

Taking a Free Ride in Morphophonemic Learning*

John J. McCarthy

Department of Linguistics
University of Massachusetts Amherst
Amherst, MA 01003 USA

Abstract

As language learners begin to analyze morphologically complex words, they face the problem of projecting underlying representations from the morphophonemic alternations that they observe. Research on learnability in Optimality Theory has started to address this problem, and this article deals with one aspect of it. When alternation data tell the learner that some surface [B]s are derived from underlying /A/s, the learner must sometimes generalize by deriving *all* [B]s, even nonalternating ones, from /A/s. An adequate learning theory must therefore incorporate a procedure that allows nonalternating [B]s to take a “free ride” on the /A/ → [B] unfaithful map.

Keywords: allophony, chain shift, learnability, morphophonemics, opacity, Optimality Theory.

*Come on and take a free ride (free ride)
(Winter 1972)*

1. Introduction

If language learning is the central problem of linguistic research, then it is important to show that learners can in principle acquire the grammar of the ambient language from the limited data available to them. Formal studies of language learnability are concerned with determining, for a particular linguistic theory, whether and under what conditions learning can succeed. Ideally, any proposed linguistic theory will be accompanied by an explicit learning model from which learnability results can be deduced.

Optimality Theory (Prince & Smolensky 2004) has been coupled with a learning model almost from the beginning. There is by now a very substantial body of work discussing the constraint demotion learning algorithm of Tesar (1995) and Tesar & Smolensky (1994, 1998, 2000), as well as alternatives. Although much of this literature deals with phonotactic learning, it has begun to address morphophonemics: how do learners use data from phonological alternations to work out a grammar and a lexicon of underlying representations? (See, e.g., Dinnsen & McGarrity 2004, Ota 2004, Pater 2004, Tesar et al. 2003, Tesar & Prince 2004, Tesar & Smolensky 2000: 77-84, Tessier 2004.) Morphophonemic learning presents problems that are not met with in phonotactic learning. This

*. Special thanks to participants in the UMass summer phonetics/phonology group (SPe/oG) for their comments and suggestions. For additional comments, I'm grateful to John Alderete, Shigeto Kawahara, Joe Pater, and particularly Alan Prince.

paper will describe one such problem and propose a solution to it.

The OT learnability literature standardly assumes that phonotactic learners posit identity maps from the lexicon to observed surface forms:¹ the learner hears [A], takes /A/ to be its underlying source, and seeks a grammar that will perform the identity map /A/ → [A]. If the learner never hears *[B], however, and this gap is nonaccidental, then his (or her) grammar must not map anything to *[B]. In particular, the /B/ → [B] identity map must not be sanctioned by the grammar. Instead, if the grammar is challenged with a /B/ input, it must treat it unfaithfully, mapping it to zero or to something that is possible in the target language.

Morphophonemic learning is different. As learners begin to analyze morphologically complex words, they discover morphophonemic alternations for which the identity map is insufficient.² For example, once a learner becomes aware of paradigmatic relationships like Egyptian Arabic *fihim/fihmu* ‘he/they understood’ and sets up underlying /fihim/, he is then committed to the unfaithful map /fihim-u/ → *fihmu*. But what happens with nonalternating forms? Is the identity map still favored?

In this article, I argue that nonalternating forms are sometimes derived by unfaithful maps as well. As I will show, there are examples with the following properties. Data from alternations show that *some* surface [B]s derive from underlying /A/s. Other evidence, some of it distributional and some involving opacity, argues that *all* surface [B]s, even the nonalternating ones, derive from underlying /A/s. I will propose a learning principle according to which learners who have discovered the /A/ → [B] unfaithful map from alternations will attempt to generalize it, projecting /A/ inputs for all surface [B]s, whether they alternate or not. In other words, the nonalternating [B]s attempt to take a “free ride” on the independently motivated /A/ → [B] map.³ The learner accepts this tentative hypothesis about underlying representations if it yields a grammar that is consistent (in the sense of Tesar 1997) and that is more restrictive (in the sense of Prince & Tesar 2004) than the grammar would be without this hypothesis. These two requirements, consistency and greater restrictiveness, are arguably sufficient to rule out inappropriate generalization of unfaithful maps.

Section 2 and the appendix of this paper document the phenomenon: there are languages where alternations show that *some* [B]s derive from /A/s and other evidence shows that *all* [B]s must derive from /A/s. Section 3 presents and applies the proposed learning algorithm. Section 4 describes a limitation of this algorithm and proposes a modification to address this limitation.

2. Exemplification and explanation of the issue

This section describes a single example, coalescence in Sanskrit. To establish the generality of the problem, however, this article also includes an appendix with several additional examples.

In Sanskrit, alternations show that some surface long mid vowels [e:] and [o:] are derived by coalescence from /ai/ and /au/ sequences, respectively.

-
1. I will not attempt to list the many works that assume the identity map in phonotactic learning, nor will I attempt to identify the origin of this idea.
 2. On the insufficiency of the identity map even in phonotactic learning, see Alderete & Tesar (2002).
 3. I’ve borrowed the term “free ride” from Zwicky (1970). The situation Zwicky has in mind is different from what I’m talking about: an analyst, instead of positing a rule /A/ → [C], instead posits a rule /A/ → [B] that is ordered before an independently motivated rule /B/ → [C]. The idea is that the output of one rule takes a free ride on the other.

(1) Sanskrit coalescence (de Haas 1988, Gnanadesikan 1997, Schane 1987, Whitney 1889)

/tava indra/	tave:ndra	'for you, Indra (voc.)'
/hita upadaijah/	hito:pade:jah	'friendly advice'

Coalescence occurs in both internal and external sandhi, and it is abundantly supported by alternations. The result of coalescence is a compromise in the height of the two input vowels. The result is long because the moras associated with the input vowels are conserved in the output. These are typical properties of coalescence cross-linguistically (de Haas 1988).

Since all surface mid vowels are long in Sanskrit, all are at least potentially the result of coalescence. For many words with mid vowels, there are alternations that support a coalescent source. But it is unlikely that relevant alternation data will exist for all mid vowels, especially in the limited experience of the language learner. When presented with an [e:] and no relevant alternations, what does the learner do? (For expositional clarity, I'll only talk about [e:] until late in section 3, when I will look at [o:] as well.) A persistent bias toward the identity map would favor positing underlying /e:/ for nonalternating [e:]. There is a good reason, however, to think that learners actually adopt a free-ride strategy, deriving nonalternating [e:] from /ai/: by deriving *all* [e:]s from /ai/s, we explain why Sanskrit has no short [e]s. If all mid vowels are derived by coalescence and if the product of vowel coalescence is always a long vowel, then there is no source for short [e].

The traditional approach to cases like this involves positing a morpheme structure constraint (MSC) that prohibits underlying mid vowels, short or long. Classic MSC's are not durable; they do not affect the output of rules. The rule of coalescence is therefore free to create (long) mid vowels. Short mid vowels are not observed in surface forms because they are barred from the lexicon by the MSC and no rule creates them. On this view, nonalternating [e:] has to be derived by coalescence because /e:/ is ruled out by the MSC.⁴ This type of coalescence is said to be non-structure-preserving (de Haas 1988), since its output is something that is not allowed in the lexicon.

In OT, richness of the base (ROTB) forbids MSCs (Prince & Smolensky 2004), but an abstractly similar analysis is possible. Gnanadesikan (1997: 139-153) analyzes Sanskrit with a chain shift: underlying /ai/ maps to [e:], while underlying /e(:)/ maps to [i(:)] or [a(:)] — it doesn't matter which. Gnanadesikan assumes a constraint set similar to (2).

(2) a. IDENT(Vowel Height) (hereafter ID(VH))

For every pair (V, V'), where V is an input vowel and V' is its output correspondent, if V and V' differ in height, assign a violation mark.

b. IDENT-ADJ(Vowel Height) (hereafter ID-ADJ(VH))

For every pair (V, V'), where V is an input vowel and V' is its output correspondent, if V and V' differ by more than one degree of height, assign a violation mark.

c. *MID

Mid vowels are prohibited.

4. The Alternation Condition of Kiparsky (1973) was specifically formulated so as not to exclude such an analysis. The Alternation Condition says, approximately, that no *neutralization* rule can apply to all instances of a morpheme. But coalescence in Sanskrit is not a neutralization rule if all [e:]s are derived from /ai/.

d. *DIPH

Diphthongs are prohibited.

e. UNIFORMITY (hereafter UNIF) (McCarthy & Prince 1995, 1999)

No output segment has multiple correspondents in the input (\approx no coalescence).

For present purposes, I will assume that this is the entire constraint set and that the only unfaithful maps that are possible are coalescence, which violates UNIF, and changes in vowel height, which violate the ID constraints.

Because mid vowels are prohibited in general, *MID must dominate ID(VH), so any input /e/ or /e:/ drawn from the rich base will map unfaithfully to a high or low vowel:

(3) *MID \gg ID(VH)

/be(:)/	*MID	ID(VH)
\rightarrow bi(:) or ba(:)		1
\sim be(:)	₁ W	L

In this comparative tableau (Prince 2002), the winning candidate's violation marks are displayed as integers in the first row below the constraints. Subsequent rows compare the winner with various losers, which are preceded by a tilde. The W indicates that *MID favors the winner over this loser; the L says that ID(VH) favors the loser over the winner. (The number of violation marks incurred by a loser is also shown in small print.) This tableau is telling us that winner-favoring *MID dominates loser-favoring ID(VH). Generally, in a properly ranked tableau, every L is dominated by some W (by the Cancellation-Domination Lemma of Prince & Smolensky 2004).

Coalescence of diphthongs requires a ranking where *DIPH dominates both UNIF and ID(VH):

(4) *DIPH \gg UNIF, ID(VH)

/va ₁ i ₂ /	*DIPH	UNIF	ID(VH)
\rightarrow ve _{1,2}		1	2
\sim va ₁ i ₂	₁ W	L	L

The winning candidate in (4) incurs two marks from ID(VH) because each of its input vowels stands in correspondence with an output vowel of a different height.

For Gnanadesikan, the key to analyzing Sanskrit is the constraint ID-ADJ(VH), since it blocks the *MID-satisfying map of /ai/ to [i:] or [a:] that would otherwise be expected. The idea is that fusing /ai/ into a vowel that is either high or low establishes correspondence between one of its input vowels and an output vowel that is two steps away from it in height. ID-ADJ(VH) specifically excludes this map even though it offers satisfaction of *MID:

(5) ID-ADJ(VH) >> *MID

/va ₁ -i ₂ /	ID-ADJ(VH)	*MID
→ ve _{1,2}		1
~ vi _{1,2} or va _{1,2}	1 W	L

As the subscripts indicate, the losing candidates in (5) are products of segmental coalescence, not segmental deletion. A segment derived by coalescence owes its allegiance, IDENT-wise, to both of its input parents. For example, [i_{1,2}] violates ID-ADJ(VH) because one of its input correspondents, /a₁/, differs from it by two degrees of height.

This analysis is a typical chain shift.⁵ Underlying /ai/ becomes [e₁], whereas underlying /e:/ and /e/ become something else, either a high vowel or a low one (it doesn't matter which). The constraint ID-ADJ(VH) blocks the fell swoop map in which /ai/ coalesces to form a high or low vowel. In any chain shift /A/ → [B] and /B/ → [C], there must be a faithfulness constraint that is violated by the forbidden /A/ → [C] map but not by the /B/ → [C] map (Moreton 2000). ID-ADJ(VH) is such a constraint.

No matter how Sanskrit is analyzed — with a MSC against mid vowels or the ranking [*MID >> ID(VH)] in (3) — it presents a learning problem. Learners who are biased toward the identity map will assume that nonalternating mid vowels are derived from underlying mid vowels. How are they able to override this natural bias and learn a grammar that excludes underlying mid vowels with a MSC or, what amounts to almost the same thing, maps all of them unfaithfully? To see this problem in more concrete terms, we will look at how learning proceeds in OT grammar. (There are no extant proposals for the learning of MSCs.)

The learner is armed with the learning procedure known as biased multirecursive constraint demotion (BCD) (Prince & Tesar 2004). In multirecursive constraint demotion (Tesar 1997), learning proceeds on the basis of a comparative tableau like those given above, though with multiple inputs considered simultaneously. (See Prince (2002) on comparative tableaux in constraint demotion.) Constraints are ranked one at a time starting from the top of the hierarchy. A constraint C is rankable if it favors no losers among the as-yet-unaccounted-for candidates. (A losing candidate has been accounted for if it is dispreferred relative to the winner by some constraint that has already been ranked.) Once some constraint has been ranked, the tableau with this new ranking is then submitted to the same procedure, recursively. The procedure terminates when all constraints have been ranked or when there are irreducible conflicts (as in (16) below), in which case no grammar can be found.

The bias in BCD favors grammars that are more restrictive (Prince & Tesar 2004) (also see Hayes 2004, Ito & Mester 1999). The bias says which constraints to rank first (=higher) when there are two or more constraints that favor no losers. By preference, markedness constraints are ranked higher, since that favors greater restrictiveness. But if there are no rankable markedness constraints, obviously a faithfulness constraint must be ranked. The faithfulness constraint that's chosen is one that will make a markedness constraint rankable. It does that by accounting for losing candidates that the markedness constraint favors. In practice, this means that the faithfulness constraint chosen for ranking has Ws wherever some markedness constraint has Ls. (This is a deliberate simplification; see

5. On chain shifts in acquisition, see Dinnsen (2004), Dinnsen and Barlow (1998), and Dinnsen, O'Connor, and Gierut (2001).

Prince & Tesar for further details that are not important in the current context.)

At the earliest stages of learning, there is little or no awareness of morphological structure and no awareness of alternations, so only the phonotactics are being learned (on this assumption, see Hayes 2004, Pater & Tessier 2003, Prince & Tesar 2004). The phonotactic learner seeks a grammar that performs identity maps – and *only* identity maps– from the inferred lexicon to the observed surface forms. In the case of Sanskrit, the ambient language contains words with the vowels [i], [i:], [a], [a:], and [e:]. A tableau is given in (6); all constraints are assumed to be unranked, though the algorithm works equally well if they have some ranking imposed on them already.

(6) Sanskrit during phonotactic learning

lexicon	cands.	*MID	*DIPH	ID(VH)	ID-ADJ(VH)	UNIF
/ba/ (similarly /ba:z/)	→ ba					
	~ bi			₁ W	₁ W	
	~ be	₁ W		₁ W		
/bi/ (similarly /bi:z/)	→ bi					
	~ ba			₁ W	₁ W	
	~ be	₁ W		₁ W		
/be:z/	→ be:	1				
	~ bi: or ba:	L		₁ W		

I will refer to this as a “support” tableau, after Tesar and Prince (2004), since it provides the support for the rankings found by BCD. In this tableau, the initially rankable constraints (i.e., the constraints that favor no losers) are *DIPH, ID(VH), ID-ADJ(VH), and UNIF. The bias in BCD favors ranking markedness constraints first, and of these only *DIPH is a markedness constraint, so it goes in the top rank of the constraint hierarchy. But because *DIPH favors no winners, ranking it does not account for any losing candidates.

On the next recursive pass through BCD, no rankable markedness constraints can be found; *MID is the only remaining markedness constraint, and it favors a loser from the input /be:z/. Therefore, BCD must rank a faithfulness constraint, and it selects the faithfulness constraint ID(VH) (see (7)). BCD chooses ID(VH) because it favors the winner in the one row where *MID favors the loser, so ranking ID(VH) will allow *MID to be ranked on the next recursion of BCD.

(7) Support tableau (6) after second ranking pass

lexicon	cands.	*DIPH	ID(VH)	*MID	ID-ADJ(VH)	UNIF
/ba/ (similarly /baː/)	→ ba					
	~ bi		₁ W		₁ W	
	~ be		₁ W	₁ W		
/bi/ (similarly /biː/)	→ bi					
	~ ba		₁ W		₁ W	
	~ be		₁ W	₁ W		
/beː/ (similarly /beː/)	→ beː			1		
	~ biː or baː		₁ W	L		

All of the loser rows have been shaded to indicate that these candidates have been accounted for by ranking ID(VH). Since no unaccounted-for candidates remain, *MID is obviously free to be ranked. Because of the markedness-high bias in BCD, the remaining faithfulness constraints go below *MID, yielding the result in (8).

(8) Sanskrit ranking after phonotactic learning

*DIPH >> ID(VH) >> *MID >> ID-ADJ(VH), UNIF

The ranking (8) is consistent with the observed vowel system of Sanskrit, but it is also consistent with a superset of Sanskrit in which short mid vowels are permitted. If this grammar is given the input /be/, it yields the output *[be]. This grammar, then, is insufficiently restrictive, allowing things that Sanskrit doesn't allow. This failure to learn the more restrictive grammar is an instance of the well-known Subset Problem (Angluin 1980, Baker 1979): learners presented with only positive evidence cannot proceed from a less restrictive grammar to a more restrictive one. The bias in BCD is intended to address this problem by ensuring that learners always favor the most restrictive grammar that is consistent with the primary data, but BCD is insufficient in this case (see also Alderete & Tesar 2002). The goal of this article is to explain how learners manage to find the more restrictive grammar anyway.

Another way to look at the problem with (8) is that Sanskrit is harmonically incomplete relative to the markedness constraints provided: it lacks short mid vowels, yet there is no markedness constraint against them. As Gnanadesikan's analysis shows, Sanskrit's harmonically incomplete system can be derived by judiciously deploying the faithfulness constraint ID-ADJ(VH), but the phonotactic learner has as yet no evidence from alternations that show mid vowels being derived from underlying diphthongs. Worse yet, even with that evidence, the morphophonemic learner also fails to acquire the more restrictive grammar, as we will see shortly.

When an unlearnable system is harmonically incomplete, there is a straightforward solution: enlarge the set of markedness constraints.⁶ If a constraint against *short* mid vowels is added, then BCD will find a grammar that is no less restrictive than Sanskrit requires. Although this constraint is not implausible, in general this is not the right way to approach the problem. Adding a constraint

6. I'm grateful to John Kingston for raising this issue.

against short mid vowels merely stipulates that which is explained in both the traditional approach with its MSC and Gnanadesikan's analysis with its chain shift: adding a constraint gives up on the idea that there is a connection between the independently motivated coalescence process and the skewed inventory of mid vowels. More seriously, this approach does not generalize to other examples; see the appendix of this article and McCarthy (2004) for further discussion. So even if this approach were right for Sanskrit, it would not help with the larger learning problem that Sanskrit is merely one example of. Therefore, I set this possibility aside and will not consider it again.

As language development proceeds, the Sanskrit learner begins to attend to morphological structure and this leads, by a process I will not be discussing (though see the references in section 1), to the discovery of unfaithful maps between input and output. Let us say, for instance, that the learner has realized that the isolation forms *tava* and *indra* are components of the phrase *tave:ndra*. The learner adds this new information about /ai/ → [e:] unfaithful maps to the support tableau, yielding (9).

(9) Sanskrit support tableau after some morphological analysis

lexicon	cands.	*DIPH	ID(VH)	*MID	ID-ADJ(VH)	UNIF
/ba/ (similarly /ba:z/)	→ ba					
	~ bi		₁ W		₁ W	
	~ be		₁ W	₁ W		
/bi/ (similarly /bi:z/)	→ bi					
	~ ba		₁ W		₁ W	
	~ be		₁ W	₁ W		
/be:z/ (similarly /be:z/)	→ be:			1		
	~ bi: or ba:		₁ W	L		
/va ₁ +i ₂ / (similarly /va ₁ +i ₂ /)	→ ve: _{1,2}		2	1		1
	~ va ₁ i ₂	₁ W	L	L		L
	~ vi: _{1,2} or va: _{1,2}		₁ L	L	₁ W	

Three comments on this tableau:

- (i) Recall that I am making the simplifying assumption that there is coalescence but no deletion (because MAX >> UNIF). The correspondence relations in the /va₁+i₂/ rows reflect this: the vowels in all candidates bear the indices of their input correspondents /a₁/ and /i₂/. ID(VH) assigns a mark for every pair of corresponding segments that differ in height. The candidate [ve:_{1,2}] has two such pairs, (a₁, e:_{1,2}) and (i₂, e:_{1,2}), whereas [vi:_{1,2}] has only the pair (a₁, i:_{1,2}). Unlike [ve:_{1,2}], however, [vi:_{1,2}] violates ID-ADJ(VH) because the pair (a₁, i:_{1,2}) contains corresponding vowels that differ by two degrees of height and not just one (see (5)).
- (ii) This tableau still contains the /be:z/ → [be:] identity map — it must, because there are forms that never alternate and/or forms whose alternations have eluded the learner's experience or attention. This lingering identity map is the crux of the learning problem: because it is there in the tableau, BCD continues to find a grammar that is insufficiently restrictive.

(iii) This tableau preserves the ranking (8) that was acquired during phonotactic learning. With the added information about $/va_1+i_2/ \rightarrow [ve_{1,2}]$, this ranking is now incorrect — the last row shows that loser-favoring constraints dominate the highest-ranking winner-favoring constraint. Because of this discrepancy between the ranking and the result, a new round of BCD is required.

On the first pass of (9) through BCD, the only markedness constraint that favors no losers is again *DIPH, so it is ranked at the top of the hierarchy:

(10) Support tableau (9) after first ranking pass. Ranking: [$*DIPH \gg rest$]

lexicon	cands.	*DIPH	*MID	ID(VH)	ID-ADJ(VH)	UNIF
/ba/ (similarly /baʔ/)	→ ba					
	~ bi			₁ W	₁ W	
	~ be		₁ W	₁ W		
/bi/ (similarly /biʔ/)	→ bi					
	~ ba			₁ W	₁ W	
	~ be		₁ W	₁ W		
/be:/	→ be:		1			
	~ bi: or ba:		L	₁ W		
/va ₁ +i ₂ /	→ ve _{1,2}		1	2		1
	~ va ₁ i ₂	₁ W	L	L		L
	~ vi _{1,2} or vaʔ _{1,2}		L	₁ L	₁ W	

Only one markedness constraint remains, *MID, and it is not yet rankable, since it favors some losers. Therefore the only choice is to rank faithfulness constraints. Two are rankable, ID-ADJ(VH) and UNIF. Ranking UNIF will not free up *MID for ranking, since UNIF favors no winners in this tableau. Therefore, ID-ADJ(VH) is ranked on this pass, yielding (11). (ID-ADJ(VH) and ID(VH) together constitute an “F-gang” (Prince & Tesar 2004: 266) that must be ranked before *MID is free for ranking.)

(11) Support tableau (10) after second ranking pass. Ranking: [$*\text{DIPH} \gg \text{ID-ADJ(VH)} \gg \text{rest}$]

lexicon	cands.	*DIPH	ID-ADJ(VH)	*MID	ID(VH)	UNIF
/ba/ (similarly /ba:ʔ/)	→ ba					
	~ bi		₁ W		₁ W	
	~ be			₁ W	₁ W	
/bi/ (similarly /bi:ʔ/)	→ bi					
	~ ba		₁ W		₁ W	
	~ be			₁ W	₁ W	
/be:ʔ/ (similarly /be:ʔ/)	→ be:ʔ			1		
	~ bi:ʔ or ba:ʔ			L	₁ W	
/va ₁ i ₂ / (similarly /va ₁ i ₂ :ʔ/)	→ ve:ʔ _{1,2}}			1	2	1
	~ va ₁ i ₂	₁ W		L	L	L
	~ vi:ʔ _{1,2} or va:ʔ_{1,2}}}		₁ W	L	₁ L	

*MID, still favors one loser, so it is not yet rankable. Both of the remaining faithfulness constraints, ID(VH) and UNIF, are rankable. The bias in BCD favors choosing the one that will allow *MID to be ranked on the next round, and that is ID(VH). Once ID(VH) has been ranked, *MID is ranked next, and UNIF goes at the bottom of the hierarchy. The final result of BCD is given in (12).

(12) Sanskrit ranking after morphophonemic learning

$*\text{DIPH} \gg \text{ID-ADJ(VH)} \gg \text{ID(VH)} \gg *MID \gg \text{UNIF}$

The problem with (12) is that it is insufficiently restrictive. Like the ranking in (8) that emerged from phonotactic learning, the ranking in (12) produces surface short mid vowels if presented with input /e/ from the rich base. Gnanadesikan's chain-shift analysis, which derives all surface mid vowels from diphthongs and maps all underlying mid vowels onto high or low vowels, cannot be learned using just BCD and this constraint set. It is unlearnable because the lingering presence of the /be:ʔ/ → [be:] identity map in the support tableau forces the ranking ID(VH) >> *MID — exactly the ranking that makes this the grammar of the superset language. This result confirms what Alderete & Tesar (2002) have argued on independent grounds: there are circumstances where a commitment to the identity map can lead to a superset language. An adequate learning theory needs to explain how learners get to the more restrictive grammar that Sanskrit requires.

3. The Free-Ride Learning Algorithm

This section presents the proposal at three different levels of detail. The first is a very brief overview; the second is a detailed textual presentation interspersed with worked examples; and the third is a formal statement of the algorithm in (24).

The morphophonemic learning problem in Sanskrit can be described as follows. Learners can infer from alternations that *some* [e:]s are derived from /ai/ sequences. To learn a sufficiently restrictive grammar, however, learners must generalize from alternating [e:]s to nonalternating [e:]s,

eliminating the /e:/ → [e:] identity map from the support tableau and instead deriving *all* [e:]s from /ai/ sequences, even in forms with no alternations.

The essence of my proposal is that learners, whenever alternations lead them to discover a new unfaithful map, always attempt to generalize that map across the entire support tableau (and thereby across the entire language). This is the free ride: nonalternating [e:]s take a free ride on the /ai/ → [e:] map that was discovered from alternations. With the /e:/ → [e:] identity map out of the way, BCD can be a more restrictive grammar. If the identity map bespeaks economy of derivation, the free ride bespeaks profligacy, but profligacy that leads to greater restrictiveness where it really matters, in the grammar. (On the trickier problem of contextually restricted free rides (e.g., all *word-final* vowels are underlying long because some are), see section 4.)

Sometimes the free ride is a trip to nowhere. There are two situations where generalizing in this way does not work out. If other, similar identity maps remain in the support tableau after taking the free ride, then BCD may yield a grammar that is unchanged, so the exercise was pointless. Or if the generalized unfaithful map was authentically neutralizing, then BCD will fail to find any grammar. (That is, BCD will be unable to rank all constraints.) In either of these situations, the free-ride generalization is rejected.

Some formal development is required to put these ideas into practice. We require a definition of an unfaithful map, a statement of how to alter the support tableau in order to take the free ride, and a way of deciding whether or not the free ride is appropriate.

The notion “unfaithful map” appears in some of my previous work on opacity (McCarthy 2003a, b). In correspondence theory (McCarthy & Prince 1995, 1999), every candidate output form specifies a correspondence relation \mathfrak{R} between itself and the input. (Thus, a candidate is really an ordered 3-tuple (input, output, $\mathfrak{R}_{(\text{input}, \text{output})}$.) Like any binary relation, \mathfrak{R} is a set of ordered pairs; it consists of a pairing of one or more input segments or zero with one or more output segments or zero.⁷ For example, if input /b₁a₂t₃/ is associated with output [p₁a₂], the indices are shorthand for the correspondence relation $\mathfrak{R}_{(/bat/, [pa])} = \{(b_1, p_1), (a_2, a_2), (t_3, \emptyset)\}$. Each ordered pair is a map between input and output. Some ordered pairs, such as (b₁, p₁) and (t₃, \emptyset), produce faithfulness violations, and others, such as (a₂, a₂), will not. An unfaithful map is one of the former: a member of a correspondence relation that is the source of a faithfulness violation.

In morphophemic learning, alternations among morphologically related forms lead the learner to posit underlying representations that require unfaithful maps. For example, the Sanskrit learner who has concluded that /tava indra/ underlies *tave:ndra* will couple this form with a correspondence relation that includes the map (a₃i₄, e:_{3,4}). This map is unfaithful because it violates ID(VH) and UNIF.

I assume that learners maintain a set of unfaithful maps \mathfrak{L} , which they have learned from analyzing paradigmatic alternation data.⁸ Whenever the lexicon column of the support tableau changes, \mathfrak{L} is updated, and if it changes, the free-ride learning algorithm (hereafter, FRLA) is invoked.

7. Here and throughout, I proceed under the assumption that correspondence is a relation between segments (or zero). The correspondence relation can obviously be extended to features, tones, moras, or other elements of phonological representation.

8. There is a systematic ambiguity in the term “unfaithful map”. In a correspondence relation, an unfaithful (or faithful) map is a pairing of specific segments in the input and output strings. In \mathfrak{L} , however, an unfaithful map is a pairing of segment types, lifted out of the particulars of the word that instantiates that map. Typographically, unfaithful maps of the former type are supplied with segmental subscripts, whereas those of the latter type have no subscripts. In actual practice, there is little danger of mixing up the two senses of the term.

FRLA proceeds as follows (see (24) for a more explicit statement). Assume that the learner has a support tableau τ_o derived from previous learning and a list \mathcal{Q} of (alternation-supported) unfaithful maps in τ_o . Start with some member of \mathcal{Q} , the unfaithful map (A, B). Make a copy of τ_o called τ_n and locate all (B, B) identity maps in τ_n . Perform “surgery” (Tesar et al. 2003) on τ_n , changing the underlying form of any morpheme that undergoes the (B, B) map, replacing it with a (A, B) map, and updating the faithfulness assessments accordingly.

Now τ_n has the same outputs as τ_o but with a different lexicon. Submit τ_n to BCD. If BCD fails to find a grammar for τ_n , discard it and go back to τ_o . If BCD finds a grammar for τ_n , compare it with the grammar for τ_o . Of these two choices, the learner adopts the one that leads to a more restrictive grammar. In other words, the learner chooses the lexicon that allows for a more restrictive grammar.

This process continues with the other unfaithful maps in \mathcal{Q} and with all of the subsets of the unfaithful maps in \mathcal{Q} . If one or more free rides are warranted, then the support tableau at the end of the process will have different underlying representations than at the beginning of the process, and the grammar will be more restrictive.

BCD’s ability to detect inconsistency plays a key role in this process. Recall that MRCD, upon which BCD is based, will fail to find a ranking when there are two or more constraints left to rank and all favor at least one loser. If no grammar can be found under the specific assumptions about underlying representation that τ_n embodies, then τ_n must have taken the free ride too far; it has eliminated at least one ineliminable identity map. By using BCD in this way, I am directly following Tesar et al. (2003) and Tesar & Prince (2004), who use inconsistency-detection to reject hypothesized underlying representations for alternating morphemes, and less directly following Kager (1999: 334), who invokes inconsistency-detection to trigger modifications in the lexicon.

If BCD succeeds in finding a ranking for τ_n , a decision must be made: which hypothesis about the lexicon is superior, the one embodied by τ_n or the one embodied by τ_o ? As the discussion of Sanskrit showed, we seek the lexicon that allows for a more restrictive grammar, thereby avoiding the Subset Problem. The choice between τ_n and τ_o therefore requires a metric for grammar restrictiveness, and one exists in the form of the *r-measure* of Prince & Tesar (2004). It is defined in (13).⁹

(13) R-measure (Prince & Tesar 2004:252)

The *r-measure* for a constraint hierarchy is determined by adding, for each faithfulness constraint in the hierarchy, the number of markedness constraints that dominate that faithfulness constraint.

Grammars with higher *r-measures* are in general more restrictive. Prince & Tesar note that grammars with equal *r-measures* may also differ in restrictiveness, but this as-yet imperfect metric appears to be sufficient for present purposes. In short, τ_n bests τ_o if τ_n supports a grammar with a higher *r-measure*.

The following examples briefly illustrate the FRLA. They are simplified, with artificially restricted support tableaux and a single unfaithful map. Three outcomes are considered: (i) giving a free ride to an identity map produces a support tableau for which no grammar can be found (devoicing in German); (ii) the free ride leads to a grammar that is the same as the one without the

9. I’m grateful to Joe Pater for suggesting this application of the *r-measure*.

free ride (hypothetical Sanskrit' with short mid vowels); and (iii) the free ride leads to a more restrictive grammar (real Sanskrit).

(i) No grammar is found. If taking a free ride leads to a wrong surmise about underlying representations, no grammar will be found, and the offending support tableau is rejected.

This will happen when a process of neutralization is the source of the unfaithful map. Imagine, for example, a learner of German whose first morphophonemic discovery is that /ra:d/ is the underlying form of [ra:t] 'wheel', so $\mathcal{L} = \{(d, t)\}$. This learner has the support tableau in (14) prior to applying the FRLA:

(14) German morphophonemic learning – τ_0

lexicon	cands.	*VCDOBST	*VCDOBSTCODA	ID(voice)
/daŋk/ 'thanks'	→ daŋk	1		
	~ taŋk	L		₁ W
/ta:t/ 'deed'	→ ta:t			
	~ da:d	₂ W	₁ W	₂ W
	~ da:t	₁ W		₁ W
/ra:d/	→ ra:t			1
	~ ra:d	₁ W	₁ W	L

Applying BCD to (14) yields the ranking [$*VCDOBSTCODA \gg ID(voice) \gg *VCDOBST$].

The FRLA makes a copy of (14) in which the lexicon is altered: all (t, t) identity maps are replaced by the (d, t) unfaithful map in \mathcal{L} . The support tableau with this free ride is given in (15).

(15) Free ride on (d, t) in German – τ_n

lexicon	cands.	*VCDOBST	*VCDOBSTCODA	ID(voice)
/daŋk/	→ daŋk	1		
	~ taŋk	L		₁ W
/da:d/	→ ta:t			2
	~ da:d	₂ W	₁ W	L
	~ da:t	₁ W		₁ L
/ra:d/	→ ra:t			1
	~ ra:d	₁ W	₁ W	L

The difference between (14) and (15) is that (14) derives [ta:t] by identity maps from /ta:t/, while (15) derives it from /da:d/ by taking a free ride on the /d/ → [t] unfaithful map. Obviously, (15) is not the outcome we seek: because /d/ → [t] is neutralizing, it's a grievous error to attempt to derive all surface [t]s from underlying /d/s. The FRLA creates (15) as a hypothesis about the German lexicon, but this hypothesis is soon rejected when it turns out that no constraint ranking is consistent with it.

The FRLA submits (15) to BCD. On the first pass, only *VCDOBSTCODA is rankable, so it goes at the top of the constraint hierarchy, thereby disposing of two losing candidates, as shown in (16).

(16) Support tableau (15) after first ranking pass

lexicon	cands.	*VCDOBSTCODA	*VCDOBST	ID(voice)
/daŋk/	→ daŋk		1	
	~ taŋk		L	₁ W
/da:d/	→ ta:t			2
	~ da:d	₁ W	₂ W	L
	~ da:t		₁ W	₁ L
/ra:d/	→ ra:t			1
	~ ra:d	₁ W	₁ W	L

What does BCD do next? Neither *VCDOBST nor ID(voice) is rankable, since each favors at least one loser. BCD has failed to find a grammar because *VCDOBST and ID(voice) are in an irreducible conflict under these specific assumptions about the lexicon. This failure is unsurprising: devoicing is a neutralization process in German, since the [d]/[t] contrast is maintained in onset position, so there is no hope of deriving *all* [t]s from /d/s. Because BCD fails to find a grammar, this new support tableau is discarded, as is the hypothesis it embodies about the German lexicon. Instead, the original support tableau (14) is correctly retained by the learner.¹⁰

Taking a free ride on a neutralization process is a bad choice. The FRLA requires learners to consider this choice, but through BCD it also supplies a way of rapidly detecting the error and recovering from it.

(ii) A grammar is found, but the ranking is unchanged. In this situation, nothing is gained in restrictiveness by taking the free ride. The FRLA does not restructure the lexicon when no improvement in grammar restrictiveness is possible – in other words, the FRLA also has a bias toward the identity map, but this bias is subordinated to restrictiveness. (See footnote 11 on why this is only an assumption rather than a solid fact.)

Imagine a language Sanskrit' that is identical to real Sanskrit except that the inventory includes short mid vowels. Suppose the learner of Sanskrit' is in the early stages of morphophonemic analysis and has just discovered the coalescence alternation, so $\mathcal{L} = \{(ai, e:)\}$. The pre-free-ride support tableau τ_0 is given in (17). The FRLA produces the support tableau (18), which is Sanskrit'

10. If the learner of German has insufficient phonotactic data at the point when the FRLA is activated, then he might reach the wrong conclusion about the lexicon. Consider, for example, the effect of removing /daŋk/ from the support tableau, so there are no words with syllable-initial voiced obstruents. In that case, a grammar – a wrong one – will be found by BCD. This does not seem like a realistic concern. Morphophonemic learning requires prior morphological analysis, whereas phonotactic learning can precede all morphological analysis. Therefore, morphophonemic learning is unlikely to begin before phonotactic learning is already quite far advanced. It is reasonable to assume that the morphophonemic learner has already recognized data that are sufficiently representative of the language's phonotactic possibilities (cf. Hayes 2004, Pater & Tessier 2003, Prince & Tesar 2004). (I'm grateful to Anne-Michelle Tessier for raising this point.)

with all (eɜ, eɛ) identity maps replaced by (ai, eɛ).

(17) Sanskrit' without free ride – τ_0

lexicon	cands.	*MID	*DIPH	ID(VH)	ID-ADJ(VH)	UNIF
/ba/ (similarly /baɜ/)	→ ba					
	~ bi			₁ W	₁ W	
	~ be	₁ W		₁ W		
/bi/ (similarly /biɜ/)	→ bi					
	~ ba			₁ W	₁ W	
	~ be	₁ W		₁ W		
/be/	→ be	1				
	~ bi or ba	L		₁ W		
/beɜ/	→ beɜ	1				
	~ biɜ or baɜ	L		₁ W		
/va ₁ i ₂ /	→ veɜ _{1,2}	1		2		1
	~ va ₁ i ₂	L	₁ W	L		L
	~ vi _{1,2} or va _{1,2}	L		₁ L	₁ W	

(18) Sanskrit' with (ai, e:) free ride – τ_n

lexicon	cands.	*MID	*DIPH	ID(VH)	ID-ADJ(VH)	UNIF
/ba/ (similarly /ba:')	→ ba					
	~ bi			₁ W	₁ W	
	~ be	₁ W		₁ W		
/bi/ (similarly /bi:')	→ bi					
	~ ba			₁ W	₁ W	
	~ be	₁ W		₁ W		
/be/	→ be	1				
	~ bi or ba	L		₁ W		
/ba ₁ i ₂ /	→ be _{1,2}	1		2		1
	~ ba ₁ i ₂	L	₁ W	L		L
	~ bi _{1,2} or ba _{1,2}	L		₁ L	₁ W	
/va ₁ +i ₂ /	→ ve _{1,2}	1		2		1
	~ va ₁ i ₂	L	₁ W	L		L
	~ vi _{1,2} or va _{1,2}	L		₁ L	₁ W	

These two support tableaux, old and new, are ranked identically by BCD: [$*DIPH \gg ID-ADJ(VH) \gg ID(VH) \gg *MID \gg UNIF$]. Nothing is gained in restrictiveness by deriving nonalternating [e:] from /ai/, since the lingering /e/ → [e] identity map in (18) will not allow *MID to dominate ID(VH). In situations like this, where restrictiveness is not at stake, we may surmise that the learner does not take the pointless free ride. Therefore, the original support tableau (17) is retained and (18) is discarded.¹¹ In Sanskrit', nonalternating [e:] is derived from /e:/, not /ai/, because [e] exists and is derived from /e/.

(iii) A grammar is found, and it is more restrictive than the previous grammar. This is the situation in real Sanskrit. Assume that the morphophonemic learner has already arrived at the stage represented by the support tableau (19), ranked as in (12).

11. It is only an assumption rather than an established fact that learners choose (17) over (18), favoring the identity map when all else is equal. Ordinary linguistic data are unhelpful in such situations, though historical or psycholinguistic evidence can be relevant. See, e.g., Kiparsky (1973) or Dresher (1981) on the Alternation Condition.

(19) Sanskrit without free ride — τ_0

lexicon	cands.	*DIPH	ID-ADJ(VH)	ID(VH)	*MID	UNIF
/ba/ (similarly /baʔ/)	→ ba					
	~ bi		₁ W	₁ W		
	~ be			₁ W	₁ W	
/bi/ (similarly /biʔ/)	→ bi					
	~ ba		₁ W	₁ W		
	~ be			₁ W	₁ W	
/beː/ (similarly /beːʔ/)	→ beː				1	
	~ biː or baː			₁ W	L	
/va ₁ +i ₂ /	→ veː _{1,2}			2	1	1
	~ va ₁ i ₂	₁ W		L	L	L
	~ viː _{1,2} or vaː _{1,2}		₁ W	₁ L	L	

At this point, $\mathcal{G} = \{(ai, eː)\}$. A copy of (19) is made, substituting the (ai, eː) unfaithful map for all (eː, eː) identity maps. The table with this free ride is shown in (20).

(20) Sanskrit with (ai, eː) free ride — τ_n

lexicon	cands.	*DIPH	ID-ADJ(VH)	ID(VH)	*MID	UNIF
/ba/ (similarly /baʔ/)	→ ba					
	~ bi		₁ W	₁ W		
	~ be			₁ W	₁ W	
/bi/ (similarly /biʔ/)	→ bi					
	~ ba		₁ W	₁ W		
	~ be			₁ W	₁ W	
/ba ₁ i ₂ /	→ beː _{1,2}			2	1	1
	~ ba ₁ i ₂	₁ W		L	L	L
	~ biː _{1,2} or baː _{1,2}		₁ W	₁ L	L	
/va ₁ +i ₂ /	→ veː _{1,2}			2	1	1
	~ va ₁ i ₂	₁ W		L	L	L
	~ viː _{1,2} or vaː _{1,2}		₁ W	₁ L	L	

BCD is applied to the new support tableau (20), yielding the ranking in (21).

(21) Ranking from τ_n (20)

*DIPH >> ID-ADJ(VH) >> *MID >> ID(VH), UNIF

Since a grammar has been found, it needs to be compared for restrictiveness with the grammar derived from the original support tableau (19). We already saw that ranking in (12), and it is repeated in (22):

(22) Ranking of τ_o (19)

*DIPH >> ID-ADJ(VH) >> ID(VH) >> *MID >> UNIF

These grammars differ in restrictiveness. The grammar for the new support tableau (21) is more restrictive because it contains the ranking $\llbracket *MID \gg ID(VH) \rrbracket$, which effectively denies faithful treatment to input mid vowels. The grammar for the old support tableau (22) has the opposite ranking of these two constraints and thereby permits surface short mid vowels. This grammar describes a superset of Sanskrit.

The rankings in (22) and (21) are competing descriptions of Sanskrit based on different assumptions about the lexicon. Both grammars are consistent with the primary data, but (21) is more restrictive. Formally, the FRLA favors (21) because it has a higher r-measure (see (13)): the r-measure of (21) is 5, while the r-measure of (22) is 4.¹² Because (21) is more restrictive, the learner adopts its support tableau (20) and the attendant hypotheses about underlying representations.

In the discussion thus far, I've pretended that a single unfaithful map, learned from alternations, is sufficient with the FRLA to lead learners to the more restrictive grammar (21). This is a simplification; Sanskrit learners cannot get to (21) until \mathcal{Q} contains at least two unfaithful maps, (ai, e:) and (au, o:). The reason for this is that taking a free ride on only one of these unfaithful maps will leave the other identity map in the support tableau. For instance, if $\mathcal{Q} = \{(ai, e:)\}$, then the support tableau in which (e:, e:) has been eliminated will still contain some (o:, o:) identity maps. And the presence of (o:, o:) in a tableau is sufficient to force ID(VH) to dominate *MID, as in the less restrictive grammar (22).

This matter is already addressed in the FRLA. The learner constructs a support tableau for each of the subsets of \mathcal{Q} (including \emptyset and \mathcal{Q} itself, of course). If the learner has not yet encountered or attended to alternation data that support the (au, o:) unfaithful map, then \mathcal{Q} will not yet contain that map and the grammar (22) will persist. Eventually, when the learner does get a handle on this other unfaithful map, he will be presented with a support tableau that, by virtue of two free rides, contains no (e:, e:) or (o:, o:) identity maps, and BCD will find the more restrictive grammar in (21).

The FRLA is defined in this way so that grammar learning itself, rather than some special extragrammatical mechanism for data reduction, is responsible for discovering the generalization that all surface mid vowels are derived from diphthongs. By itself, an unfaithful map is the meagerest sort of generalization: (A, B) pairs the input segment /A/ with the output segment [B], abstracting away only from the actual form(s) in which this pairing occurs (see footnote 8). Unfaithful maps do not include information about the context in which they occur, nor do they state higher-order linguistic generalizations — e.g., there is no procedure for generalizing the (ai, e:) map to all mid vowels. Rather, the capacity for learning this or any other linguistic generalization is reserved to grammar learning by BCD. In this way, we avoid an unwelcome redundancy, where the learning mechanism preprocesses the data in a way that is closely attuned to what the constraints are looking for.

12. The r-measure of (21), for instance, is the sum of the number of markedness constraints that dominate ID-ADJ(VH) (that is, 1) and the number of markedness constraints that dominate ID(VH) and UNIF (that is, 2 each). The maximum r-measure for this constraint set is 6.

A final point. Once the learner has adopted the lexicon and grammar based on the free ride, he ought to use this knowledge to inform future learning. In Sanskrit, a free ride on the (ai, e:) map is possible because this map is invertible or biunique: upon encountering [e:], the learner can safely infer underlying /ai/. Information about binunique unfaithful maps, obtained from the FRLA, can be applied to all data that the learner subsequently encounters. As yet, however, I have no proposal for how learners might exploit this undoubtedly useful information.

* * *

This section concludes with a formal statement of the FRLA in Algol/C-inspired pseudocode.

(23) Definitions

a. (Support) tableau

A (support) tableau τ is a set of ordered 6-tuples (input, winner, $\mathfrak{R}_{(\text{input}, \text{winner})}$, loser, $\mathfrak{R}_{(\text{input}, \text{loser})}$, tableau-row), where *input* is an underlying representation, *winner* and *loser* are output candidates, $\mathfrak{R}_{(\text{input}, \text{winner})}$ and $\mathfrak{R}_{(\text{input}, \text{loser})}$ are their respective correspondence relations with the input, and *tableau-row* is an ordered n -tuple of elements chosen from $\{W, L, \emptyset\}$, with one element for each constraint in CON.

b. Unfaithful map in τ

An unfaithful map in the (support) tableau τ is a member (α_i, β_i) of $\mathfrak{R}_{(\text{input}, \text{winner})}$ in some member of τ , where $\alpha \neq \beta$. The set of all unfaithful maps in τ is designated by \mathcal{Q}_τ .

(24) Free-ride learning algorithm

Given: an initial support tableau τ and a set \mathcal{Q}_τ of all unfaithful maps in τ .

FRLA(τ, \mathcal{Q}_τ)

```

 $\tau_o := \tau$                                      /*initialize result*/
For every  $p_i \in \wp(\mathcal{Q}_\tau)$                    /* $\wp(\mathcal{Q}_\tau)$ =powerset of  $\mathcal{Q}_\tau$ */
     $\tau_n := \mathbf{Subst}(\tau_o, p_i)$            /*substitute unfaithful maps in  $\mathcal{Q}$ 
                                                subset  $p_i$  for identity maps in  $\tau_o$ */
    if BCD( $\tau_n$ ) = undefined                 /*if BCD fails, go to next  $p_i$ */
        continue
    else if R-measure(BCD( $\tau_n$ )) > R-measure(BCD( $\tau_o$ ))
         $\tau_o := \tau_n$                        /*if  $\tau_n$  has a higher r-measure, then
                                                it's now the one to beat*/
    endif
endfor
return( $\tau_o$ )

```

Remarks:

- **FRLA**(τ, \mathcal{Q}) accepts as input a tableau and a list of unfaithful maps. It returns a tableau that may or may not be different from the input tableau.
- **Subst**(τ, \mathcal{Q}) accepts as input a tableau and a list of unfaithful maps. For every unfaithful map (α, β) in \mathcal{Q} , it locates any winners in τ with (β, β) identity maps (i.e., it checks whether $(\beta, \beta) \in \mathfrak{R}_{(\text{input}, \text{winner})}$). The morphemes underlying these winners are altered by substituting α for β . As in “surgery” (Tesar et al. 2003), these changes are carried over to all forms in τ that contain the affected morphemes, and constraint assessments are updated to reflect the altered

inputs.

- **BCD**(τ) is defined by Prince & Tesar (2004). It accepts a tableau as input and returns a constraint hierarchy or *undefined* if no hierarchy can be found.
- **R-measure**(\mathcal{H}) is defined Prince & Tesar (2004) and in (13) above. It accepts as input a constraint hierarchy (with constraints tagged as markedness or faithfulness), and it returns a positive integer or zero.

The complexity of this learning algorithm is, in the worst case, proportional to the cardinality of $\wp(\mathcal{A})$ — that is, 2^N , where N is the number of unfaithful maps in \mathcal{A} . Exponentially increasing complexity always looks bad, but in this case it isn't as bad as it seems. As the learning problem scales up from examples of single processes to whole languages, or as the learner processes more alternation data, N grows very slowly. Most segments are mapped faithfully most of the time; unfaithful maps are very much the exception. There is no fixed upper limit on N , but in practice languages with more than, say, 20 distinct unfaithful maps are probably not common. There is, then, little danger of a massive combinatorial explosion in more realistic learning situations than have been contemplated here.

Furthermore, there may be ways of improving the FRLA's efficiency in actual practice. The various members of $\wp(\mathcal{A})$ are not equally likely to yield a grammar on their respective passes through the main loop of $\text{FRLA}(\tau, \mathcal{A})$. Because finding a grammar is roughly equivalent to finding a linguistically significant generalization, grammars are most likely to be found for free rides on those subsets of \mathcal{A} that (i) violate the same faithfulness constraint(s) and (ii) are maximal subject to (i). For example, the Sanskrit learner would be well advised to consider $p_i = \{(ai, e:), (au, o:)\}$ on an earlier pass through the loop than $p_i = \{(ai, e:)\}$ or $p_i = \{(au, o:)\}$, since $(ai, e:)$ and $(au, o:)$ violate the same faithfulness constraint, UNIF. Once the free ride has been found, FRLA continues in search of other free rides, but the complexity of the problem is reduced significantly: because there are no remaining $(e:, e:)$ or $(o:, o:)$ identity maps in the support tableau, it is no longer necessary to examine members of $\wp(\mathcal{A})$ that contain any unfaithful maps with $[e:]$ or $[o:]$ as output.

4. Limitations of and Extensions to the FRLA

The FRLA has a significant limitation: it does not work for contextually restricted free rides. In Sanskrit, the free ride is context-free: surface $[e:]$, regardless of its environment, is derived from underlying $/ai/$. But in cases like colloquial Arabic, the free ride is limited to a particular context.

Colloquial Arabic word-final vowels are always short (for details of the analysis sketched here, see McCarthy (2004)). Alternations like $[ˈjilqa] \sim [ˈjilqaːni]$ 'he will find ~ he will find me' indicate that the final $/a/$ of $[ˈjilqa]$ is derived from an underlying long vowel $/aː/$. Shortening of long vowels in word-final position is common cross-linguistically. The responsible markedness constraints are NONFINALITY (NONFIN), which prohibits word-final stress, and WEIGHT-TO-STRESS (WSP), which requires heavy syllables to be stressed. By dominating ID(long), they map underlying $/jilqaː/$ to $[ˈjilqa]$ rather than faithful $*[ˈjilqaː]$ (which violates NONFIN) or $*[ˈjilqaː]$ (which violates WSP).

Significantly, colloquial Arabic has no stems like $*[ˈjitba] \sim [ˈjitbani]$, with a stem-final vowel that is short before suffixes. Since stem-final vowels are always long before suffixes, stem-final vowels are always derived from underlying long vowels. A hypothetical stem like $/jitba/$ must be treated unfaithfully, perhaps by deleting the final short vowel. The ranking arguments in (25) and

(26) account for this. The constraint FINAL-C in (25) requires words to end in a consonant, and by dominating MAX(V) it compels deletion of word-final short vowels. But the constraint MAX(V:), for which there is solid support elsewhere in the language, blocks enforcement of FINAL-C on underlying final long vowels.¹³

(25) FINAL-C >> MAX(V) (hypothetical example)

/jitba/	FINAL-C	MAX(V)
→ 'jitb		1
a. ~ 'jitba	₁ W	L

(26) MAX(V:) >> FINAL-C

/jilqa:/	MAX(V:)	FINAL-C
→ 'jilqa		1
a. ~ 'jilq	₁ W	L

Colloquial Arabic therefore has a chain shift: /V:#/ → [V#] and /V#/ → [Ø#]. (MAX(V:) prevents the fell-swoop map /V:#/ → [Ø#].) This chain shift, unlike the one in Sanskrit, is limited to a particular context, word-final position. Alternation data will tell the learner that *some* [V#]s derive from /V:#/s, but alternation data won't help with other [V#]s, either because the learner has not heard the suffixed form or because there is no suffixed form (e.g., ['huwwa] 'he'). The FRLA ought to get the learner over this hump to conclude that *all* [V#]s derive from /V:#/s. If the learner fails to get over that hump, then he will have learned a less restrictive grammar includes the ranking [[MAX(V) >> FINAL-C]. This grammar wrongly permits stems like *['jitba]~['jitbani].

Unfortunately, the FRLA in its present form is not up to this task. It can deal with an across-the-board free ride like Sanskrit's, but not with a contextually limited one. Alternation data tell the learner that there is a (a:, a) unfaithful map in the language. But replacing all (a, a) identity maps with (a:, a) maps is a gross error: it's not possible to derive, say, ['katab] 'he wrote' from /ka:ta:b/, since the language does not have a general shortening process (cf. ['ki'ta:b] 'book', ['ka:tib] 'writing'). What is required is a way for learners to discover contextually limited free rides like the one in Colloquial Arabic: only word-final (a, a) identity maps should be replaced by (a:, a). Identity maps in other positions must be left intact.

There are two ways to fix this problem with the FRLA. One way is to greatly increase the power of the currently quite limited theory of unfaithful maps, so that they can be distinguished by the contexts in which the maps occur. For reasons already discussed (see the end of section 3), this seems like a poor idea: it requires a special auxiliary learning mechanism, independent of the grammar, that recapitulates the interactional possibilities of NONFIN and WSP. The contexts in which unfaithful maps occur are defined by the grammar, and any special learning mechanism charged with discovering and recording those contexts would simply duplicate the effect of the grammar.

The other way of improving the FRLA is to provide it with a wider range of hypotheses about the lexicon. One of those hypotheses is a lexicon where only word-final [a]s take a free ride on

13. On why there are no stems like *['jitb]~['jitbani], see McCarthy (2004).

(a:, a). That is, the learner must consider lexicons in which various subsets of the (a, a) identity maps in the support tableau have been replaced by (a:, a). These possibilities are listed in (27) for a dataset that contains a form with underlying /a:/ determined from alternations ([ˈjilqa]~[jilˈqa:ni]), two forms with underlying nonfinal /a:/ ([kiˈta:b], [ˈka:tib]), and two forms with nonalternating surface [a] that might be eligible for a free ride ([ˈhuwwa], [ˈkatab])

(27) Hypothesized lexicons

- a. /jilqa:/, /kita:b/, /ka:tib/, /huwwa/, /katab/
- b. /jilqa:/, /kita:b/, /ka:tib/, /huwwa/, /ka:tab/
- c. /jilqa:/, /kita:b/, /ka:tib/, /huwwa/, /kata:b/
- d. /jilqa:/, /kita:b/, /ka:tib/, /huwwa/, /ka:ta:b/
- e. /jilqa:/, /kita:b/, /ka:tib/, /huwwa:/, /katab/
- f. /jilqa:/, /kita:b/, /ka:tib/, /huwwa:/, /ka:tab/
- g. /jilqa:/, /kita:b/, /ka:tib/, /huwwa:/, /kata:b/
- h. /jilqa:/, /kita:b/, /ka:tib/, /huwwa:/, /ka:ta:b/

Lexicon (27a) has no free rides. Lexicon (27h) has an across-the-board free ride: all [a]s everywhere are derived from /a:/. Lexicon (27e) is the desired result: only word-final [a] takes a free ride on the (a:, a) unfaithful map. Other lexicons represent other choices about which [a]s to grant a free ride to.

When presented with this wider range of possible lexicons, the FRLA will select the correct one, (27e). The short explanation is that BCD cannot find a grammar for (27b, c, d, f, g, h), and (27e) is preferred to (27a) because its grammar has a higher r-measure. The long explanation is that no grammar can be found for (27b, c, d, f, g, h) because no ranking of these constraints can derive [ˈkatab] from /ka:tab/, /kata:b/, or /ka:ta:b/ while also preserving the long vowels of /kita:b/ and /ka:tib/. For example, (28) is a support tableau that presupposes the lexicon in (27h), where all [a]s are derived from /a:/.

(28) Across-the-board (a:, a) free ride (= lexicon (27h))¹⁴

lexicon	cands.	NONFIN	WSP	ID(long)	FINAL-C	MAX(V:)	MAX(V)
/jilqa:/	→ 'jilqa			1	1		
	~ 'jilqa:		₁ W	L	₁		
	~ jil'qa:	₁ W	₁ W	L	₁		
	~ 'jilq			L	L	₁ W	₁ W
/kita:b/	→ ki'ta:b						
	~ 'kitab			₁ W			
/ka:tib/	→ 'ka:tib						
	~ 'katib			₁ W			
/huwwa:/	→ 'huwwa			1	1		
	~ 'huwwa:		₁ W	L	₁		
	~ huw'wa:	₁ W	₁ W	L	₁		
	~ 'huww			L	L	₁ W	₁ W
/ka:ta:b/	→ 'katab			2			
	~ ka'ta:b			₁ L			
	~ 'ka:tab			₁ L			
	~ 'ka:ta:b		₁ W	L			

The first two /ka:ta:b/ loser rows contain L but no W, so there is no ranking for BCD to find. Other markedness constraints could be introduced to deal with this, but then they would interfere with /kita:b/ and /ka:tib/.

BCD succeeds in finding a grammar for (27a) and (27e). The support tableau for the latter is given in (29).

14. The constraint assessments in the support tableau (28) makes certain standard assumptions about Arabic prosody (see, e.g., Hayes (1995), where they go under the rubric of final consonant extrametricality): final superheavy syllables actually consist of a heavy syllable followed by a word appendix, which is sufficient to satisfy NONFIN; and word-final consonants do not project a mora, so final CVC is light.

(29) Word-final (a:, a) free ride (= lexicon (27e))

lexicon	cands.	NONFIN	WSP	ID(long)	FINAL-C	MAX(V:)	MAX(V)
/jilqa:/	→ 'jilqa			1	1		
	~ 'jilqa:		₁ W	L	₁		
	~ jil'qa:	₁ W	₁ W	L	₁		
	~ 'jilq			L	L	₁ W	₁ W
/kita:b/	→ ki'ta:b						
	~ 'kitab			₁ W			
/ka:tib/	→ 'ka:tib						
	~ 'katib			₁ W			
/huwwa:/	→ 'huwwa			1	1		
	~ 'huwwa:		₁ W	L	₁		
	~ huw'wa:	₁ W	₁ W	L	₁		
	~ 'huww			L	L	₁ W	₁ W
/katab/	→ 'katab						
	~ ka'ta:b			₁ W			
	~ 'ka:tab			₁ W			
	~ 'ka:ta:b		₁ W	₁ W			

For this support tableau, BCD finds the grammar [NONFIN, WSP >> MAX(V:) >> FINAL-C >> ID(long), MAX(V)], with an r-measure of 8.¹⁵

When BCD is applied to the support tableau for the lexicon with no free rides, (27a), it finds a different grammar, [NONFIN, WSP >> MAX(V:) >> MAX(V) >> FINAL-C >> ID(long)]. This grammar has a lower r-measure, 7, and is in fact less restrictive, since it allows stems of the forbidden *[jɪtba]~[jɪtbaɪ] type. The FRLA therefore favors the lexicon in (27e), in which only word-final [a] gets a free ride on (a:, a).

The general idea, then, is that the learner discovers a contextually limited free ride on the unfaithful map (A, B) by testing all of the logically possible combinations of (B, B) identity maps throughout the support tableau, relying on BCD to determine which of these possibilities is linguistically meaningful. BCD can do this because random, linguistically meaningless combinations of free rides are not analyzeable with the given constraint-set, so BCD quickly fails to find a grammar. In Arabic, the free ride is limited to word-final position because that is where the length contrast is neutralized, as a result of NONFIN and WSP. The learner discovers this by considering various lexicons in which the free ride on the shortening process occurs in word-final position and in other, sometimes random positions as well. Most of this space of possibilities is occupied by lexicons for which no grammar is possible.

15. MAX(V:) is ranked before MAX(V) under the assumption that their ranking is fixed universally (cf. Beckman 1998). Furthermore, additional data involving syncope in other positions require this ranking (McCarthy 2004).

The worst-case complexity of the FRLA as defined in section 3 is exponential in the size of \mathcal{L} . The complexity of the system I've just sketched is exponential in the number of morphemes in the support tableau — a much more serious matter. The size of the support tableau (i.e., the number of winner~loser pairs) that is required for ranking is, in the worst case, quadratic in the number of constraints that must be ranked (Tesar & Smolensky 2000: 99). The overall complexity of this system is therefore proportional to 2^{N^2} , where N is the number of constraints in CON. This is deeply problematic, particularly since a reasonable value for N is probably an order of magnitude larger than a reasonable value for the cardinality of \mathcal{L} .

It may be possible to avoid this combinatorial explosion with an iterative procedure for lexicon building. Start with a support tableau that includes only the forms that are not eligible for the (A, B) free ride, either because they contain no [B]s or because any [B]s that they do contain are known from alternations to be derived from /A/s. (In the Arabic example above, this would mean starting with just /jilqa:/, /kita:b/, and /ka:tib/, the three constants among the lexicons in (27).) Now, take each of the remaining [B]-containing forms one at a time and construct its possible underlying representations with and without the free ride. (In Arabic, these possibilities include /huwwa:/ and /huwwa/ for ['huwwa] and /ka:ta:b/, /ka:tib/, /kata:b/, and /katab/ for ['katab].) Add each of these possibilities separately to the known, constant support tableau and apply BCD, which will determine which possibilities work at all and whether any lead to a more restrictive grammar. The result will be the same as the computationally intractable algorithm, but the worst-case complexity is much more modest: it is exponential in the largest number of [B]s that any word contains.

5. Conclusion

In morphophonemic learning, the underlying representations influence the grammar and the grammar influences the underlying representations. Constraint demotion already ensures that underlying representations will influence the grammar: constraints are ranked on the basis of the input→output maps in a mark-data table. The main challenge for a learning algorithm, then, is how to get information going in the other direction as well.

In this article, I have pursued the idea that learners simultaneously consider various hypotheses about underlying representation, rejecting any for which no grammar is possible and preferring the one that allows the most restrictive grammar. The inconsistency-detection property of multi-recursive constraint demotion (Tesar 1997) is crucial in eliminating the hypotheses for which there is no grammar, and the r-measure (Prince & Tesar 2004) is crucial in determining which grammar is the most restrictive.

Morphophonemic learning is a complex problem, and I have addressed only a small piece of it. The cases examined here involve free-ride effects: alternation data show that some [B]s derive from /A/s, and other evidence shows that even nonalternating [B]s must also derive from /A/s. In other words, the grammar maps /A/ to [B], and this information propagates back to the lexicon from nonalternating [B]s. The learner considers lexicons in which nonalternating [B]s are and are not derived from /A/s; if a more restrictive grammar can be found, then the former is chosen.

Morphophonemic learning necessarily intersects with phonological opacity (see the appendix), which presents its own learning challenges. It remains to be seen whether the proposal developed here can be modified and extended to address this notably thornier problem.

Appendix: Further Free-Ride Examples

1. Coalescence in Choctaw

In Choctaw (Lombardi & McCarthy 1991, Nicklas 1974, 1975, Ulrich 1986), preconsonantal vowel-nasal sequences coalesce to form a long nasalized vowel:

(30) Choctaw coalescence

/am-pala/	ã:pala	'my lamp'	cf. amissi 'my deer'
/am-tabɪ/	ã:tabɪ	'my cane'	
/am-miŋko/	ã:mĩ:ko	'my chief'	
/onsi/	õ:si	'eagle'	

In a closed syllable, coalescence yields a short nasalized vowel (/ta-N-kçi/ → [tãkçi] 'to tie' with a nasal infix). This is an expected consequence of Choctaw's syllable canons. Word-final nasals coalesce only if they are affixal; the resulting vowel is also short for prosodic reasons. Geminate nasals that are tautomorphemic or derived by assimilation do not coalesce: [onna] 'dawn', /ta-l-na-h/ → [tannah] 'to be woven'. These are typical geminate inalterability effects (Guerssel 1977, 1978, Hayes 1986, Schein & Steriade 1986).

In Choctaw, the vowel inventory has a systematic gap: except in environments where long vowels are prohibited (closed syllables and word-finally), nasalized vowels are always long. In Ulrich's (1986: 60ff.) analysis, this gap is explained by prohibiting underlying nasal vowels with a MSC and deriving all surface nasal vowels by coalescence of VN sequences. Even nonalternating nasalized vowels must be derived by coalescence since they too are consistently long (e.g., 'chief' and 'eagle' in (30)). This analysis not only explains the length of nasalized vowels in terms of an independently motivated process, but it also provides an elegant explanation for an otherwise puzzling observation: Choctaw has no nasalized vowels followed by tautomorphemic nasal consonant *[õ:na] (Ulrich 1986: 62). The reason: geminate inalterability protects a nasal consonant from coalescence when it is followed by a tautomorphemic nasal, so there is no way of deriving a sequence of a nasalized vowel and a tautomorphemic nasal.

In OT with richness of the base, it is not possible to ban nasalized vowels from inputs. The closest alternative is a chain-shift analysis similar to Gnanadesikan's account of Sanskrit: /...aNC.../ → [...ã:C...] but /ã(:)/ → [a(:)]. In this way, input nasalized vowels are pushed out of the way and all surface nasalized vowels are derived by coalescence. This analysis requires the FRLA like Sanskrit. On the basis of alternating morphemes like /am/, the learner discovers the unfaithful coalescent map. The free-ride algorithm permits the learner to extend this map to nonalternating morphemes like [õ:si], thereby accommodating the more restrictive grammar of a language that prohibits nasalized vowels that are short (in non-final open syllables) or that are followed by tautomorphemic nasals.

In the discussion of Sanskrit, I mentioned an alternative to the chain-shift analysis: introduce a markedness constraint against short mid vowels. A similar alternative might be proposed for Choctaw: if there is a markedness constraint against short nasalized vowels, couldn't it explain the same facts without the chain shift and consequently without free-ride learning? The problem with this approach is that it will not account for Ulrich's observation that nasalized vowels are never followed by tautomorphemic nasal consonants. Indeed, it's hard to imagine any reasonable markedness constraint that could give direct expression to this generalization.

2. Coalescence in Rotuman

De Haas (1988) cites Rotuman as an example of non-structure-preserving coalescence. It requires free-ride learning in much the same way as Sanskrit and Choctaw. The primary source on Rotuman is Churchward (1940). For a bibliography of the extensive secondary literature, see McCarthy (2000).

Rotuman distinguishes two “phases”, complete and incomplete. The distribution of the two phases is based on prosodic structure; my focus here is on the way that the phase difference is realized in stems. A stem that is in the incomplete phase always ends in a stressed heavy syllable; heavy syllables are otherwise prohibited (except finally in monosyllables and recent loans). Stem-final /CV₁CV₂/ sequences are turned into heavy syllables in three different ways, depending of the quality of V₁ and V₂. If V₁ is a back rounded vowel and V₂ is front and no lower than V₁, they coalesce to form the front rounded counterpart of V₁:

(31) Coalescence in Rotuman

/hoti/	høt	‘to embark’
/mose/	møs	‘to sleep’
/futi/	fyt	‘to pull’

The result of coalescence is a short vowel because syllables are maximally bimoraic. The underlying forms show up unaltered in the complete phase of these stems: [‘hoti], [‘mose], [‘futi].

Front rounded vowels have no other source in Rotuman, and so they are only ever found in final closed syllables of morphemes in the incomplete phase.¹⁶ For example, *[‘høti] and *[‘møse] are not possible words of this language. As in the chain-shift analyses of Sanskrit and Choctaw, all surface front rounded vowels will be derived by coalescence, while any input front rounded vowel maps unfaithfully to something else, such as its nonround counterpart. The FRLA is required so the learner can deal with any morphemes that always appear in the incomplete phase or that have only been observed in the incomplete phase.

There is no plausible markedness alternative to the chain shift and the free ride. To explain the distribution of front rounded vowels without deriving all of them by coalescence, it would be necessary to posit a markedness constraint banning front rounded vowels from open syllables. Such a constraint enjoys no independent typological support and is purely descriptive.

3. Opacity and allophony in Japanese

Ito & Mester (2001: 280-281, 2003) identify a class of cases, which they dub “allophonic masking”, that present problems for input-based theories of opacity in OT. The interaction in Japanese between nasalization of voiced velar stops and *rendaku* voicing is a perfect example.

The voiced velar stop in initial position regularly alternates with its nasal counterpart in medial position: [go] ‘the game of go’, [okiŋo] ‘handicap go’. As is well known, medial voiced obstruents block *rendaku* voicing of stem-initial obstruents: [kamikaze] ‘divine wind’, *[kamigaze].

16. An exception: front rounded vowels are also derived by harmony with a front rounded vowel that is itself derived by coalescence: /pulufi/ → [pylyf] ‘to adhere’.

In this respect, medial [ŋ] acts just like a voiced obstruent, even if it never alternates with [g]: [sakatoŋe] ‘reverse thorn’, *[sakadoŋe]. (Additionally, [g]s derived by *rendaku* also nasalize: [oriŋami] ‘paper folding’ from /ori-kami/).

In a traditional analysis, Japanese requires a MSC excluding /ŋ/ from underlying representations. The *rendaku* rule is ordered early enough to see the underlying /g/ in /saka-toŋe/, and so voicing of the /t/ is correctly blocked. By a later rule, medial /g/ becomes [ŋ].

Because of ROTB, OT has no MSC’s, leading Ito & Mester to propose a stratal OT analysis that uses an early stratum to approximate the effect of the MSC. Although the rich base offers both /g/ and /ŋ/ in initial and medial positions, the phonology of the first stratum wipes out this potential distinction, mapping /ŋ/ to [g] in all contexts. *Rendaku* also occurs in the first stratum, where it sees only [g]. In the second stratum, medial [g] that is the output of the first stratum becomes nasalized.

This example and others like it challenge approaches to opacity in OT that rely on properties of the input, such as sympathy (McCarthy 1999), comparative markedness (McCarthy 2003a), or turbidity (Goldrick 2000, Goldrick & Smolensky 1999). These approaches would refer to the /g/ that underlies surface [ŋ] to explain why there is no *rendaku* voicing in [sakatoŋe]. For example, one might assume in turbid fashion that the [+voice] feature of underlying /g/ remains present in surface structure, though unpronounced, and its presence is sufficient to block *rendaku*. The problem with this analysis is that, under ROTB, there is no guarantee that surface [ŋ] derives from /g/; it might simply derive from /ŋ/, in which case *rendaku* would not be blocked. In fact, the analysis wrongly predicts a contrast between medial [ŋ]s that do and do not block *rendaku*, the former derived from /g/ and the latter from /ŋ/. There is no such contrast, however. The MSC-based analysis and Ito & Mester’s stratal approach do not predict this contrast because all surface [g]s and [ŋ]s derive from earlier gs.

Gnanadesikan’s analysis of Sanskrit suggests a different way of looking at Japanese, however. Suppose that all surface [ŋ]s are derived from underlying /g/s and that the grammar maps input /ŋ/ to something other than /g/, such as [n]. This is a chain shift: /g/ → [ŋ] (medially) and /ŋ/ → [n] (everywhere). Abstractly, the learning problem presented by this chain shift is the same as we saw with Sankrit.

For concreteness, I will borrow several markedness constraints from Ito & Mester’s (2003) analysis.

- (32) a. *VgV
Intervocalic g is prohibited (a lenition constraint).
b. *g
Voiced velar stops are prohibited (cf. Ohala 1983: 196-197).
c. *ŋ
Velar nasals are prohibited.

We will look at their interaction among themselves and with two faithfulness constraints, IDENT(Nasal) and IDENT(Place).

For the phonotactic learner of Japanese, the support tableau contains observed initial [g]s and medial [ŋ]s, both of which are derived by the identity map.

(33) Support tableau at phonotactic stage

lexicon	cands.	*VgV	*g	*ŋ	Id(Nas)	Id(Pl)
/ga/	→ ga		1			
	~ da		L			₁ W
	~ ŋa		L	₁ W	₁ W	
	~ na		L		₁ W	₁ W
/aŋa/	→ aŋa			1		
	~ aga	₁ W	₁ W	L	₁ W	
	~ ana			L		₁ W

On the first pass through BCD, *VgV is identified as the only rankable markedness constraint. On the next pass, we seek a faithfulness constraint that can be ranked so as to free up one of the remaining markedness constraints for ranking. Id(Pl) frees up *ŋ, which accounts for the remaining mark-data pairs. The residual constraints, *g and Id(Nas), are ranked so as to conform to the markedness-high bias. The resulting ranking is given in (34).

(34) Result of phonotactic learning for Japanese

*VgV >> Id(Pl) >> *ŋ >> *g >> Id(Nas)

As the learner grows aware of data from alternations, he learns that there are unfaithful maps /g/ → [ŋ] that occur when stem-initial /g/ becomes medial as a result of the morphology. This unfaithful map is added to the support tableau and BCD is applied once again.

(35) Support tableau after some morphophonemic learning

lexicon	cands.	*VgV	*g	*ŋ	Id(Nas)	Id(Pl)
/ga/	→ ga		1			
	~ da		L			₁ W
	~ ŋa		L	₁ W	₁ W	
	~ na		L		₁ W	₁ W
/aŋa/	→ aŋa			1		
	~ aga	₁ W	₁ W	L	₁ W	
	~ ana			L		₁ W
/a-ga/	→ aŋa			1	1	
	~ aga	₁ W	₁ W	L	L	
	~ ana			L		₁ W

BCD gives this support tableau the same ranking, (34). The commitment to the identity map for nonalternating forms means that [aŋa] must derive from /aŋa/ — and the presence of the (ŋ, ŋ)

identity map in the support tableau forces the ranking $ID(PI) \gg * \eta$. The proposed chain shift, where underlying / η / maps to [n], is apparently not learnable under these assumptions about the input. It is learnable with the FRLA, however, since its basic character is similar to Sanskrit.

In the discussion of Sanskrit, I noted that a local solution to the learning problem was possible by adding another markedness constraint. No such move will solve the problem in Japanese, however. Surface [n]s all act like they are derived from underlying /g/s with respect to blocking *rendaku*. Fiddling with the constraints will not somehow magically subvert the commitment to the identity map.

References

- Alderete, John; Tesar, Bruce (2002). "Learning Covert Phonological Interaction: An Analysis of the Problem Posed by the Interaction of Stress and Epenthesis". Rutgers University Center for Cognitive Science, New Brunswick, NJ. Report 72. [Available on Rutgers Optimality Archive, ROA-543.]
- Angluin, Dana (1980). "Inductive Inference of Formal Languages from Positive Data". *Information and Control* 45: 117-135.
- Baker, C. L. (1979). "Syntactic Theory and the Projection Problem". *Linguistic Inquiry* 10: 533-581.
- Beckman, Jill (1998). *Positional Faithfulness*. University of Massachusetts Amherst, Doctoral dissertation. [Available on Rutgers Optimality Archive, ROA-234.]
- Churchward, C. M. (1940). *Rotuman Grammar and Dictionary*. Sydney: Australasia Medical Publishing Co. [Repr. 1978, AMS Press, New York.]
- de Haas, Wim (1988). *A Formal Theory of Vowel Coalescence: A Case Study of Ancient Greek*. Dordrecht: Foris.
- Dinnsen, Daniel A. (2004). "A Typology of Opacity Effects in Acquisition". Ms., Indiana University, Bloomington, IN.
- Dinnsen, Daniel A.; Barlow, Jessica A. (1998). "On the Characterization of a Chain Shift in Normal and Delayed Phonological Acquisition". *Journal of Child Language* 25: 61-94.
- Dinnsen, Daniel A.; McGarrity, Laura W. (2004). "On the Nature of Alternations in Phonological Acquisition". *Studies in Phonetics, Phonology, and Morphology* 11: 23-42. [Available at <http://www.indiana.edu/~sndlrng/DinnsenMcGarrity%2004.pdf>.]
- Dinnsen, Daniel A.; O'Connor, Kathleen; Gierut, Judith (2001). An Optimality Theoretic Solution to the Puzzle-Puddle-Pickle Problem. Handout of paper presented to Annual Meeting, Linguistic Society of America, Washington, DC.
- Dresher, B. Elan (1981). "On the Learnability of Abstract Phonology". In: Baker, C. L.; McCarthy, John (eds.). *The Logical Problem of Language Acquisition*. Cambridge, MA: MIT Press, pp. 188-210.
- Gnanadesikan, Amalia (1997). *Phonology with Ternary Scales*. University of Massachusetts at Amherst, Doctoral dissertation. [Available on Rutgers Optimality Archive, ROA-195.]
- Goldrick, Matthew (2000). "Turbid Output Representations and the Unity of Opacity". In: Hirotani, Masako (ed.) *Proceedings of the North East Linguistics Society* 30. Amherst, MA: GLSA Publications, pp. 231-246.
- Goldrick, Matthew; Smolensky, Paul (1999). "Opacity, Turbid Representations, and Output-Based Explanation". Talk presented at Workshop on the Lexicon in Phonetics and Phonology, Edmonton, Alberta.

- Guerssel, Mohammed (1977). "Constraints on Phonological Rules". *Language* 3: 267-305.
- Guerssel, Mohammed (1978). "A Condition on Assimilation Rules". *Linguistic Analysis* 4: 225-254.
- Hayes, Bruce (1986). "Inalterability in CV Phonology". *Language* 62: 321-251.
- Hayes, Bruce (1995). *Metrical Stress Theory: Principles and Case Studies*. Chicago: The University of Chicago Press.
- Hayes, Bruce (2004). "Phonological Acquisition in Optimality Theory: The Early Stages". In: Kager, René; Pater, Joe; Zonneveld, Wim (eds.). *Fixing Priorities: Constraints in Phonological Acquisition*. Cambridge: Cambridge University Press, pp. 158-203. [Available on Rutgers Optimality Archive.]
- Ito, Junko; Mester, Armin (1999). "The Phonological Lexicon". In: Tsujimura, Natsuko (ed.) *The Handbook of Japanese Linguistics*. Oxford: Blackwell, pp. 62-100.
- Ito, Junko; Mester, Armin (2001). "Structure Preservation and Stratal Opacity in German". In: Lombardi, Linda (ed.) *Segmental Phonology in Optimality Theory*. Cambridge: Cambridge University Press, pp. 261-295.
- Ito, Junko; Mester, Armin (2003). "Lexical and Postlexical Phonology in Optimality Theory: Evidence from Japanese". *Linguistische Berichte Sonderheft 11: Resolving conflicts in grammars*: 183-207. [Available from <http://people.ucsc.edu/~ito/PAPERS/lexpostlex.pdf>.]
- Kager, René (1999). *Optimality Theory*. Cambridge: Cambridge University Press.
- Kiparsky, Paul (1973). "Phonological Representations". In: Fujimura, O. (ed.) *Three Dimensions of Linguistic Theory*. Tokyo: TEC, pp. 3-136.
- Lombardi, Linda; McCarthy, John (1991). "Prosodic Circumscription in Choctaw Morphology". *Phonology* 8: 37-71.
- McCarthy, John J. (1999). "Sympathy and Phonological Opacity". *Phonology* 16: 331-399.
- McCarthy, John J. (2000). "The Prosody of Phase in Rotuman". *Natural Language and Linguistic Theory* 18: 147-197.
- McCarthy, John J. (2003a). "Comparative Markedness". *Theoretical Linguistics* 29: 1-51.
- McCarthy, John J. (2003b). "Sympathy, Cumulativity, and the Duke-of-York Gambit". In: Féry, Caroline; Vijver, Ruben van de (eds.). *The Syllable in Optimality Theory*. Cambridge: Cambridge University Press, pp. 23-76.
- McCarthy, John J. (2004). "The Length of Stem-Final Vowels in Colloquial Arabic". Ms., University of Massachusetts Amherst, Amherst, MA. [Available on Rutgers Optimality Archive, ROA-616. To appear in *Perspectives on Arabic Linguistics* 18, ed. Dilworth B. Parkinson. Amsterdam: John Benjamins.]
- McCarthy, John J.; Prince, Alan (1995). "Faithfulness and Reduplicative Identity". In: Beckman, Jill; Walsh Dickey, Laura; Urbanczyk, Suzanne (eds.). *University of Massachusetts Occasional Papers in Linguistics 18*. Amherst, MA: GLSA Publications, pp. 249-384. [Available on Rutgers Optimality Archive, ROA-103.]
- McCarthy, John J.; Prince, Alan (1999). "Faithfulness and Identity in Prosodic Morphology". In: Kager, René; van der Hulst, Harry; Zonneveld, Wim (eds.). *The Prosody-Morphology Interface*. Cambridge: Cambridge University Press, pp. 218-309. [Excerpted in *Optimality Theory in Phonology: A Reader*, ed. by John J. McCarthy, Malden, MA and Oxford, Blackwell (2004).]
- Moreton, Elliott (2000). Faithfulness and Potential. Ms., University of Massachusetts Amherst, Amherst, MA.
- Nicklas, Thurston (1974). *The Elements of Choctaw*. University of Michigan, Doctoral dissertation.

- Nicklas, Thurston (1975). "Choctaw Morphophonemics". In: Crawford, J. (ed.) *Studies in Southeastern Indian Languages*. Athens, GA: University of Georgia Press, pp. 237-250.
- Ohala, John (1983). "The Origin of Sound Patterns in Vocal Tract Constraints". In: MacNeilage, Peter (ed.) *The Production of Speech*. New York: Springer-Verlag, pp. 189-216.
- Ota, Mitsuhiro (2004). "The Learnability of the Stratified Phonological Lexicon". *Journal of Japanese Linguistics* 20. [Available on Rutgers Optimality Archive, ROA-668.]
- Pater, Joe (2004). Course Handout. University of Massachusetts Amherst, Amherst, MA.
- Pater, Joe; Tessier, Anne-Michelle (2003). "Phonotactic Knowledge and the Acquisition of Alternations". In: Solé, M. J.; Recasens, D.; Romero, J. (eds.). *Proceedings of the 15th International Congress of Phonetic Sciences*. Barcelona: Universitat Autònoma de Barcelona, pp. 1177-1180. [Available at <http://people.umass.edu/pater/pater-tessier.pdf>.]
- Prince, Alan (2002). "Arguing Optimality". Ms., Rutgers University, New Brunswick, NJ. [Available on Rutgers Optimality Archive, ROA-562.]
- Prince, Alan; Smolensky, Paul (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, Mass., and Oxford, UK: Blackwell. [Revision of 1993 technical report, Rutgers University Center for Cognitive Science. Available on Rutgers Optimality Archive, ROA-537.]
- Prince, Alan; Tesar, Bruce (2004). "Learning Phonotactic Distributions". In: Kager, René; Pater, Joe; Zonneveld, Wim (eds.). *Fixing Priorities: Constraints in Phonological Acquisition*. Cambridge: Cambridge University Press, pp. 245-291. [Available on Rutgers Optimality Archive.]
- Schane, Sanford (1987). "The Resolution of Hiatus". In: Bosch, A.; Need, B.; Schiller, E. (eds.). *CLS 23: Parasession on Autosegmental and Metrical Phonology*. Chicago: Chicago Linguistic Society, pp. 279-290.
- Schein, Barry; Steriade, Donca (1986). "On Geminates". *Linguistic Inquiry* 17: 691-744.
- Tesar, Bruce (1995). *Computational Optimality Theory*. University of Colorado, Doctoral dissertation. [Available on Rutgers Optimality Archive.]
- Tesar, Bruce (1997). "Multi-Recursive Constraint Demotion". Ms., Rutgers University, New Brunswick, NJ. [Available on Rutgers Optimality Archive, ROA-197.]
- Tesar, Bruce; Alderete, John; Horwood, Graham; Merchant, Nazarré; Nishitani, Koichi; Prince, Alan (2003). "Surgery in Language Learning". In: Garding, G.; Tsujimura, M. (eds.). *Wccfl 22 Proceedings*. Somerville, MA: Cascadilla Press, pp. 477-490.
- Tesar, Bruce; Prince, Alan (2004). "Using Phonotactics to Learn Phonological Alternations". In: *CLS 39, Part II: The Panels*. Chicago: Chicago Linguistic Society. [Available on Rutgers Optimality Archive, ROA-620.]
- Tesar, Bruce; Smolensky, Paul (1994). "The Learnability of Optimality Theory". In: Aranovich, Raul; Byrne, William; Preuss, Susanne; Senturia, Martha (eds.). *Proceedings of the Thirteenth West Coast Conference on Formal Linguistics*. Stanford, CA: CSLI Publications, pp. 122-137.
- Tesar, Bruce; Smolensky, Paul (1998). "Learnability in Optimality Theory". *Linguistic Inquiry* 29: 229-268.
- Tesar, Bruce; Smolensky, Paul (2000). *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.
- Tessier, Anne-Michelle (2004). Looking for Oo-Faithfulness in Phonological Acquisition. Ms., University of Massachusetts Amherst, Amherst, MA.

- Ulrich, Charles (1986). *Choctaw Morphophonology*. UCLA, Doctoral dissertation.
- Whitney, W. D. (1889). *Sanskrit Grammar*. Cambridge, MA: Harvard University Press.
- Winter, Edgar. 1972. *Free Ride: Epic Records*.
- Zwicky, Arnold M. (1970). "The Free-Ride Principle and Two Rules of Complete Assimilation in English". In: *Papers from the Sixth Regional Meeting Chicago Linguistic Society*. Chicago: Chicago Linguistic Society, pp. 579-588.