

Lenition and morphology in Finnish

Kaden Holladay*

Abstract. This paper investigates an interaction between consonant lenition and morphology in Finnish. The language has a process of consonant lenition whereby underlying geminate consonants at syllable boundaries lenite (degeminate) when the addition of an affix makes the post-geminate rime bimoraic. A small class of possessor agreement affixes do not condition lenition, even if they create the appropriate phonological environment. A puzzling interaction emerges when possessor agreement affixes are stacked on top of certain lenition-conditioning affixes. I account for this interaction in a way that improves on Kiparsky’s (2003) analysis. In doing so, I extend Pater’s (2010) method for modeling exceptional phonology via lexically-indexed constraints.

Keywords. Finnish; lenition; morphology; phonology; lexically-indexed constraints

1. Background and the puzzle. Finnish has a process of consonant lenition (also termed *gradation* in the Uralic literature) whereby underlying geminate stops at syllable boundaries degeminate when the post-geminate rime becomes bimoraic (Antilla 1997, Kiparsky 2003). Various types of rime are bimoraic in Finnish, but in this paper I’ll focus only on VC rimes for simplicity. (See Kiparsky 2003 for an analysis of the interactions between stress, syneresis, and moraicity in the case of surface VV rimes.)

Examples (1a) and (2a) below¹ provide a baseline, where an underlying geminate surfaces faithfully. The (b) variants illustrate the lenition process, with the conditioning environment resulting from the addition of /-C/ or /-CCV/ affixes.

- (1) a. [hat.tu]²
 /hattu/
 hat
- b. [hɑ.tun]
 /hattu-n/
 hat-GEN.SG

*Many thanks to Andrew Lamont, Paul Kiparsky, Gaja Jarosz, Joe Pater, Max Nelson, members of the UMass Amherst Sound Workshop, and attendees of Morphology II at LSA 2019 for comments on this project. Thanks also to Saara Kaikkonen and Laura Maaranen for Finnish judgments. Authors: Kaden Holladay, University of Massachusetts at Amherst (kholladay@linguist.umass.edu).

¹All data in this paper, except when noted otherwise, come from Kiparsky (2003), Karlson’s (1983) grammar, or judgments elicited informally from two native speakers.

²I follow standard IPA conventions here, except that geminate consonants and long vowels are represented with doubling (e.g., “CC” instead of “C:”). Glossing abbreviations: ACC = accusative, CND = conditional, DIM = diminutive, GEN = genitive, LOC = locative, NOM = nominative, POSS = possessor agreement, PRV = privative, PST = past tense.

- (2) a. [ras.su**k**.ka]
 /rassu-**kk**A³/
 fat-DIM
- b. [ras.su.kal.la]
 /rassu-**kk**A-lla/
 fat-DIM-LOC

In (1a), the post-geminate rime is monomoraic, and no lenition is observed. In (1b) by contrast, the rime has become bimoraic by virtue of containing GEN.SG /-n/, so the /tt/ geminate in the root lenites. In (2a), the rime following the geminate stop is likewise monomoraic, and the word surfaces faithfully. But in (2b), post-geminate rime has become bimoraic by the addition of (the first consonant of) the locative /-lla/, so the /kk/ geminate of the diminutive suffix lenites. (Note that the /ss/ of the root doesn't lenite, as it isn't a stop.)

1.1 POSSESSOR AGREEMENT AFFIXES DO NOT CONDITION LENITION. Finnish has six possessor agreement affixes (POSS), given in (3) below (where V=vowel).

(3)

	SG	PL
1	-ni	-mme
2	-si	-nne
3	-nsA~-Vn	

When present, POSS affixes occur at the right edge of the phonological word (occurring at least more peripherally than both derivational affixes and case markers; I ignore clitics here). While POSS affixes can create the environment for lenition, they exceptionally do not condition it, as (4) and (5) below show.

- (4) [hat.tum.me]
 /hattu-mme/
 hat-POSS:1PL

- (5) [tak.kan.sa]
 /takka-nsA/
 hat-POSS:3

Example (4) contains an underlying geminate (/tt/) in the root. Despite the fact that the possessor agreement affix /-mme/ closes the post-geminate rime (making it bimoraic), no lenition is observed, standing in contrast to examples like (1b) and (2b). Example (5) illustrates the same point.

1.2 AFFIX STACKING AND LENITION: THE GENERAL CASE. What happens lenition-wise when a POSS affix is stacked on top of a more internal, lenition-conditioning affix? Usually, things go as expected given the data presented above.

³/A/ is a low vowel underspecified for backness; it receives its backness specification via vowel harmony.

- (6) [ha.tut.tam.me]⁴
 /hattu-ttA-mme/
 hat-PRV-POSS:1PL

In (6), lenition on the root geminate is conditioned by the internal affix (the privative), as expected given (1b) and (2b). The privative itself contains another geminate, which does *not* lenite because the following rime is closed by a possessor agreement affix. This too goes as expected, given examples (4) and (5).

1.3 THE PUZZLE: AFFIX STACKING WHEN THE INTERNAL AFFIX IS /-C/. A puzzling interaction emerges when a possessor agreement affix is stacked on top of a mono-consonantal (/C/) affix.

- (7) [hat.tum.me]
 /hattu-n-mme/
 hat-GEN.SG-POSS:1PL

- (8) [hat.tu.si]
 /hattu-t-si/
 hat-NOM/ACC.PL-POSS:2SG

There are two crucial descriptive facts that stand out about examples (7) and (8) above. Fact [1]: *No lenition occurs on the root geminate*. Note that the root /hattu/ surfaces with its geminate intact in both (7) and (8). Fact [2]: *The (otherwise mono-consonantal) internal affix doesn't surface*. (Compare (7) with (1b).)

Importantly, the generalization from §1.2 about POSS affixes not conditioning lenition is surface-transparent in (7-8). That is, on the surface, POSS is adjacent to the root, and the former does not condition lenition on the latter.

From a derivational perspective, however, this might look like an under-application of lenition. Why don't the affixes that usually instantiate GEN.SG and NOM/ACC.PL (/n/ and /t/) condition lenition on the root? We saw from example (6) that affixes situated between the root and POSS can indeed condition lenition on the root geminate, so long as they provide the appropriate environment by making a post-geminate rime bimoraic. The GEN.SG and NOM/ACC.PL case affixes in (7) and (8) occur in the same context (that is, they're situated between the root and POSS), and they also have the appropriate shape to condition lenition: /-C/ (cf. 1b).

Kiparsky (2003: p. 153) makes an important observation about the data in (7) and (8), namely that facts [1] and [2] seem to be related. That is, perhaps the absence of an exponent for the internal affixes bleeds lenition on the root geminate. For example, if there were no phonological exponent of the GEN.SG in (7), then there would be nothing to condition lenition on /hattu/, since /-mme/ doesn't condition lenition anyway.

In the next section I review Kiparsky's (2003) particular strategy for tying together facts

⁴Kiparsky (p.c.) points out that the combination of privative case and POSS sounds peculiar, though my consultants had accepted it in informal consultation. Phonologically similar examples can be found that illustrate the same point.

[1] and [2]. I offer two related arguments against his analysis, and in §3 I offer an alternative.

2. Kiparsky's (2003) Stratal OT analysis. We turn attention now to Kiparsky's solution for the puzzle from §1.3; his analysis implicates a morpho-phonological conspiracy (see Kisseberth 1970). A phonological markedness constraint (the Stem Constraint, introduced below) drives omission of the morphemes /-n/ and /-t/, which usually instantiate the GEN.SG and NOM/ACC.PL case categories in examples like (7) and (8). The same markedness constraint drives a substantively different syncope process in other contexts; a discussion of one of these contexts will come below.

The morpho-phonological component is assumed in Kiparsky (2003) to have a Stratal OT architecture (Stratal Optimality Theory; see Kiparsky 2000 for an overview). In Stratal OT, morphs are indexed to particular strata. At stratum n , morphs indexed to n are inserted, and a stratum-specific phonology – with a stratum-specific constraint ranking – applies. The output of stratum n then feeds stratum $n+1$, where morphs indexed to $n+1$ are inserted, and a potentially distinct phonology applies once again. This process applies cyclically through all the strata in a given language.

For Kiparsky, Finnish possessor agreement affixes are indeed to the word stratum (ω); all other affixes relevant here are indexed to the stem stratum (α). The output of the stem stratum feeds the word stratum. Consonant lenition is considered a stem-level phonological process, but not a word-level process. That is, at α , the markedness constraint driving lenition outranks the faithfulness constraint militating against lenition ($M \gg F$), but at ω this ranking is reversed ($F \gg M$).

I revisit example (6), repeated below as (9), to illustrate how this system works in practice.

- (9) [ha.tut.tam.me]
 /hattu-ttA-mme/
 hat-PRV-POSS:1PL

Recall that the internal affix conditions lenition on the root, but the possessor agreement affix does not condition lenition on the internal affix. For Kiparsky, the derivation of (9) can be represented as (10) below, where brackets delineate strata.

- (10) [hattu -ttA] _{α} → ha.tut.ta (output of the stem level)
 [[ha.tut.ta] _{α} -mme] _{ω} → ha.tut.tam.me

At the stem stratum, /hattu/ and /-ttA/ are inserted. The constraint ranking at α , with $M \gg F$, ensures that lenition occurs. The output feeds ω , where /-mme/ is inserted. At ω , no lenition occurs, because now the faithfulness constraint against lenition outranks the markedness constraint driving lenition.

This system works well to derive examples like (1b), (2b), (4-5), and (6/9). However, it runs into a problem with the puzzling examples (7) and (8). To see why, consider a first-pass derivation of (7), which I offer in (11) below.

- (11) [hattu -n] _{α} → ha.tun (output of the stem level)
 [[ha.tun] _{α} -mme] _{ω} → *ha.tunm.me (cf. observed [hat.tum.me])

The first-pass derivation I offer here fails to capture the observed surface form in two respects. First, it predicts lenition on the root geminate. If /hattu/ and /-n/ are inserted at α , with an $M \gg F$ ranking, root lenition is predicted, since the GEN.SG affix /-n/ would make the second syllable's rime bimoraic. However, no root lenition is observed (fact [1] from §1.3). Second, an exponent of the intermediate affix (here GEN.SG /-n/) is predicted to appear, while in fact no such exponent is observed (fact [2] from §1.3). To resolve the mismatch between predicted and observed forms in examples like (7), Kiparsky puts forward that the GEN.SG category is not realized as an affix at all – rather, it is omitted morphologically (p. 151). Morphological omission is driven by the Stem Constraint ($*-C]_{\alpha}$), a markedness constraint which penalizes stems not ending in a vowel.

- (12) $*-C]_{\alpha}$ = stems must end in -V
(Kiparsky 2003: 153)

This constraint is assumed to be undominated in Finnish. The competing (but lower-ranked) faithfulness constraint relevant here is what Kiparsky calls DEP-CAT(EGORY)⁵.

- (13) DEP-CAT = A grammatical category must be morphologically expressed.
(Kiparsky 2003: 154)

Violations of this constraint are incurred when a grammatical category (in this instance, case) is omitted from the derivation: that is, not granted a morphological representation to begin with. By the ranking $*-C]_{\alpha} \gg \text{DEP-CAT}$, categories otherwise expressed as an affix with a /-C/ shape will have no expression the Stratal morpho-phonological derivation at all. Since $*-C]_{\alpha}$ outranks DEP-CAT, including an underlying /-n/ in the derivation of (7) is non-optimizing to begin with, as it incurs a violation of the Stem Constraint. Thus the derivation of (7) is, for Kiparsky, represented as (14).

- (14) $[\text{hattu } \emptyset]_{\alpha} \rightarrow \text{hat.tu}$
 $[[\text{hat.tu}]_{\alpha} \text{-mme}]_{\omega} \rightarrow \text{hat.tum.me}$

Since there are no morphemes to instantiate the GEN.SG, NOM.PL, and ACC.PL categories (/n/ and /-t/ are, in his words, “simply not present”, p. 151), the absence of lenition on the root geminate and the non-exponence of GEN.SG and NOM/ACC.PL categories are simultaneously derived. Kiparsky's intuition that non-exponence is related to the lack of root lenition is crucial, and is something we'll make use of in the proposed re-analysis in §3. However, the exact strategy by which he relates these facts is less than ideal, as I spell out below.

2.1 PROBLEMS WITH MORPHOLOGICAL OMISSION. The Stem Constraint is considered by Kiparsky to be undominated, and its status as such means that Finnish stems will always be vowel-final. One way the grammar enforces this is through morphological omission under $*-C]_{\alpha} \gg \text{DEP-CAT}$, as detailed above. But there are other process that enforce conformity to the Stem Constraint as well.

One such process is as follows. When a multi-segmental, consonant-final affix is situated

⁵A MAX constraint seems to me more appropriate than a DEP constraint here, but I'll retain Kiparsky's original name.

left-adjacent to a POSS affix in Finnish, it surfaces without its final consonant (Kiparsky 2003: p. 149). In contrast to (14) above, the entire morpheme is not omitted. Some examples are offered below.

- (15) a. [talo**ihin**]
 /talo-i-hin/
 house-PL-LOC
 b. [talo**ihimme**]
 /talo-i-hin-mme/
 house-PL-LOC-POSS:1PL
- (16) a. [vapaa**siin**]
 /vapaa-i-siin/
 free-PL-LOC
 b. [vapaa**siisi**]
 /vapaa-i-siin-si/
 free-PL-LOC-POSS:2SG

(14) and (15) feature suppletive allomorphs of a locative case (the illative): /-hin/ and /-siin/. In the (a) examples, when no possessor agreement is present, the morphs surface with their underlying final consonants. But in the (b) examples, when left-adjacent to possessor agreement, the final consonant does not surface. Kiparsky treats this as a case of final consonant deletion: consonant-final affixes like /-hin/ and /-siin/ “lose their last consonant before possessive suffixes” (p. 149). Though the faithfulness constraint he has in mind is not stated explicitly, the relevant ranking seems to be $*-C]_{\alpha} \gg \text{MAX}(C)$.

Recall, however, that the non-exponence of the mono-consonantal case affixes (/n/ and /t/) was derived through a different ranking: $*-C]_{\alpha} \gg \text{DEP-CAT}$. This treatment divorces the behavior of *multi*-segmental, consonant-final affixes from the behavior of *mono*-segmental, consonant-final affixes. The former are subject to phonological syncope, while the latter are subject to morphological omission. For Kiparsky, independent processes are acting on these two types of affix, resulting in a morpho-phonological conspiracy driven by the Stem Constraint. There’s nothing wrong with analyzing a set of facts as a conspiracy per se; indeed, one of the great advantages of Optimality Theory over rule-based approaches is that the former *can* model conspiracies in a satisfying way. Still, a more parsimonious account of the Finnish facts would be one where the behavior both types of affix (mono-consonantal and multi-segmental) was derived from a single $M \gg F$ ranking.

There’s another, related problem for Kiparsky’s analysis as well. Consider the following question: *why* are mono-consonantal affixes subject to morphological omission, while multi-segmental, consonant-final affixes are subject to phonological syncope? In principle, phonological syncope could satisfy the Stem Constraint for either type of affix, as could morphological omission. To resolve this question, Kiparsky recruits another constraint, MAX-MORPHEME.

- (17) MAX-MORPHEME = A morpheme must be phonologically expressed.
 (Kiparsky 2003: 153)

Crucially, violations of this constraint are incurred only when a morpheme already present in

the morpho-phonological derivation surfaces with no exponent. That is, violations are *not* incurred if a morpheme is omitted entirely.

The following tableaux contain Kiparsky's ranking of these constraints⁶ and illustrates how they interact with some of the examples I introduced above. Note that candidates are now input-output pairs of morphemes and surface forms. MAX-MORPHEME is abbreviated at MAX-M, and each morpheme is marked with a subscript indicating the stratum it's indexed to. Tableau (18) illustrates how morphological omission is derived.

(18)

	*-C] _α	MAX-M	DEP-CAT	MAX(C)
a. /hattu _α -n _α -mme _ω /~[ha.tunm.me]	*			
b. /hattu _α -n _α -mme _ω /~[ha.tum.me]		*		*
☞ c. /hattu _α -∅ _α -mme _ω /~[hat.tum.me]			*	

Here candidate (a) expresses the category GEN.SG with the morpheme /-n/. This /-n/ is not subject to phonological syncope, so no violation of MAX(C) is incurred. The output of the stem level for this candidate is [ha.tun], which incurs a violation of *-C]_α, the Stem Constraint. Candidate (b) likewise expresses the category GEN.SG with the morpheme /-n/, but now this morpheme is subject to syncope, incurring a critical violation of MAX-MORPHEME (and incurring a violation of MAX(C) as well). Candidate (c) opts for morphological omission, and does not express GEN.SG at all. This incurs a violation of DEP-CAT, but satisfies the higher-ranked constraints *-C]_α and MAX-MORPHEME.

Now consider (19), whose winning candidate features the other process implicated in the conspiracy – phonological syncope.

(19)

	*-C] _α	MAX-M	DEP-CAT	MAX(C)
a. /vapaa _α -i _α -siin _α -si _ω /~[va.pai.siin.si]	*			
☞ b. /vapaa _α -i _α -siin _α -si _ω /~[va.pai.sii.si]				*
c. /vapaa _α -i _α -∅ _α -si _ω /~[va.pai.si]			*	

Candidate (a) expresses the category LOC with the morpheme /-siin/. The final consonant of this candidate is not subject to phonological syncope, so no violation of MAX(C) is incurred. The output of the stem level for this candidate is [va.pai.siin], which incurs a critical violation of *-C]_α, the Stem Constraint. Candidate (b) also expresses the category LOC with the morpheme /-siin/, but with the final consonant being subject to syncope. This incurs a violation of MAX(C), but crucially not MAX-MORPHEME, as the morpheme is still phonologically expressed. Candidate (c) opts for morphological omission and does not express LOC at all; this incurs a violation of DEP-CAT.

While this analysis seems to work at first glance, there's a lookahead problem lurking in the shadows. Recall that the output of stem stratum feeds the word stratum. What this means is that stem-level constraint evaluation should have no information available about whether any word-level affixes will be added later in the derivation. But consider the contrast between example (1b) and the derivation in (14/18c). On Kiparsky's treatment, the GEN.SG category in (1b) is expressed as the morpheme /-n/, while in (14/18c), the morpheme corresponding to the

⁶Kiparsky does not include MAX(C) in the relevant tableaux, but I've added it here for completeness.

GEN.SG category is omitted. Omission, or lack thereof, is sensitive to whether there *will* be a POSS morpheme added at a derivationally later stratum.

The central issue with the constraints utilized in tableaux (18-19) is that violations of $*-C]_{\alpha}$, MAX-MORPHEME, DEP-CAT, and MAX(C) are being evaluated across strata. To my understanding, this contradicts the cyclicity assumption of Stratal OT, and instead gestures toward a new analysis of these facts. My goal in the next section is to offer such an analysis.

3. Re-analysis: exceptional phonology and lexically-indexed markedness. In this section I offer a re-analysis of the Finnish lenition facts. This analysis unifies the treatment of multi-segmental, consonant-final affixes with the treatment of mono-consonantal affixes, in that I derive the fact that POSS affixes always surface right-adjacent to vowels with a single $M \gg F$ ranking. (Recall that Kiparsky had two: $*-C]_{\alpha} \gg \text{MAX(C)}$ and $*-C]_{\alpha} \gg \text{DEP-CAT}$.) My analysis obviates the need for morphological omission and avoids the lookahead problem of Kiparsky's system. Instead, I employ and expand on Pater's (2010) means of capturing exceptional phonological behavior with lexically-indexed constraints.

3.1 CONSONANT DELETION. The first task is to unify the treatments of multi-segmental, consonant-final affixes (like /-hin/ and /-siin/ from (15) and (16)) with the treatment of mono-consonantal affixes (like /-n/ and /-t/ from (7) and (8)). Since both types of affix only lose a consonant when left-adjacent to POSS affixes, I propose that we should consider the POSS affixes *themselves* as the culprits responsible for deletion (rather than the Stem Constraint). I'll model consonant deletion with phonological alignment constraint (ALIGN; McCarthy & Prince 1993, see Hyde 2012), defined in (20).

(20) ALIGN(L, POSS, R, V) = one violation for every left edge of a POSS affix that is not aligned to the right edge of a vowel

If the ALIGN constraint outranks a faithfulness constraint militating against consonant deletion, the final consonants of affixes left-adjacent to POSS affixes will be subject to syncope. The relevant faithfulness constraint is defined in (21), and tableau (22) illustrates the ranking required to derive (16b).

(21) MAX(C) = one violation for every input consonant not present in the output

(22)

/vapaa-i-siin-si/	ALIGN	MAX(C)
 a. vapaisiisi		*
b. vapaisiinsi	*W	L

In candidate (b), the /n/ input segment is present in the output. This candidate incurs a violation of ALIGN, and does not incur violation of MAX(C). In candidate (a), a violation of MAX(C) is incurred, but ALIGN is satisfied. Since ALIGN outranks MAX(C), candidate (a) is optimal. This particular ranking (ALIGN \gg MAX(C)) will be important later, but for the moment, it will be put aside.

3.2 EXCEPTIONAL NON-TRIGGERS. Next, we need a way of capturing the lenition facts. Specifically, we need a way of modeling the contrast between, on the one hand, POSS affixes, which exceptionally do not condition (or *trigger*) lenition, and on the other, the rest of the af-

fixes of the language, which do trigger lenition. I make use of Pater’s (2010) model, which captures exceptional phonological behavior by way of lexically indexing constraints.

As an illustration of the type of phenomenon lexical indexing is supposed to model, consider Finnish /a-i/ from Pater (2010). When /a/ and /i/ are brought together at morpheme boundaries, they can surface as either [ai] or [oi], the choice of which depends arbitrarily on the morpheme containing /i/. Some morphs containing /i/, such as the conditional /-isi/, allow /a-i/ to surface faithfully as [ai]. When other other morphs containing /i/ are involved, such as the past tense /-i/, [oi] surfaces:

- (23) [ant**ai**si]
/anta-isi/
give-CND
- (24) [ant**oi**]
/anta-i/
give-PST

The contrast between examples (23) and (24) can be captured if there are actually two markedness constraints involved in their derivations. The first is a general constraint that incurs violations for any and all [ai] strings. The second, *[ai]_L, incurs violations only when a member of a lexically-indexed (L) class of morphs is part of the [ai] string.

- (25) *[ai] = one violation for every output [ai] string
- (26) *[ai]_L = one violation for every output [ai] string that contains⁷ some morpheme lexically indexed to L, the exceptional class of morphs

Assume past tense /-i/ belongs to the lexically indexed class. With a faithfulness constraint in the mix (in this case IDENT([LOW])), the attested pattern can be derived on the ranking given in tableaux (27-28) below.

(27)

/anta-isi/	*[ai] _L	IDENT([LOW])	*[ai]
☞ a. antaisi			*
b. antoisi		*	

(28)

/anta-i _L /	*[ai] _L	IDENT([LOW])	*[ai]
☞ a. antoi		*	
b. antai	*		*

In tableau (27), no morpheme indexed to L is present, so violations of *[ai]_L cannot be incurred. The IDENT constraint outranks the general *[ai] constraint, so a change in vowel height, as in candidate (b), is non-optimizing. In (28) by contrast, a morpheme lexically indexed to L is present, namely /-i/. Candidate (b) here incurs violations of both markedness constraints (but crucially *[ai]_L); candidate (a) wins.

Pater’s system can model cases where an exceptional class of morphs are *triggers* of a

⁷A string will contain a morpheme if and only if some exponent of the morpheme occurs within the string. See Jarosz (2018) for a discussion of related locality restrictions.

process (or non-undergoers, if faithfulness constraints are lexically indexed). To model the lenition facts from §1, additional machinery will be needed; the reasons for this deserve some comment.

A process in Optimality Theory can be modeled with a markedness-over-faithfulness ranking. When a select class of morphs are exceptional triggers of a process, a markedness constraint indexed to that class can be posited above the relevant faithfulness constraint; the process will apply when the relevant structural configuration is met so long as it contains a morpheme indexed to the exceptional class. If the general (i.e., non-indexed) markedness constraint is ranked below the faithfulness constraint, the process cannot apply to strings not containing an exceptional morph. In this characteristic $M_L \gg F \gg M$ ranking, the indexed markedness constraint, crucially, outranks the non-indexed markedness constraint.

Now recall the lenition pattern from §1: all morphs *except* the possessor agreement affixes can trigger lenition. Here, an exceptional class of morphs are *non-triggers* of a process. No ranking containing a lexically-indexed markedness constraint, a non-indexed (general) markedness constraint, and a faithfulness constraint will derive this pattern; the specific/indexed vs. general/non-indexed dichotomy is too narrow. For illustration, assume that the exceptional (lexically-indexed) class contains all and only POSS affixes. If the ranking were $M_L \gg F \gg M$, lenition would only occur under conditioning by POSS affixes, the exact opposite of the attested pattern. If the ranking were $M \gg F \gg M_L$, lenition would apply everywhere, again contra fact.

This problem stems from the fact that Pater’s (2010) lexically-indexed markedness constraint definition features only existential quantification over morphs (as I put it, “one violation for every output [ai] string that contains *some* morpheme lexically indexed to L. . .”⁸). That is, the original schematic for lexically-indexed constraints is something like the following (cf. (26) above):

- (29) $*[S]_{\exists, n}$ = one violation for every output string [S] such that some (\exists) morpheme contained in [S] is indexed to class n

Here S stands for the locus of violation – some marked string. Subscript \exists indicates that this markedness constraint is sensitive to whether *some* morpheme of a particular class falls within the string, while subscript n tracks class membership (I use n instead of L here because in what follows I’ll be indexing morphemes to multiple classes.)

Intuitively, what’s needed to model Finnish lenition is for the process to apply just in case *all* morphs in the marked string are not POSS morphs. That is, the constraint definitions need to be able to reference universal quantification – whether all morphs contained in a string are of a certain variety. If they are, lenition should apply; if they aren’t, it shouldn’t. Accordingly, I propose that the following type of constraint definition should also be possible.

- (30) $*[S]_{\forall, n}$ = one violation for every output string [S] such that all (\forall) morphs contained in [S] are indexed to class n

I’ll make use of both existential and universal varieties in the analysis below (though the exis-

⁸Pater’s original definition is $*X_L$ = Assign a violation mark to any instance of X that contains a phonological exponent of a morpheme specified as L.

tential variety of the constraint could be substituted for a general/non-indexed constraint without complications). Before going there, we need to spell out exactly what the n and S variables in the constraint definition schematic above should correspond to for the task at hand.

n is a variable over classes of morphs. I assume the following classes of Finnish morphs: class 1 contains all morphs except the possessor agreement morphs, while class 2 contains all and only possessor agreement morphs.⁹

- (31) Class 1 = { x | x is not in Class 2}
 Class 2 = { $-ni$, $-si$, $-mme$, $-nne$, nsA , $-Vn$ }

S is a variable over types of phonological strings. Following Antilla (1997), the type of phonological string that drives lenition is assumed here to be $[\sigma_{\mu\mu}\sigma_{\mu\mu}]$ – i.e., a string consisting of adjacent, bimoraic syllables.¹⁰ This is a “quantitative dissimilation constraint” (Kiparsky 2003: 131) active in the Finnish metrical domain. Recall that geminate lenition is driven by the weight of the post-geminate rime, occurring when the post-geminate rime is bimoraic. The repair is to degeminate, shedding a mora from the first syllable of the string.

What’s relevant for the current analysis are two combinations of constraint variety and morpheme class:

- (32) $*[\sigma_{\mu\mu}\sigma_{\mu\mu}]_{\exists,2}$ = one violation for every output string $[\sigma_{\mu\mu}\sigma_{\mu\mu}]$ such that **some morph** contained in $[\sigma_{\mu\mu}\sigma_{\mu\mu}]$ is indexed to **class 2**
 (33) $*[\sigma_{\mu\mu}\sigma_{\mu\mu}]_{\forall,1}$ = one violation for every output string $[\sigma_{\mu\mu}\sigma_{\mu\mu}]$ such that **all morphs** contained in $[\sigma_{\mu\mu}\sigma_{\mu\mu}]$ is indexed to **class 1**

The faithfulness constraint against degemination is assumed to be $\text{MAX}(\mu)$.¹¹

- (34) $\text{MAX}(\mu)$ = one violation for every input mora with no output correspondent

With constraints (32-34), as well as ALIGN and $\text{MAX}(C)$ as defined in (20-21), it’s possible capture all the Finnish lenition facts without reference to morphological omission. Consider first (1b), repeated here as (35).

- (35) [ha.tun]
 /hattu-n/
 hat-GEN.SG

(35) represents a case where lenition is conditioned by the presence of $/-n/$, which results in a string of adjacent bimoraic syllables. Since neither /hattu/ nor $/-n/$ are POSS morphs, both are indexed to class 1, as indicated by subscript in the input of tableau (36) below.¹²

⁹The lexically-indexed markedness constraints I use here are indexed to either a natural class (class 2: POSS affixes) or the complement class of that natural class. Indexing to a complement class is a complication to the lexical indexing system, but seems to be required to get the Finnish facts right on a lexical indexing approach.

¹⁰Anttila calls his constraint $*HH$, where “H” means “heavy syllable”.

¹¹In fact, it needs to be more specific than $\text{MAX}(\mu)$, as only geminate stops lenite – geminate sonorants don’t. I ignore this complication here.

¹²My analysis is couched in Parallel OT (Prince & Smolensky 1993/2004), mainly for expository reasons. In all examples presented in this paper, there are only ever two geminate stops adjacent on the consonant tier. The lenition facts can captured equally well in Parallel OT or Harmonic Serialism (ibid.) so long as this is the case. But observe

(36)

/hattu ₁ -n ₁ /	MAX(C)	*[σ _{μμ} σ _{μμ}] _{∀,1}	MAX(μ)	*[σ _{μμ} σ _{μμ}] _{∃,2}
☞ a. ha.tun			*	
b. hat.tun		*		
c. hat.tu	*		*	

Note that because all morphs in the input are in class 1, no violations of *[σ_{μμ}σ_{μμ}]_{∃,2} can be incurred throughout the tableau. Let's examine the candidates one by one. Candidate (36b) is fully faithful – it incurs a violation of *[σ_{μμ}σ_{μμ}]_{∀,1}, since it hosts a string of adjacent bi-moraic syllables containing only class 1 morphs. Candidate (36c) ameliorates the violation of *[σ_{μμ}σ_{μμ}]_{∀,1}, but in the wrong way: it deletes /n/, incurring a violation of both MAX(μ) and MAX(C). The optimal candidate, (36a), is one that avoids a violation of *[σ_{μμ}σ_{μμ}]_{∀,1} by way of degemination, which incurs a violation of MAX(μ) but not MAX(C).

Next, consider (4), repeated here as (37). Here, no lenition is observed (recall that /-mme/ is a possessor agreement morpheme and exceptionally does not condition lenition).

- (37) [hat.tum.me]
/hattu-mme/
hat-POSS:1PL

Though the root /hattu/ is in class 1, the POSS morpheme /-mme/ is in class 2. This significantly affects the way that markedness violations are assigned, as (38) shows.

(38)

/hattu ₁ -mme ₂ /	MAX(C)	*[σ _{μμ} σ _{μμ}] _{∀,1}	MAX(μ)	*[σ _{μμ} σ _{μμ}] _{∃,2}
☞ a. hat.tum.me				*
b. ha.tum.me			*	
c. hat.tu.me ¹³			*	

Because the marked [σ_{μμ}σ_{μμ}] string – consisting of the first two syllables – now contains morphs of both classes, violations of the universal variety of the constraint cannot be incurred, though violations of the existential variety can. Because the existential variety of the constraint is out-

what happens when there are *more than two* geminate stops in the input (ex. from Jensen 1973):

- (i) /rokko-tt-UttA-vA/ → [ro.ko.tut.ta.va]

(37) undergoes two instances of lenition, once on the root and once on the most internal affix. On the consonant tier it features /...GGG.../ → [...NNG...], where G = geminate and N = non-geminate. If all morphs are present before phonological evaluation, Parallel OT and vanilla Harmonic Serialism only predict lenition on the medial geminate: /...GGG.../ → [...GNG...]. This is because the /...GGG.../ → [...GNG...] mapping gets rid of both [σ_{μμ}σ_{μμ}] loci of violation with minimal faithfulness violation. Andrew Lamont is developing a directional variety of Harmonic Serialism (Lamont 2018, forthcoming) which can evaluate markedness violations from left-to-right. Such directional evaluation can capture this Finnish pattern in conjunction with Serial Markedness Reduction constraints (Jarosz 2014). The pattern in (37) could be captured a different way if phonological optimization is interleaved with morpheme insertion (add an affix, lenite, add another affix, lenite, etc.). This sort of architecture is utilized in, e.g., Wolf (2008) and Baković (2002). Such an approach would require reference to both morphology and phonology, so is a bit less parsimonious than Directional HS.

¹³Candidate (38c) is in fact harmonically bounded by (38b), since moraic obstruents are more marked than moraic sonorants. For cases like /hattu₁-tta₁/ → [ha.tut.ta], *[hat.tu.ta], we need to have a mechanism for directionality. It's best to use Directional HS – see Lamont (forthcoming).

ranked by $\text{MAX}(\mu)$, however, lenition is non-optimizing. This rules out candidates (38c) and (38d); all said and done, (38a) wins.

With these cases out of the way, let's return to the puzzle from §1.3, which was illustrated in part with (7), repeated here as (39).

- (39) **hat.tum.me**
 hattu-n-mme
 hat-GEN.SG-POSS:1PL

The two facts that require explanation are the non-exponence of GEN.SG /-n/ and the lack of lenition on the root. Recall that in tableaux (22), I posited an $\text{ALIGN} \gg \text{MAX}(\text{C})$ ranking to explain why multi-segmental, consonant-final affixes *and* mono-consonantal affixes surface without their final consonants when left-adjacent to POSS. I invoke that ranking again in tableau (40) below.

(40)

/hattu ₁ -n ₁ -mme ₂ /	ALIGN	MAX(C)	*[σ _{μμ} σ _{μμ}] _{∀,1}	MAX(μ)	*[σ _{μμ} σ _{μμ}] _{∃,2}
☞ a. hat.tum.me		*		*	*
b. ha.tum.me		*		**W	L
c. hat.tunm.me	*W	L		L	*
d. ha.tunm.me	*W	L		*	L

Candidates (40c) and (40d) here contain [n], supplied by GEN.SG. This consonant is situated between the POSS affix and a vowel, so a violation of ALIGN is incurred for both candidates, which are accordingly ruled out regardless of whether lenition occurs on the underlying root geminate. Candidate (b) satisfies the ALIGN constraint, but also lenites the root geminate. This incurs a pointless violation of $\text{MAX}(\mu)$ (in addition to the $\text{MAX}(\mu)$ violation incurred by the deletion of /n/). Candidate (40a) is optimal. This candidate satisfies the ALIGN constraint (no [n] sits between POSS and a vowel), and does not feature lenition, a process which is non-optimizing here, since class 1 and class 2 morphs are both present in the marked [σ_{μμ}σ_{μμ}] string, just like in tableau (40). In this way, the interaction between the undominated ALIGN CONSTRAINT and the lexically-indexed markedness constraints neatly captures the intuition that the surface absence of the intermediate affix is bleeding consonant lenition on the root.

The constraints and ranking in the tableaux above thus derive the behavior of Finnish geminate stops with regard to lenition. Crucially, they also derive the non-exponence of the internal affixes in examples like (7/35) and (8), and do so in a way that dovetails with the treatment of multi-segmental, consonantal-final affixes: both are subject to phonological syncope.

4. Concluding remarks. In this paper I offered a reanalysis of a morpho-phonological puzzle in Finnish noted by Kiparsky (2003). My analysis shifted the locus of explanation for the constellation of affixation and lenition facts entirely into the phonology. It can derive the attested and non-attested patterns of lenition, and reduces Kiparsky's (2003) cases of morpheme omission to syncope – a welcome move, given that syncope occurs in Finnish in identical environments. This syncope process was modeled with a single ranking: an ALIGN constraint (relativized to POSS affixes) over MAX(C).

Modeling the lenition facts with lexically indexed constraints required two extensions to

Pater's (2010) system: (i) indexed markedness constraint definitions referenced universal as well as existential quantification over morphs, and (ii) constraint indexation to the complement class of a natural class was allowed. It's important, of course, that these extensions to the system be validated by empirical coverage of related phenomena.

References

- Anttila, Arto. 1997. *Variation in Finnish phonology and morphology*. Stanford, CA: Stanford University dissertation.
- Bakovic, Eric. 2002. Vowel harmony and cyclicity in Eastern Nilotic. *Berkeley Linguistics Society (BLS)* 27. 1-12. <https://doi.org/10.3765/bls.v27i1.1116>.
- Hyde, Brett. 2012. Alignment constraints. *Natural Language and Linguistic Theory* 30. 789-836. <https://doi.org/10.1007/s11049-012-9167-3>.
- Jarosz, Gaja. 2014. Serial markedness reduction. *Proceedings of 2013 Annual Meetings on Phonology*. <http://dx.doi.org/10.3765/amp.v1i1.40>.
- Karlsón, Fred. 1983. *Finnish grammar*. Helsinki: WSOY.
- Kisseberth, Charles. 1970. On the functional unity of phonological rules. *Linguistic Inquiry* 1. 291-306. <https://www.jstor.org/stable/4177568>.
- Kiparsky, Paul, 2000. Stratal OT: A synopsis and FAQs. In Yuchau E. Hsiao and Lian-Hee Wee (eds.), *Capturing phonological shades*. Cambridge Scholars Publishing,
- Kiparsky, Paul. 2003. Finnish noun inflection. In Diane Nelson and Satu Manninen (eds.), *Generative approaches to Finnic linguistics*. Stanford: CSLI.
- Lamont, Andrew. To appear. Majority rule in Harmonic Serialism. *Proceedings of 2018 Annual Meetings on Phonology*.
- Lamont, Andrew. Forthcoming. Directional Harmonic Serialism. MS.
- McCarthy, John J. & Prince, Alan. 1993. Generalized Alignment. In Geert Booij and Jaap van Marle (eds.), *Yearbook of Morphology*. 79-153. Dordrecht: Kluwer.
- Pater, Joe. 2010. Morpheme-specific phonology: Constraint indexation and inconsistency resolution. In Steve Parker (ed.), *Phonological argumentation: Essays on evidence and motivation*. 123-154. London: Equinox.
- Wolf, Matthew. 2008. *Optimal interleaving: Serial phonology-morphology interaction in a constraint-based model*. Amherst, MA: University of Massachusetts Amherst dissertation.