

STRATIFIED FAITHFULNESS IN HARMONIC GRAMMAR AND EMERGENT CORE-PERIPHERY STRUCTURE*

JENNIFER L. SMITH

University of North Carolina at Chapel Hill

Foundational work by Ito & Mester (1995ab, 1999, 2001, 2008) connects core-periphery phonology, in which different lexical strata in a language tolerate different degrees of ‘foreign’ phonological structure, to a ranked hierarchy of markedness constraints against which faithfulness constraints are ranked differently for different strata. Implementing a version of the Stratified Faithfulness model in Harmonic Grammar takes advantage of cumulative interaction between specific and general faithfulness constraints (Jesney & Tessier 2011) to solve two remaining problems: how to make core-periphery structure a soft bias rather than an absolute requirement, and how to formalize the consistency of faithfulness rankings across strata that is a necessary condition for productive core-periphery phonology.

Keywords: loanword phonology, stratified lexicon, Harmonic Grammar, cumulative constraint interaction, indexed constraints

1 Introduction

Early in the development of Optimality Theory (OT; Prince & Smolensky 1993/2004; McCarthy & Prince 1995), work pioneered by Ito & Mester (1995ab, 1999) established a key conceptual insight: a language with a stratified lexicon that has a phonologically productive CORE-PERIPHERY STRUCTURE, where successive lexical strata tolerate increasingly ‘foreign’ phonological properties, can be modeled in terms of a HIERARCHY OF MARKEDNESS CONSTRAINTS (surface well-formedness constraints). Constraints against properties that are ‘less foreign,’ or more CORE, are ranked lower, and so the relevant structures are more easily tolerated in loanword phonology. Constraints against properties that are ‘more foreign,’ or more PERIPHERAL, are ranked higher, and so the structures that violate these constraints are more aggressively nativized.

While the notion of a markedness-constraint hierarchy as the backbone of a core-periphery phonology is intuitively appealing, however, the formal implementation of this insight in a constraint-based grammar has proven not to be entirely straightforward. In particular, Fukazawa, Kitahara, & Ota (1998) argue that a language with a stratified phonology is best modeled with distinct sets of faithfulness constraints indexed to each stratum, so as to allow for different faithfulness versus markedness rankings in different strata simultaneously—rather than, for example, with a cogrammars model in which a single set of constraints is literally reranked for different strata (as in Ito & Mester 1995a; see also, e.g., Inkelas & Zoll 2007). However, Fukazawa et al. (1998) and Ito & Mester (1999) go on to demonstrate that this OT STRATIFIED FAITHFULNESS model requires additional ranking stipulations if it is to enforce a strict core-periphery structure. Given a markedness hierarchy $M1 \gg M2 \gg M3$, there is a logically possible faithfulness ranking (see §2) that leads to satisfaction of *low*-ranking $M3$ but violation of *high*-ranking $M1$, producing surface forms that fall outside the intended core-periphery patterning—that is, IMPOSSIBLE NATIVIZATIONS (Ito & Mester 1999, 2001) that preserve a ‘more-foreign’ property while nativizing a ‘less-

*Thanks to Elliott Moreton, Brian Hsu, and the Phonetics/Phonology Research Group at UNC-CH for comments and suggestions. Many thanks also to Junko Ito and Armin Mester for their always illuminating and generously encouraging discussions on loanword phonology, Japanese linguistics, and a range of other topics in phonological theory.

foreign' one. In other words, without additional metaconditions on possible rankings, the OT Stratified Faithfulness model actually allows for non-core-periphery patterns, even in the presence of a strict markedness hierarchy.

Hsu & Jesney (2017a) develop an alternative approach to core-periphery phonology in the framework of Harmonic Grammar (HG; Legendre, Miyata, & Smolensky 1990; Pater 2009). Their WEIGHTED SCALAR CONSTRAINTS model likewise incorporates a markedness-constraint hierarchy to establish which of the phonologically marked properties are 'more foreign' and 'less foreign.' But this model proposes only a single set of faithfulness constraints, plus a scaling factor that boosts the weight of each faithfulness constraint in proportion to the 'distance' of the form it is evaluating from the lexical core. Such an approach allows faithfulness to take increasingly greater priority over markedness requirements as forms become more peripheral, without positing multiple sets of stratum-specific faithfulness constraints. Crucially, in Hsu & Jesney's (2017a) version of the Weighted Scalar Constraints model,¹ the relative weighting relations among faithfulness constraints can never change across strata. This guarantees that no surface form can ever satisfy a lower-weighted markedness constraint while violating a higher-weighted one. Impossible nativizations are excluded, and core-periphery phonology is strictly enforced, without any need for extrinsic stipulations on the relative weighting of faithfulness constraints.

In summary, the OT Stratified Faithfulness model *cannot enforce* a strict core-periphery structure without additional metaconditions on faithfulness rankings across strata, while the Weighted Scalar Constraints model predicts that every stratified lexicon *necessarily* has a strict core-periphery structure. But empirical evidence suggests that what is really needed is an intermediate position. On the one hand, at least some speakers of a number of languages, including Japanese, judge productively that impossible nativizations—which fall outside a strict core-periphery structure—are dispreferred (see Ito & Mester (1999, 2001) on loanwords in Japanese and German and the continuum of registers in Jamaican Creole English; Pinta (2013), Smith & Pinta (2017) on loanwords in Paraguayan Guarani). On the other hand, Fukazawa et al. (1998) summarize a number of markedness implicational relations in the Japanese lexicon which, taken together, show that the lexical strata in Japanese do not in fact form a strict core-periphery structure across all dimensions of markedness; see also Ito & Mester (1995b), who note that the Mimetic stratum and the Sino-Japanese stratum do not form a subset/superset relation, and Kawahara, Nishimura, & Ono (2003), who argue that the Sino-Japanese stratum is even less marked than the Native stratum in certain respects.

This paper implements a version of the Stratified Faithfulness model in Harmonic Grammar that builds on the insights of previous approaches, but has two advantages. First, the HG Stratified Faithfulness model is able to account for both kinds of stratified phonologies: those that do, and those that do not, have core-periphery structure. Second, the formal properties of HG make it simple for this model to encode an EMERGENT PREFERENCE for core-periphery structure, which accounts for the existence of productive impossible-nativization effects without treating core-periphery structure as a universal requirement on stratified phonologies. The formal implementation of this emergent preference takes advantage of CUMULATIVE CONSTRAINT INTERACTION between specific and general faithfulness constraints in HG, in the tradition of work by Jesney & Tessier (2011) on other types of specific/general faithfulness interactions.

First, §2 presents background on core-periphery phonologies, the OT Stratified Faithfulness model, and the role of consistent faithfulness rankings across strata in enforcing strict core-periphery structure. The properties of the HG Stratified Faithfulness model are presented in §3, and its predictions are explored in


¹Hsu & Jesney (2017b) introduce a revised version of the Weighted Scalar Constraints model that allows for limited divergence from a strict core-periphery structure; essentially, since each constraint can have its own scaling factor, the relative priority of any two constraints can change places across strata, but no more than one time.

a series of learning simulations in §4. The emergent preference for core-periphery structure is formally modeled in §5. Finally, conclusions and implications are discussed in §6.

2 Core-periphery structure and consistent faithfulness rankings

As Ito & Mester (1995ab, 1999, 2001, 2008) observe, a language with a phonologically stratified lexicon—in which there are lexical classes that differ in their phonological characteristics—often displays a CORE-PERIPHERY STRUCTURE: there is a phonologically restricted subset of the lexicon at the core, with increasingly less-restricted strata toward the periphery. This situation is illustrated in the following Venn diagram for (part of) the Japanese lexicon, based on the discussion in Ito & Mester (1999), which represents the subsets of lexical forms for which each of the markedness constraints is enforced.

(1) Lexical strata in a core-periphery structure



In the core stratum, ‘Native’, all three markedness constraints are enforced: NoNT, NoP, and NoTI. In each subsequent stratum, progressively fewer markedness constraints are enforced: NoP and NoTI in ‘Sino-Japanese’ (old loans from Chinese languages), only NoTI in ‘Assimilated Foreign’ (older and/or more nativized modern loans, chiefly from European languages), and none of the three in ‘Unassimilated Foreign’ (newer and/or less nativized modern loans). These constraints are defined in (2), following Ito & Mester (1995a, 1999); see these works for examples and discussion. See also Irwin (2011) for a recent overview of the history and synchronic characteristics of lexical strata in Japanese.

(2) Markedness constraints in the Japanese stratified lexicon

- (a) NoNT Assign one violation for every sequence of [+nasal] [–voice]
(‘No nasal–voiceless obstruent sequences’); Hayes (1999), Pater (2001)
- (b) NoP Assign one violation for every singleton (non-geminate) [p]
- (c) NoTI Assign one violation for every sequence of [Coronal, –son] [i]
(‘Coronal obstruents are palatal before [i]’)

In a constraint-based framework, such as OT or HG, a phonological RESTRICTION (predictable pattern; lack of contrast) is enforced by highly ranked or weighted markedness constraints (M), whereas a phonological CONTRAST (unpredictable pattern; lack of restriction) is enforced by highly ranked or weighted faithfulness constraints (F). A language in which different lexical strata have distinct but productive patterns of restriction and contrast is therefore particularly interesting: in such a language, the relative domination hierarchy between markedness and faithfulness constraints differs in the different

strata. Specifically, when lexical stratum A has a productive restriction that lexical stratum B does not have, we conclude that $M \gg F$ for A but $F \gg M$ for B . In the OT Stratified Faithfulness model (Fukazawa 1997; Fukazawa, Kitahara, & Ota 1998; Ito & Mester 1999, 2001, 2008), there is a distinct set of faithfulness constraints indexed to each lexical stratum. This allows for the grammar to simultaneously specify $M \gg FA$ for stratum A , enforcing a restriction, but $FB \gg M$ for stratum B , allowing contrast, so that the full ranking for the language has $FB \gg M \gg FA$.

The faithfulness constraints that conflict with the markedness constraints in (2), and take priority over them in the more peripheral strata of Japanese, are defined in (3).

- (3) Faithfulness constraints in the Japanese stratified lexicon
- (a) IDENT[voi] Assign one violation for every pair of corresponding segments that differ in their value for $[\pm\text{voice}]$ (McCarthy & Prince 1995)
 - (b) IDENT[p] *informally*: Assign one violation for every [p] that surfaces as [h]
(formally, this might be IDENT[Labial], or IDENT $[\pm\text{continuant}]$, or a cumulative effect of the two if implemented in HG)
 - (c) IDENT[ant] Assign one violation for every pair of corresponding segments that differ in their value for $[\pm\text{anterior}]$

In the OT Stratified Faithfulness model, the lexical strata in (1) can be analyzed in terms of the markedness ranking $\text{NoTI} \gg \text{NoP} \gg \text{NoNT}$, stratum-specific versions of the faithfulness constraints in (3) for strata U , A , S , and N ,² and rankings between opposing markedness and faithfulness constraints as in (4) (Ito & Mester 1999). The markedness constraints, which enforce phonological restrictions if undominated, are shown in bold.

- (4) OT Stratified Faithfulness rankings for the core-periphery phonology in (1)
- (a) Unassimilated Foreign:
IDENT[ant] $_U \gg \mathbf{NoTI}$ IDENT[p] $_U \gg \mathbf{NoP}$ IDENT[voi] $_U \gg \mathbf{NoNT}$
 - (b) Assimilated Foreign:
 $\mathbf{NoTI} \gg \text{IDENT[ant]}_A$ IDENT[p] $_A \gg \mathbf{NoP}$ IDENT[voi] $_A \gg \mathbf{NoNT}$
 - (c) Sino-Japanese:
 $\mathbf{NoTI} \gg \text{IDENT[ant]}_S$ $\mathbf{NoP} \gg \text{IDENT[p]}_S$ IDENT[voi] $_S \gg \mathbf{NoNT}$
 - (d) Native:
 $\mathbf{NoTI} \gg \text{IDENT[ant]}_N$ $\mathbf{NoP} \gg \text{IDENT[p]}_N$ $\mathbf{NoNT} \gg \text{IDENT[voi]}_N$

While the OT Stratified Faithfulness model is capable of representing a language with a strict core-periphery structure, as in the subset of the Japanese lexicon seen in (1) and (4), it is not capable of excluding a language that falls outside this structure (Fukazawa et al. 1998; Ito & Mester 1999). Even with the markedness constraints ranked in the relevant domination hierarchy $\text{NoTI} \gg \text{NoP} \gg \text{NoNT}$, nothing systematically excludes the existence of an additional stratum X in which outputs *satisfy* the low-ranking NoNT but *violate* the high-ranking NoTI . Such a stratum would have the ranking in (5); the

²Alternatively, a set of non-indexed, general faithfulness constraints could be used in place of a set of faithfulness constraints indexed to the core stratum (Native); see Ito & Mester (1999: note 38) for discussion.

effects of this ranking are seen in (6). Unfaithful structures in output forms are shown with **underline**, and (relevant) faithful structures are shown with *italic underline*.

(5) Constraint ranking for the non-core-periphery stratum X

$\text{IDENT}[\text{ant}]_X \gg \text{NoTI} \gg \text{NoNT} \gg \text{IDENT}[\text{voi}]_X$

(6) Stratum X optimal candidate violates high-ranking NoTI, satisfies low-ranking NoNT

$/\text{tinta}/_X$	$\text{IDENT}[\text{ant}]_X$	NoTI	NoNT	$\text{IDENT}[\text{voi}]_X$
→ a. <u>tinda</u>		*		*
b. <u>ɬinda</u>	* _W	L		*
c. <u>tɪnɬa</u>		*	* _W	L
d. <u>ɬɪnɬa</u>	* _W	L	*(_W)	L

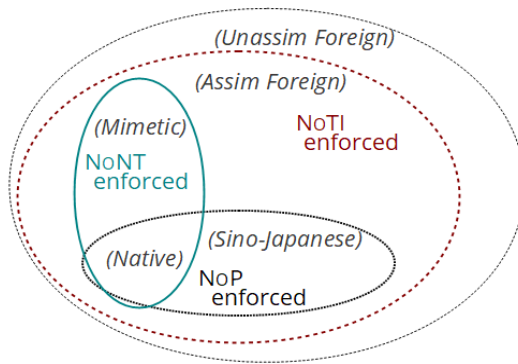
This non-core-periphery stratum X is possible as long as the faithfulness constraints that conflict with the ranked markedness constraints NoTI \gg NoNT are themselves independently rankable. In the OT Stratified Faithfulness model, there is no simple non-stipulative way to prevent a high-ranking $\text{IDENT}[\text{ant}]_X$ from dominating NoTI, rendering it inactive, while $\text{IDENT}[\text{voi}]_X$ is still ranked below NoNT.

Ito & Mester (1999) and Fukazawa et al. (1998) explore metaconditions on faithfulness rankings that would prevent a stratum like X from arising, and Hsu & Jesney's (2017a) Weighted Scalar Constraints model ensures that such a stratum is not formally possible. The key insight behind both the Stratified Faithfulness ranking metaconditions and the Weighted Scalar Constraints approach is to ensure that *the hierarchy among faithfulness constraints is consistent* from stratum to stratum. In terms of the example in (6), this restriction would exclude any stratum X in which $\text{IDENT}[\text{ant}]_X \gg \text{IDENT}[\text{voi}]_X$, which in turn excludes the ranking in (5) that creates the non-core-periphery pattern.

However, the outright formal exclusion of a language with a non-core-periphery pattern is not in fact desirable (see, e.g., Fukazawa et al. 1998; Inkelas & Zoll 2007). For example, a fuller picture of the lexicon of Japanese includes the Mimetic stratum, consisting of sound-symbolic and other expressive forms. Crucially, the Mimetic and the Sino-Japanese strata do not stand in a subset/superset relationship: NoNT applies only to Mimetic (and Native) forms, and NoP applies only to Sino-Japanese (and Native) forms.³ This situation is shown in the Venn diagram in (7), after Ito & Mester (1995b: 823).

³Ito & Mester (1995a: 190) actually argue that Mimetic forms do not in fact have a ranking equivalent to $\text{IDENT}[\text{p}]_M \gg \text{NoP}$, because [p] can occur only in initial position in this stratum, making the relevant faithfulness constraint a positional one (e.g., Beckman 1999), along the lines of $\text{IDENT}[\text{p}]-\sigma 1_M$. However, the simplified analysis as given in (7) is still useful for the purposes of the learning simulation in §4.4, where it provides an example of a schematic language with a non-core-periphery structure. As for the phonological analysis of lexical strata in actual Japanese, the broader point still remains that the Mimetic and Sino-Japanese strata do not stand in a subset/superset relation.

(7) Non-core-periphery structure in the Japanese lexicon, with Mimetic forms



On the other hand, a model of stratified phonology that assigns no preference at all to strict core-periphery structure is too weak. Fukazawa et al. (1998) suggest that existing cases of core-periphery structure are merely epiphenomenal, in that new strata are only created when markedness-violating loanwords are encountered, so diachronically later strata typically happen to be more phonotactically permissive. But this explanation does not account for the existence of *productive* impossible-nativization effects (Ito & Mester 1999, 2001; Pinta 2013, Smith & Pinta 2017)—for at least some speakers of some languages, nativization patterns that fall outside a strict core-periphery structure are actively dispreferred. Such effects show that the phonological grammar does, in some way, *prefer* to maintain a core-periphery structure for a stratified lexicon.

The HG Stratified Faithfulness model, introduced in the following section, finds a middle ground. In learning simulations exploring this model, not only does the learner acquire the markedness hierarchy underpinning a core-periphery structure. Crucially, the basis for a preference for a hierarchy of cumulative faithfulness weights that is consistent across lexical strata likewise emerges automatically, unless there is overt evidence to the contrary. This approach thus accounts both for the existence of stratified phonologies that do not have a strict core-periphery structure, and also for the existence of speakers with productive impossible-nativization effects.

3 The HG Stratified Faithfulness model

Harmonic Grammar (HG; Legendre, Miyata, & Smolensky 1990) differs from Optimality Theory (Prince & Smolensky 1993/2004) in that HG constraints are weighted rather than strictly ranked. This difference allows HG to show CUMULATIVE CONSTRAINT INTERACTION, also known as ‘gang effects’: the violations of a set of lower-weighted constraints can, under the right conditions (Pater 2009, 2016), ‘gang up’ and assign a higher overall penalty than the violation of a higher-weighted constraint.

Jesney & Tessier (2011) demonstrate that such cumulative constraint interaction plays a fundamental role in establishing the overall influence of specific and general constraints in the grammar of a language. For example, consider a language in which stressed syllables resist a place-assimilation process that targets unstressed syllables. Place assimilation of a segment in a stressed syllable would violate both a general faithfulness constraint, IDENT[Place], and its positional version indexed to stressed syllables (Beckman 1999), IDENT[Place]-σ. Crucially, in HG, neither of these IDENT constraints alone actually needs to outweigh the constraint driving place assimilation, as long as the *cumulative* weight of the two is higher than that of the assimilation constraint.

The HG STRATIFIED FAITHFULNESS model proposed here extends this insight to account for stratified phonology, with or without a core-periphery structure. The model includes general faithfulness constraints, not indexed to any stratum, along with stratum-specific faithfulness constraints. Because the learner begins with the Initial State weighting relation of $w(M) > w(F)$ (Smolensky 1996; Jesney & Tessier 2011), any learning datum whose target surface form shows faithfulness effects will initially produce a non-target output (i.e., an error) for the learner, triggering an incremental increase in weight for all relevant faithfulness constraints and a decrease for all relevant markedness constraints. For example, a learning datum such as /inta/s→[inta]s, indexed to the Sino-Japanese stratum, provides evidence for promoting faithfulness to voicing above the markedness constraint NoNT. But this will raise the weight not only of IDENT[voi]s, but also of general IDENT[voi], until the cumulative weights of the two constraints are enough to overcome NoNT.

The crucial consequence of modeling stratified phonology with a cumulative interaction involving general faithfulness constraints is that, because such constraints are relevant to all strata, the relative frequency with which any given general faithfulness constraint is promoted over the course of grammar-learning depends on the proportion of strata in which it is satisfied. This in turn means that, when the learner is exposed to a language that has a strict core-periphery structure, not only the markedness constraints *but also the general faithfulness constraints* end up with a relative weighting that reflects this structure. As discussed in §5 below, thanks to cumulative constraint interaction in HG, it is this relative hierarchy among the general faithfulness constraints that biases the grammar toward adherence to strict core-periphery structure *even for a potential novel stratum*, giving rise to impossible-nativization effects. Before this discussion of the emergent core-periphery bias, however, the next section (§4) first presents a series of learning simulations, to confirm that a simulated learner with an HG Stratified Faithfulness grammar behaves as predicted when it is exposed to a core-periphery phonology and a non-core-periphery phonology.

4 HG-GLA learning simulations

Learning simulations were carried out in the HG Stratified Faithfulness model for three schematic languages: a language with a strict core-periphery structure in which all strata contain the same number of lexemes; a language with a strict core-periphery structure in which the most peripheral stratum has a higher proportion of lexemes; and a language that has a non-core-periphery structure.

A version of the Gradual Learning Algorithm (Boersma & Hayes 2001) implemented for Harmonic Grammar (the HG-GLA; Jesney & Tessier 2011, Boersma & Pater 2016) was trained on each schematic language in order to simulate the acquisition of the constraint weights needed for each grammar.⁴ To preview the results, both strict core-periphery languages ended up with a general-faithfulness hierarchy that supports a bias against impossible nativizations (as discussed in §5 below), and even the non-core-periphery language was successfully acquired.

⁴These learning simulations make two key simplifying assumptions—as, implicitly, do Ito & Mester (1995ab, 1999), and Hsu & Jesney (2017a)—namely, that the learner knows which lexemes belong to which strata, and that the learner has access to the correct underlying forms for unfaithful outputs. For further discussion of how the strata themselves might be identified and modeled by a learner, see Fukazawa et al. 1998; Pater 2005, 2010; and Hayes (2016).

4.1 The schematic target languages

The three schematic languages used in the learning simulations are simplified versions of the grammar of Japanese as analyzed by Ito & Mester (1995ab, 1999). Each language has either four or five lexical strata, which are distinguished by their patterns of enforcement of the three markedness constraints defined in (2) above. The lexical strata are summarized in (8), where ‘*’ indicates that a constraint is *not enforced* in the stratum in question and ‘✓’ indicates that a constraint is *enforced*.


(8) Lexical strata in the schematic target languages

Stratum	NoNT	NoP	NoTI
<i>U</i> Unassimilated Foreign	*	*	*
<i>A</i> Assimilated Foreign	*	*	✓
<i>S</i> Sino-Japanese	*	✓	✓
<i>M</i> Mimetic (where relevant)	✓	*	✓
<i>N</i> Native	✓	✓	✓

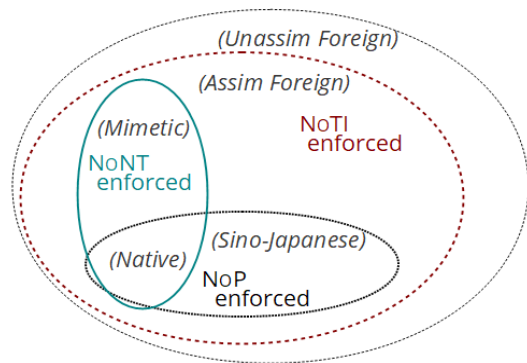
As discussed in §2, a language with all five strata does not form a strict core-periphery structure, because neither the Mimetic nor the Sino-Japanese stratum is a subset of the other in terms of markedness constraint domains. The learning simulations included both strict core-periphery languages with no Mimetic stratum, as in (9)(a) (repeated from (1)), as well as the full non-core-periphery system, as in (9)(b) (repeated from (7)).

(9) Structure of the stratified lexicon in the schematic target languages

(a) Strict core-periphery structure: Mimetic stratum removed (languages #1, #2)



(b) Non-core-periphery structure: Mimetic stratum included (language #3)



The faithfulness constraints used in the learning simulations are those defined in (3) above. As discussed in §3, in the HG Stratified Faithfulness model, each faithfulness constraint has a general version as well as a specific version indexed to each stratum. The constraint set used here includes a set of stratal faithfulness constraints for the Native (core) stratum, although in the schematic languages under consideration, this is not formally necessary. Since the Native stratum satisfies all three markedness constraints, showing no effect of faithfulness for any of the properties under discussion, this stratum could be modeled using the general faithfulness constraints only (and indeed, the weights of the faithfulness constraints indexed to the Native stratum remain zero in all of the learning simulations).

The complete constraint set used in the learning simulations is therefore as given in (10).

(10) Constraint set for learning simulations

(a) Markedness constraints	NoNT	NoP	NoTI
(b) General faithfulness constraints	IDENT[voi]	IDENT[p]	IDENT[ant]
(c) Stratal faithfulness constraints			
• Unassimilated Foreign	IDENT[voi] <i>U</i>	IDENT[p] <i>U</i>	IDENT[ant] <i>U</i>
• Assimilated Foreign	IDENT[voi] <i>A</i>	IDENT[p] <i>A</i>	IDENT[ant] <i>A</i>
• Sino-Japanese	IDENT[voi] <i>S</i>	IDENT[p] <i>S</i>	IDENT[ant] <i>S</i>
• Mimetic (where relevant)	IDENT[voi] <i>M</i>	IDENT[p] <i>M</i>	IDENT[ant] <i>M</i>
• Native	IDENT[voi] <i>N</i>	IDENT[p] <i>N</i>	IDENT[ant] <i>N</i>

The schematic languages presented to the learner consist of three words assigned to each stratum, for a total of twelve words (in the four-stratum, strict core-periphery languages) or fifteen words (in the five-stratum, non-core-periphery language). Each word has exactly one structure that violates one of the markedness constraints under discussion—a [nt] sequence, a singleton [p], or a [ti] sequence—in its input (underlying) form. Whether or not this structure surfaces faithfully in the target language depends on the stratum to which the word is assigned, as summarized in (8) above. Inputs and outputs for each word are given in (11); unfaithful structures in output forms are shown with **bold underline**, while (relevant) faithful structures are shown with *italic underline*.

(11) Words in the schematic target languages

Stratum	/nt/ sequence	singleton /p/	/ti/ sequence
<i>U</i> Unassimilated Foreign	/inta/ <i>U</i> → [in <u>ɾ</u> a] <i>U</i>	/paku/ <i>U</i> → [<u>p</u> aku] <i>U</i>	/mati/ <i>U</i> → [ma <u>t</u> i] <i>U</i>
<i>A</i> Assimilated Foreign	/inta/ <i>A</i> → [in <u>ɾ</u> a] <i>A</i>	/paku/ <i>A</i> → [<u>p</u> aku] <i>A</i>	/mati/ <i>A</i> → [ma <u>t</u> ɛi] <i>A</i>
<i>S</i> Sino-Japanese	/inta/ <i>S</i> → [in <u>ɾ</u> a] <i>S</i>	/paku/ <i>S</i> → [<u>h</u> aku] <i>S</i>	/mati/ <i>S</i> → [ma <u>t</u> ɛi] <i>S</i>
<i>M</i> Mimetic (where relevant)	/inta/ <i>M</i> → [in <u>ɾ</u> a] <i>M</i>	/paku/ <i>M</i> → [<u>p</u> aku] <i>M</i>	/mati/ <i>M</i> → [ma <u>t</u> ɛi] <i>M</i>
<i>N</i> Native	/inta/ <i>N</i> → [in <u>ɾ</u> a] <i>N</i>	/paku/ <i>N</i> → [<u>h</u> aku] <i>N</i>	/mati/ <i>N</i> → [ma <u>t</u> ɛi] <i>N</i>

Three different learning scenarios were considered. The first (§4.2), as a baseline case, was a strict core-periphery language (no Mimetic stratum) with a uniform distribution of lexical items across strata. The second (§4.3) was still a strict core-periphery language, but had more lexical items in Unassimilated Foreign than in the other strata, in order to explore the effect of a non-uniform distribution of lexical items on the relative weights of the constraints. The third (§4.4) was a language with all five strata,

including Mimetic, and a uniform distribution of items across strata; this language was included in order to confirm that a non-core-periphery structure could be acquired, and if so, to determine its effect on the relationships among the constraint weights.

HG-GLA learning simulations were carried out in Praat (version 5.4.16; Boersma & Weenink 2015) and were structured like those of Jesney & Tessier (2011). Initial weights were set at 100 for the markedness constraints and at 0 for the faithfulness constraints. This difference encodes the $M \gg F$ initial-state bias required for learning restrictive grammars in the absence of overt alternations (Smolensky 1996) and assigns the faithfulness constraints weights that are low enough to avoid unintended cumulative effects (Jesney & Tessier 2011). Plasticity, the amount by which a constraint's weight is raised or lowered at each learning step where the target output is not yet selected, was set at 1.0 for the markedness constraints and 0.2 for the faithfulness constraints; Jesney & Tessier (2011) demonstrate that weights must change more quickly for markedness constraints than for faithfulness constraints in order once again to prevent unwanted types of cumulative constraint interaction involving faithfulness. Finally, since there is no variation in the schematic languages under consideration, evaluation noise was set at 0, there was no plasticity decrement ('number of plasticities' was set at 1), and relative plasticity spreading was set at 0. The decision strategy used was 'LinearOT', which excludes negative values for weights, and the reranking method was 'Symmetric All', which means that each time a learner's output failed to match the target-language output, all constraints favoring the learner's current output had their weights lowered and all constraints favoring the target-language output had their weights raised (according to the relevant plasticity settings). There were 100,000 learning data presented in each learning simulation, chosen randomly according to the frequency distribution of forms in the target language. Five separate simulations were run for each schematic language to confirm that the grammars were converging consistently on the end-state pattern; as is discussed in more detail below, the resulting grammars were indeed consistent.

4.2 Schematic language #1: Strict core-periphery structure, uniform distribution

The first schematic language tested has a strict core-periphery structure among lexical strata, and its lexical items are evenly distributed among strata so that information from each stratum has the same degree of influence on the learning trajectory. This target language corresponds to (9)(a), with no Mimetic stratum; thus, there are no faithfulness constraints for stratum M , and no lexical items with the M index. The relative frequency of all 12 remaining lexical items (see (11)) was set at 1.

A representative set of HG-GLA weights learned for each constraint (the results from one of the five simulations) is shown in (12). Also shown is the sum of the weight of each stratal faithfulness constraint and the weight of its associated general faithfulness constraint, representing the cumulative faithfulness interaction, as well as the difference between this cumulative faithfulness weight and the weight of the antagonistic markedness constraint ($F-M$) for each phonological pattern. It is this last value that determines whether a particular structure is realized faithfully or nativized in a given stratum. If the $F-M$ value is positive, then the faithfulness constraints (cumulatively) are weighted higher than the conflicting markedness constraint, and the relevant structure is faithfully preserved. If the $F-M$ value is negative, then the markedness constraint's weight is higher than the combined weights of the general and stratal faithfulness constraints, and the relevant structure is nativized.

Final constraint weights might in principle differ somewhat across multiple learning simulations, since the order in which forms are encountered by the learner might result in weights being distributed differently among constraints. As it turns out, however, the results of the five simulations for this language scenario were very consistent (the range of weights for each constraint across the five simulations is shown in (12)). There was almost no variation for either the markedness constraints or the

general faithfulness constraints: NoTI, NoP, IDENT[ant], and IDENT[p] had the same weights in all five simulations; NoNT differed by 1 and IDENT[voi] by 0.2 (one learning step each) across simulations. The stratal faithfulness constraints showed a little more variability; specifically, in cases where multiple strata are exempt from the same markedness requirement, the relative weight of the (non-zero) stratal faithfulness constraints will vary a bit from simulation to simulation. This happens because individual learning simulations differ in exactly how many forms from each stratum are encountered by the learner, and in which order, during the phase before the general-faithfulness weights are raised enough to help the stratal faithfulness constraints overcome the markedness constraints—leading to differences across simulations in how far the different stratal faithfulness constraints have their weights raised before target forms are always produced and learning stops. But even here, the most variation seen, for IDENT[voi] constraints, was 2.6 for Assimilated Foreign, 2 for Unassimilated Foreign, and 1.8 for Sino-Japanese, or 13, 10, and 9 learning steps respectively. As a comparison, the constraints whose weights changed the most were NoNT and IDENT[voi], each changing by 80 or 81 learning steps per simulation.


(12) HG-GLA outcome: Core-periphery language, uniform distribution

<i>category</i>	<i>constraint</i>	<i>representative weight</i>	<i>range</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
markedness constraints	NoTI	28	0			
	NoP	22	0	—	—	—
	NoNT	19	1			
general faithfulness constraints	IDENT[ant]	14.4	0			
	IDENT[p]	15.6	0	—	—	—
	IDENT[voi]	16.2	0.2			
Unassimilated Foreign	IDENT[ant] <i>U</i>	14.4	0	28.8	0.8	<i>faithful</i> [ti] <i>U</i>
	IDENT[p] <i>U</i>	7	1.4	22.6	0.6	<i>faithful</i> [p] <i>U</i>
	IDENT[voi] <i>U</i>	7	2	23.2	4.2	<i>faithful</i> [nt] <i>U</i>
Assimilated Foreign	IDENT[ant] <i>A</i>	0	0	14.4	–13.6	nativized [tei] <i>A</i>
	IDENT[p] <i>A</i>	8.6	1.4	24.2	2.2	<i>faithful</i> [p] <i>A</i>
	IDENT[voi] <i>A</i>	5	2.6	21.2	2.2	<i>faithful</i> [nt] <i>A</i>
Sino-Japanese	IDENT[ant] <i>S</i>	0	0	14.4	–13.6	nativized [tei] <i>S</i>
	IDENT[p] <i>S</i>	0	0	15.6	–6.4	nativized [h] <i>S</i>
	IDENT[voi] <i>S</i>	4.2	1.8	20.4	1.4	<i>faithful</i> [nt] <i>S</i>
Native	IDENT[ant] <i>N</i>	0	0	14.4	–13.6	nativized [tei] <i>N</i>
	IDENT[p] <i>N</i>	0	0	15.6	–6.4	nativized [h] <i>N</i>
	IDENT[voi] <i>N</i>	0	0	16.2	–2.8	nativized [nd] <i>N</i>

The constraint weights from (12) are plotted in (13), with lexical strata along the *x* axis and constraint weights along the *y* axis. Markedness constraints and general faithfulness constraints are labeled at the left edge of the plot; markedness constraints are plotted with filled symbols, and their antagonistic faithfulness constraints are plotted with the corresponding open symbols. Where a stratal faithfulness constraint has a non-zero weight, that weight is plotted as an *addition to* the value of the weight of the general faithfulness constraint, representing their cumulative constraint interaction. This cumulative faithfulness weight is plotted with the same symbol as the general version of the faithfulness constraint,

and is connected to the general-faithfulness weight by a vertical line. Finally, a horizontal line is plotted at the weight of each markedness constraint in order to emphasize the value that the cumulative weight of the opposing faithfulness constraints (general and stratal) would have to surpass in order for the structure in question to be realized faithfully.

(13) Relationships among weights: Core-periphery language, uniform distribution



The results in (12) and (13) show that the relative weights among the markedness constraints are equivalent to the markedness rankings that would be proposed in the original OT Stratal Faithfulness model—a domination hierarchy that parallels the markedness subset/superset relation among the lexical strata. In other words, the markedness constraint that is violated in the most strata, NoNT, is weighted lowest, and the one that is enforced in the most strata, NoTI, is weighted highest: $w(\text{NoTI}) > w(\text{NoP}) > w(\text{NoNT})$. Likewise, the general faithfulness constraints are ordered in exactly the relationship that would be enforced by a metacondition on cross-stratum faithfulness ranking in the original OT Stratified Faithfulness model. $\text{IDENT}[\text{voi}]$, the general faithfulness constraint that is relevant to the contrast preserved in the most strata, is weighted highest, and $\text{IDENT}[\text{ant}]$, relevant to the contrast that is preserved only in the most peripheral stratum, is weighted lowest: $w(\text{IDENT}[\text{voi}]) > w(\text{IDENT}[\text{p}]) > w(\text{IDENT}[\text{ant}])$. Conceptually, it is clear why the HG-GLA produces these results. The markedness constraint that is violated in the most strata will undergo the most demotion in weight over the course of learning, and will therefore end up lowest. Conversely, the general faithfulness constraint that is satisfied in the most strata will undergo the most promotion in weight, and will end up highest.

The role of stratal faithfulness constraints in the current model, on the other hand, is very different from the role that they play in the OT Stratified Faithfulness model. Here, the stratal faithfulness constraints enter into gang effects with the general faithfulness constraints, and this cumulative interaction affects how the specific faithfulness constraints are weighted. Unsurprisingly, the stratal versions of faithfulness constraints that are violated in a particular stratum are weighted at zero, since the learner sees no

evidence that they are ever active. However, the stratal faithfulness constraints with non-zero weights show a pattern that is essentially the *inverse* of that among the general faithfulness constraints: $w(\text{IDENT}[\text{ant}]) > w(\text{IDENT}[\text{p}]) > w(\text{IDENT}[\text{voi}])$. Again, given the nature of the HG-GLA, it is clear why this is so. As discussed above, general $\text{IDENT}[\text{voi}]$ is weighted highest among the general faithfulness constraints, and the conflicting markedness constraint NoNT is weighted lowest among the markedness constraints—target forms from all but the Native stratum realize voicing faithfully in /nt/ clusters, so $\text{IDENT}[\text{voi}]$ is promoted the most and NoNT is demoted the most. This means that the additional ‘boost’ needed for any given stratum-specific version of $\text{IDENT}[\text{voi}]$ to make the cumulative voicing-faithfulness weight higher than that of NoNT is smaller than that needed for $\text{IDENT}[\text{p}]$, which in turn is smaller than that for $\text{IDENT}[\text{ant}]$.

In summary, when an HG-GLA learner is exposed to a language with a stratified lexicon that has a strict core-periphery structure and a uniform distribution of forms across strata, the markedness constraints form a domination hierarchy that mirrors the core-periphery structure, and the general faithfulness constraints fall into a reverse hierarchy so that the faithfulness constraint supporting the most-marked property (according to the markedness hierarchy) is ranked the lowest. This general-faithfulness hierarchy plays a key role in the emergent core-periphery preference, as discussed in §5 below.

4.3 Schematic language #2: Core-periphery structure, non-uniform distribution

The second schematic language is designed to examine the effect of a case in which a peripheral stratum makes up the majority of the lexicon, to see whether this changes the relationships among the markedness constraints or among the general faithfulness constraints as compared to the baseline case with an even distribution of forms across strata. The target language once again corresponds to (9)(a), with no Mimetic stratum. This time, however, the relative frequency of lexical items across strata was manipulated so that the words in the most peripheral stratum, Unassimilated Foreign, were presented to the learner five times as frequently as those from the other three strata.


Results are shown in (14) and (15). Here again, there was little variability in the final weights assigned to each constraint across the five learning simulations, so the values for one representative simulation are presented in the discussion (and the range for each constraint across simulations is included in (14)).

(14) HG-GLA outcome: Core-periphery language, non-uniform distribution

<i>category</i>	<i>constraint</i>	<i>representative weight</i>	<i>range</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
markedness constraints	NoTI	28	0			
	NoP	20	1	—	—	—
	NoNT	18	2			
general faithfulness constraints	IDENT[ant]	14.4	0			
	IDENT[p]	16	0.2	—	—	—
	IDENT[voi]	16.4	0.4			
Unassimilated Foreign	IDENT[ant] <i>U</i>	14.4	0	28.8	0.8	<i>faithful</i> [ti] <i>U</i>
	IDENT[p] <i>U</i>	11.8	1.4	27.8	7.8	<i>faithful</i> [p] <i>U</i>
	IDENT[voi] <i>U</i>	10.4	2.2	26.8	8.8	<i>faithful</i> [nt] <i>U</i>

<i>category</i>	<i>constraint</i>	<i>representative weight</i>	<i>range</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
Assimilated Foreign	IDENT[ant]A	0	0	14.4	–13.6	nativized [tɛi]A
	IDENT[p]A	4.2	1.2	20.2	0.2	<i>faithful</i> [p]A
	IDENT[voi]A	2.8	2.2	19.2	1.2	<i>faithful</i> [nt]A
Sino-Japanese	IDENT[ant]S	0	0	14.4	–13.6	nativized [tɛi]S
	IDENT[p]S	0	0	16	–4	nativized [h]S
	IDENT[voi]S	3.2	0.8	19.6	1.6	<i>faithful</i> [nt]S
Native	IDENT[ant]N	0	0	14.4	–13.6	nativized [tɛi]N
	IDENT[p]N	0	0	16	–4	nativized [h]N
	IDENT[voi]N	0	0	16.4	–1.6	nativized [nd]N

(15) Relationships among weights: Core-periphery language, non-uniform distribution



The results of this set of simulations are not qualitatively different from those discussed in §4.2, even though this time the majority of the lexicon belongs to the Unassimilated Foreign stratum, and so the learner was exposed to more faithful than unfaithful forms for all three phonological structures at hand.

The main difference for this schematic language seems to be that the weights for all faithfulness constraints specific to the Unassimilated Foreign stratum are relatively high—even those for IDENT[voi]U and IDENT[p]U, which have ended up with much higher weights than are objectively needed, given the final weights of general IDENT[voi] and IDENT[p]. Conceptually, forms from stratum U were encountered so frequently in the course of the learning simulation that the weights of IDENT[voi]U and IDENT[p]U were

increased nearly every time those of their general counterparts were increased; as a result, the U faithfulness constraints contribute about half of the general+stratal cumulative faithfulness weight that would be needed to outweigh the conflicting markedness constraints for Unassimilated Foreign forms. As learning progressed, however, encounters with target forms from other strata—where the stratal faithfulness constraints had not had their own weights raised so quickly, and so non-target forms were still being selected—would continue to raise the weights of general IDENT[voi] and IDENT[p].

For the patterns of greatest interest here, however, there is essentially no difference between this core-periphery language and the one with a uniform distribution of lexical items. Among the markedness constraints we see $w(\text{NoTI}) > w(\text{NoP}) > w(\text{NoNT})$, while among the general faithfulness constraints we see $w(\text{IDENT}[\text{voi}]) > w(\text{IDENT}[\text{p}]) > w(\text{IDENT}[\text{ant}])$, just as for the first core-periphery language. This comparison shows that under the HG Stratified Faithfulness approach, any language with a strict core-periphery structure, regardless of the relative sizes of the different strata, will give the learner an HG grammar in which both the markedness hierarchy and the general faithfulness hierarchy reflect that core-periphery structure.

4.4 Schematic language #3: Non-core-periphery structure

The final set of learning simulations is designed to examine the constraint weights acquired for a language with a non-core-periphery structure. This time, the target language corresponds to (9)(b), including the Mimetic stratum. There are five lexical strata and five sets of stratum-specific faithfulness constraints, but the Mimetic and Sino-Japanese strata do not stand in a subset/superset relationship: as shown above in (8) and (11), singleton [p] is tolerated in Mimetic forms, and nasal–voiceless obstruent sequences are tolerated in Sino-Japanese forms, but not vice versa. The relative frequency of all lexical items for this schematic language was set at 1, so the lexicon was evenly distributed among the lexical strata.


Results are shown in (16) and (17); as above, one representative set of weights is reported here, and the range of weights assigned across the five simulations is also given in (16).

(16) HG-GLA outcome: Non-core-periphery language

<i>category</i>	<i>constraint</i>	<i>representative weight</i>	<i>range</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
markedness constraints	NoTI	28	0			
	NoP	20	1	—	—	—
	NoNT	20	0			
general faithfulness constraints	IDENT[ant]	14.4	0			
	IDENT[p]	16	0.2	—	—	—
	IDENT[voi]	16	0			
Unassimilated Foreign	IDENT[ant] U	14.4	0	28.8	0.8	<i>faithful</i> [ti] U
	IDENT[p] U	5.4	1.8	21.4	1.4	<i>faithful</i> [p] U
	IDENT[voi] U	5.8	1	21.8	1.8	<i>faithful</i> [nt] U
Assimilated Foreign	IDENT[ant] A	0	0	14.4	–13.6	nativized [tɛi] A
	IDENT[p] A	5.2	2	21.2	1.2	<i>faithful</i> [p] A
	IDENT[voi] A	5.8	1.2	21.8	1.8	<i>faithful</i> [nt] A

<i>category</i>	<i>constraint</i>	<i>representative weight</i>	<i>range</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
Sino-Japanese	IDENT[ant]S	0	0	14.4	–13.6	nativized [tei]S
	IDENT[p]S	0	0	16	–4	nativized [h]S
	IDENT[voi]S	4.4	1.8	20.4	0.4	<i>faithful</i> [nt]S
Mimetic	IDENT[ant]M	0	0	14.4	–13.6	nativized [tei]M
	IDENT[p]M	5.4	1.8	21.4	1.4	<i>faithful</i> [p]M
	IDENT[voi]M	0	0	16	–4	nativized [nd]M
Native	IDENT[ant]N	0	0	14.4	–13.6	nativized [tei]N
	IDENT[p]N	0	0	16	–4	nativized [h]N
	IDENT[voi]N	0	0	16	–4	nativized [nd]N

(17) Relationships among weights: Non-core-periphery language



NoTI has the highest weight among the markedness constraints, and IDENT[ant] has the lowest weight among the general faithfulness constraints, as in the previous learning scenarios. This time, however, the general constraints pertaining to singleton [p] and to post-nasal voicing are tied: NoP and NoNT have the same weight, as do IDENT[p] and IDENT[voi]. The ties have come about because singleton [p] and faithful [nt] are permitted in the same number of strata: [p] in Mimetic, Assimilated Foreign, and Unassimilated Foreign, and [nt] in Sino-Japanese, Assimilated Foreign, and Unassimilated Foreign. As for the stratal faithfulness constraints, in each case where they are relevant (have a weight greater than zero), they have been apportioned just enough weight to allow for the cumulative effects between the general and the stratum-specific faithfulness constraints to overcome the opposing markedness constraints.

Thus, the HG Stratified Faithfulness model presented here allows an HG-GLA learner to acquire a set of weights for markedness, general faithfulness, and stratal faithfulness constraints that accurately models even a non-core-periphery language. When there is no strict core-periphery structure in the language being acquired, there will not be a strict domination hierarchy among the markedness constraints or among the general faithfulness constraints. Still, even this schematic language shows a final ranking that is as nearly consistent with core-periphery structure as the learning data will permit.

4.5 Summary: Learning simulations

The learning simulations described in this section have tested the HG Stratified Faithfulness model against a set of schematic languages with stratified phonologies. Results show that not only the stratum-defining markedness constraints, but also the relevant general faithfulness constraints, are ordered by the learner in a domination hierarchy that is determined by the number of lexical strata in which each constraint is enforced.

One implication of these findings is that even a language with a *non*-core-periphery phonology, in which not all lexical strata stand in a subset/superset relation, can be learned—a desired outcome, since languages with this pattern are attested.

For languages that do have a core-periphery structure, this direct connection between the number of strata in which a constraint is enforced and the relative weight of that constraint has additional implications. First, because no two of the stratum-distinguishing markedness constraints are enforced in the same number of strata, no two of them will have the same weight.⁵ This establishes a strict domination hierarchy among the markedness constraints, capturing Ito & Mester’s (1995ab, 1999) fundamental insight that a core-periphery phonology has such a markedness hierarchy as its underpinning.

The second implication for the learning of a strict core-periphery language is the most novel contribution of the HG Stratified Faithfulness model: the general versions of the faithfulness constraints that are relevant for distinguishing among strata are likewise ordered by the learner in a strict domination hierarchy. As demonstrated in the following section, this general-faithfulness hierarchy provides the basis for the grammar’s emergent preference for core-periphery structure, thereby accounting for productive impossible-nativization effects, without invoking the formally complex type of metacondition on faithfulness rankings that is necessary in the OT Stratified Faithfulness model. Moreover, the current model predicts that even languages with only a partial core-periphery structure can still show limited impossible-nativization effects, which is consistent with reports of such effects in Japanese by Ito & Mester (1999).

5 Core-periphery structure as an emergent bias

The results of §4 show that a HG Stratified Faithfulness learner, when exposed to a stratified phonology with a core-periphery structure, acquires a domination hierarchy among the general faithfulness constraints that reflects that core-periphery structure. This section now proposes a means by which such a general-faithfulness domination hierarchy creates an EMERGENT BIAS toward maintaining strict core-

⁵This wording is a minor simplification for ease of exposition. In a language where two markedness (or two faithfulness) constraints are both satisfied in the same *set* of strata—not just the same *number* of strata—the learner would assign them the same weight, but this scenario is still consistent with the language having a strict core-periphery phonology, and is clearly distinct from the case of NoP and NoNT in (9)(b) and the associated learning simulation in §4.4.

periphery structure in any new lexical strata that might be added to the language. This emergent bias, made possible by cumulative constraint interaction in HG, correctly predicts the existence of productive impossible-nativization effects.

An impossible-nativization effect, as discussed by Ito & Mester (1999, 2001), is the rejection or dispreference by native speakers of a form in which a lower-ranked (lower-weighted) markedness constraint is enforced but a higher-ranked (higher-weighted) one is not, resulting in the *nativization* of a ‘less-foreign’ property along with the *faithful preservation* of a ‘more-foreign’ property. Impossible-nativization candidates are inconsistent with strict core-periphery structure, but as shown in (6) above, repeated here as (18), the OT Stratified Faithfulness model cannot on its own rule out the introduction of a novel stratum X in which an impossible-nativization candidate is optimal.

(18) Stratum X optimal candidate violates high-ranking NoTI, satisfies low-ranking NoNT

/tinta/ X	IDENT[ant] X	NoTI	NoNT	IDENT[voi] X
→ a. <u>t</u> inda		*		*
b. <u>t</u> inda	* _W	L		*
c. <u>t</u> inta		*	* _W	L
d. <u>t</u> inta	* _W	L	*(_W)	L

Under OT Stratified Faithfulness, the only way to prevent any lexical strata from ever allowing the mapping /tinta/ X → [tinda] X —an impossible nativization, given the established markedness ranking of NoTI » NoNT—is to introduce a metacondition that keeps the ranking among the faithfulness constraints consistent across strata. In (18), this would prevent IDENT[ant] X from dominating IDENT[voi] X , given evidence for IDENT[voi] » IDENT[ant] in other strata. Without IDENT[ant] X » IDENT[voi] X , candidate (a) cannot win.

Unfortunately, it becomes complex and stipulative to formulate a requirement (or even, for better empirical accuracy, a soft bias or defeasible preference) that IDENT[voi] Z » IDENT[ant] Z for all strata Z merely because there is some stratum Q where IDENT[voi] Q » IDENT[ant] Q . (See Ito & Mester (1999) and Fukazawa et al. (1998) for two different formal implementations of such a metacondition.) The problem is that the entire effect of faithfulness to the feature [\pm anterior] for any given stratum Z is due to one constraint, IDENT[ant] Z , and likewise for [\pm voice] and IDENT[voi] Z .⁶ Crucially, there is no intrinsic relationship between IDENT[ant] Z for stratum Z and IDENT[ant] Q for any other stratum Q ; they operate entirely independently in the constraint hierarchy. As a result, defining and enforcing a metacondition on ranking relationships across strata poses a thorny problem.

In the HG Stratified Faithfulness model introduced here, by contrast, there is a direct formal relationship between the relative degree of faithfulness to [\pm voice] versus [\pm anterior] in stratum Z and that in stratum Q . First, if the preservation of /ti/ is a property of only the most-peripheral stratum, but the preservation of /nt/ is a property of all but the most-core stratum, then *general* IDENT[voi] is satisfied in more strata than *general* IDENT[ant], and so the grammar *already includes* the general-faithfulness hierarchy $w(\text{IDENT}[\text{voi}]) > w(\text{IDENT}[\text{ant}])$ (see §3 and §4). Second, because of cumulative constraint interaction in HG, these general faithfulness constraints actually contribute to the outcome of the M-vs-F competition

⁶In the language under consideration, a general IDENT[ant] or IDENT[voi] constraint would be dominated by the antagonistic markedness constraint, NoTI or NoNT respectively, and would therefore play no role in the phonology. If these general faithfulness constraints were not dominated, then /t/-palatalization and post-nasal voicing would never occur, and so would not be diagnostic of lexical strata in the first place.

in *all* strata. That is, the weight of the ‘faithfulness effect’ for $[\pm\text{voice}]$ in stratum Z is the sum of the weights of stratal $\text{IDENT}[\text{voi}]Z$ and *general* $\text{IDENT}[\text{voi}]$, and likewise for $[\pm\text{anterior}]$.

With these pieces in place, the only additional proposal needed in order for the model to encode a bias toward core-periphery structure is a requirement that, if any novel stratum is added to the grammar after the ordinary process of language acquisition is complete—as might occur in a situation of novel language contact (a new wave of loanwords), or perhaps in an experimental setting (a nonce-loan nativization experiment)—then ALL FAITHFULNESS CONSTRAINTS INDEXED TO THAT STRATUM ARE ASSIGNED EQUAL WEIGHT, in the absence of overt evidence to the contrary. In terms of the current example, if all FAITHZ constraints for a novel stratum Z are assigned the same weight, then the domination hierarchy $w(\text{IDENT}[\text{voi}]) > w(\text{IDENT}[\text{ant}])$ among the general faithfulness constraints is carried over to the faithfulness relations for stratum Z : $(w(\text{IDENT}[\text{voi}]) + w(\text{IDENT}[\text{voi}]Z)) > (w(\text{IDENT}[\text{ant}]) + w(\text{IDENT}[\text{ant}]Z))$, because $w(\text{IDENT}[\text{voi}]Z) = w(\text{IDENT}[\text{ant}]Z)$. This principle of Uniform Weight by Stratum, along with cumulative constraint interaction and the emergent general-faithfulness hierarchy, thus ensures that faithfulness domination relations remain consistent even in a novel stratum (in the absence of overt evidence to the contrary).

As an example, consider schematic language #1, with a strict core-periphery structure, which was discussed in the learning simulation reported in §4.2. The end-state weights for the markedness constraints NoTI and NoNT, and the general faithfulness constraints $\text{IDENT}[\text{voi}]$ and $\text{IDENT}[\text{ant}]$, are repeated in (19) (from (13)), but now a new stratum Z is added to the language, and its possible effects are considered. By the principle of Uniform Weight by Stratum, all faithfulness constraints indexed to novel stratum Z are assigned the same weight, but three different scenarios can be distinguished (Z_1 – Z_3).

(19) Novel stratum Z introduced for Language #1

<i>category</i>	<i>constraint</i>	<i>weight</i>	<i>cumulative faithfulness</i>	<i>F–M</i>	<i>outcome</i>
markedness	NoTI	28.0	—	—	—
	NoNT	19.0	—	—	—
general faithfulness	$\text{IDENT}[\text{ant}]$	14.4	—	—	—
	$\text{IDENT}[\text{voi}]$	16.2	—	—	—
Stratum Z_1	$\text{IDENT}[\text{ant}]Z_1$	$w < 2.8$	<i>c.f.</i> < 17.2	< 0	/tinta/ Z_1 → [te inda] Z_1
	$\text{IDENT}[\text{voi}]Z_1$		<i>c.f.</i> < 19.0	< 0	
Stratum Z_2	$\text{IDENT}[\text{ant}]Z_2$	$2.8 < w < 13.6$	$17.2 < \textit{c.f.} < 28.0$	< 0	/tinta/ Z_2 → [te inta] Z_2
	$\text{IDENT}[\text{voi}]Z_2$		$19.0 < \textit{c.f.} < 29.8$	> 0	
Stratum Z_3	$\text{IDENT}[\text{ant}]Z_3$	$w > 13.6$	<i>c.f.</i> > 28.0	> 0	/tinta/ Z_3 → [ɹinta] Z_3
	$\text{IDENT}[\text{voi}]Z_3$		<i>c.f.</i> > 29.8	> 0	

If the novel stratum Z is introduced with a weight for all stratal faithfulness constraints that is some value less than 2.8 (stratum ‘ Z_1 ’), then the cumulative weights of both $\text{IDENT}[\text{ant}] + \text{IDENT}[\text{ant}]Z_1$ and $\text{IDENT}[\text{voi}] + \text{IDENT}[\text{voi}]Z_1$ are low enough for the markedness constraints NoTI and NoNT to take priority, leading to nativization of both ‘foreign’ properties: /tinta/ Z_1 → [**te**inda] Z_1 . If the weights for stratum Z faithfulness constraints have some value greater than 13.6 (stratum ‘ Z_3 ’), then the cumulative weights of both $\text{IDENT}[\text{ant}] + \text{IDENT}[\text{ant}]Z_3$ and $\text{IDENT}[\text{voi}] + \text{IDENT}[\text{voi}]Z_3$ are high enough to overcome the antagonistic markedness constraints, leading to preservation of both properties: /tinta/ Z_3 → [ɹinta] Z_3 . Finally, if the stratum Z faithfulness weights are given a value between these two points (stratum ‘ Z_2 ’), then $\text{IDENT}[\text{voi}] + \text{IDENT}[\text{voi}]Z_2$ will overcome NoNT, but $\text{IDENT}[\text{ant}] + \text{IDENT}[\text{ant}]Z_2$ will still be outweighed by NoTI,

leading to nativization of the ‘more-foreign’ /ti/ sequence but preservation of the ‘less-foreign’ /nt/ sequence: /tinta/Z2 → [tɛɪnta]Z2.

As desired, the impossible-nativization mapping /tinta/ → [ɹɪndɪa] is not a possible outcome for stratum Z. The only way for this form to be chosen would be if IDENT[ant]Z had a weight greater than 13.6 while IDENT[voi]Z had a weight less than 2.8. While such weights could be acquired in the presence of overt evidence, this is not a possible state of affairs for a newly introduced or hypothetical stratum under Uniform Weight by Stratum. Thus, the model correctly accounts for the fact that impossible-nativization effects — a dispreference for impossible-nativization mappings—can be productive.

In fact, even a language without a strict core-periphery structure, such as schematic language #3 (§4.4), is predicted to show limited impossible-nativization effects. As discussed above, there is no domination relation between the general faithfulness constraints IDENT[p] and IDENT[voi] in this language, because of the non-superset/subset relation between the Mimetic and Sino-Japanese strata. However, even this language has an end state in which IDENT[ant] has a higher weight than either IDENT[p] or IDENT[voi]. This predicts that speakers of language #3 would in fact find the mapping /tinta/ → [ɹɪndɪa] to be an impossible nativization, since this requires a reversed hierarchy in which faithfulness to [±anterior] takes preference over faithfulness to [±voice]. This prediction is consistent with the fact that some speakers of Japanese, which has only a partial core-periphery structure along the lines of language #3, do show impossible-nativization effects (Ito & Mester 1999; see also Smith & Muratani in prep.).

6 Conclusions and implications

The HG Stratified Faithfulness model successfully allows both core-periphery and non-core-periphery structure to be acquired for a stratified phonology, depending on the patterns in the learning data. In addition, because Harmonic Grammar allows for cumulative constraint interaction, the HG Stratified Faithfulness approach provides a much simpler way to enforce the consistent faithfulness ranking across strata that is required for a grammar with a strict core-periphery structure. A domination hierarchy among the general faithfulness constraints, which emerges automatically during the acquisition of a language with a core-periphery phonology, can be straightforwardly projected to any novel stratum—thereby preserving strict core-periphery structure—by means of the formally simple principle of Uniform Weight by Stratum.

The schematic examples considered in this paper are all quite simple, so it is left to future work to explore interesting questions of increased complexity. For example, what happens to the markedness and general-faithfulness domination hierarchies when the same faithfulness constraint conflicts with multiple markedness constraints at different points in the hierarchy, such as IDENT[voi] in Japanese, which actually conflicts not only with low-priority NoNT, but also with higher-priority NoDD (Ito & Mester 1995ab, 1999)?

More investigation is also needed into just how learners assign forms to strata in the course of language acquisition, but for promising directions to pursue, see Fukazawa et al. (1998), Pater (2005, 2010), and, using weighted constraints and a Maximum Entropy learner to assign forms to strata in English, Hayes (2016). Ito & Mester (1995b: 824, 1999:70) note that, while some of the lexical strata in Japanese are highly cohesive, others show more gradient or fuzzy boundaries; this observation certainly has implications for the acquisition of stratified phonology as well.

The wide-ranging implications and general relevance of these remaining questions go far beyond the phonological analysis of a loanword-rich lexicon (although that topic is interesting in its own right), thereby highlighting the foundational nature of Ito & Mester’s original insight, at the beginning of

constraint-based phonology, relating core-periphery structure and its ‘hierarchy of foreignness’ (Kiparsky 1968: 132) to an explicit hierarchy of markedness constraints in a constraint-based formal grammar.

References

- Beckman, Jill N. 1999. *Positional faithfulness*. New York: Garland.
- Boersma, Paul, & Bruce Hayes. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32: 45–86.
- Boersma, Paul, & Joe Pater. 2016. Convergence properties of a gradual learning algorithm for Harmonic Grammar. In John McCarthy and Joe Pater (eds.), *Harmonic Grammar and Harmonic Serialism*, 389–434. London: Equinox Press.
- Boersma, Paul & David Weenink. 2015. Praat [software], version 5.4.16. Retrieved August 15, 2015 [<http://www.praat.org/>].
- Fukazawa, Haruka. 1997. A Correspondence Theoretic approach to Japanese vocabulary strata. *Proceedings of MALC (Mid-America Linguistics Conference) 1997*.
- Fukazawa, Haruka, Mafuyu Kitahara & Mitsuhiko Ota. 1998. Lexical stratification and ranking invariance in constraint-based grammars. *CLS* 34(2): 47–62.
- Hayes, Bruce. 1999. Phonetically driven phonology: The role of Optimality Theory and inductive grounding. In Michael Darnell et al. (eds.), *Formalism and functionalism in linguistics*, vol. I, 243–285. Amsterdam: Benjamins.
- Hayes, Bruce. 2016. Comparative phonotactics. *CLS* 50.
- Hsu, Brian & Karen Jesney. 2017a. Loanword adaptation in Québec French: Evidence for weighted scalar constraints. *WCCFL* 34: 249–258.
- Hsu, Brian & Karen Jesney. 2017b. Weighted scalar constraints capture the typology of loanword adaptation. *2017 Annual Meeting on Phonology* [<http://dx.doi.org/10.3765/amp.v5i0.4246>].
- Inkelas, Sharon & Cheryl Zoll. 2007. Is grammar dependence real? A comparison between cophological and indexed constraint approaches to morphologically conditioned phonology. *Linguistics* 45: 133–171.
- Irwin, Mark. 2011. *Loanwords in Japanese*. Amsterdam: Benjamins.
- Ito, Junko & Armin Mester. 1995a. The core-periphery structure of the lexicon and constraints on reranking. In Jill Beckman, Laura Walsh Dickey, & Suzanne Urbanczyk (eds.), *Papers in Optimality Theory*, 181–209. Amherst: GLSA.
- Ito, Junko & Armin Mester. 1995b. Japanese phonology. In John Goldsmith (ed.), *The handbook of phonological theory*, 817–838. Cambridge, Mass.: Blackwell.
- Ito, Junko & Armin Mester. 1999. The structure of the phonological lexicon. In Natsuko Tsujimura (ed.), *The handbook of Japanese linguistics*, 62–100. Malden, Mass.: Blackwell.
- Ito, Junko & Armin Mester. 2001. Covert generalizations in Optimality Theory: the role of stratal faithfulness constraints. *Studies in Phonetics, Phonology, and Morphology* 7: 3–33.
- Ito, Junko, & Armin Mester. 2008. Lexical classes in phonology. In Shigeru Miyagawa & Mamoru Saito (eds.), *The Oxford handbook of Japanese linguistics*, 84–106. Oxford: Oxford University Press.
- Jesney, Karen, and Anne-Michelle Tessier. 2011. Biases in Harmonic Grammar: The road to restrictive learning. *NLLT* 29: 251–290.
- Kawahara, Shigeto, Kohei Nishimura, & Hajime Ono. 2003. Unveiling the unmarkedness of Sino-Japanese. In William McClure (ed.), *Japanese/Korean Linguistics 12*, 140–151. Stanford: CSLI Publications.

- Kiparsky, Paul. 1968. How abstract is phonology? In Osamu Fujimura (ed.), *Three dimensions of linguistic theory*, 5-56. Tokyo.
- Legendre, Géraldine, Yoshiro Miyata, & Paul Smolensky. 1990. Can connectionism contribute to syntax?: Harmonic grammar, with an application. Report CU-CS-485-90. Computer Science Department, University of Colorado at Boulder.
- McCarthy, John & Alan Prince. 1995. Faithfulness and reduplicative identity. In Jill Beckman, Laura Walsh Dickey, & Suzanne Urbanczyk (eds.), *Papers in Optimality Theory*, 250–384. Amherst: GLSA.
- Pater, Joe. 2001. Austronesian nasal substitution revisited. In Linda Lombardi (ed.), *Segmental phonology in Optimality Theory: Constraints and representations*, 159–182. Cambridge: CUP.
- Pater, Joe. 2005. Learning a stratified grammar. *BUCLD 29*: 482–492.
- Pater, Joe. 2009. Weighted constraints in generative linguistics. *Cognitive Science* 33: 999–1035.
- 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.
- Pater, Joe. 2016. Universal Grammar with weighted constraints. In John McCarthy & Joe Pater (eds.), *Harmonic Grammar and Harmonic Serialism*, 1–46. London: Equinox Press.
- Pinta, Justin. 2013. *Lexical strata in loanword phonology: Spanish loans in Guarani*. MA thesis, UNC-CH.
- Prince, Alan & Paul Smolensky. 1993. *Optimality Theory: Constraint interaction in generative grammar*. Ms., Rutgers U. and U. of Colorado, Boulder. [Malden, Mass.: Blackwell, 2004.]
- Smith, Jennifer L., and Yuka Muratani. In preparation. Nonce-loan judgments and impossible-nativization effects in Japanese. Ms., UNC-CH.
- Smith, Jennifer L., & Justin Pinta. 2017. Aggressive core-periphery phonology in Guarani. Ms., UNC-CH and The Ohio State University.
- Smolensky, Paul. 1996. The initial state and ‘richness of the base’ in Optimality Theory. Technical Report JHU-CogSci-96-4. Baltimore: The Johns Hopkins University.