43.

Steriade, Donca. 1987a. Locality conditions and feature geometry. *North East Linguistic Society (NELS)* 17. 595–617.

Steriade, Donca. 1987b. Redundant values. *Chicago Linguistic Society (CLS)* 23(2). 339–362.

Walker, Rachel & Fidèle Mpiranya. 2005. On triggers and opacity in coronal harmony. *Berkeley Linguistic Society (BLS)* 31. 383–394.

Walker, Rachel, Dani Byrd & Fidèle Mpiranya. 2008. An articulatory view of Kinyarwanda coronal harmony. *Phonology* 25(3). 499–535.