Phonology **23** (2006) 337–385. © 2006 Cambridge University Press
doi:10.1017/S0952675706000984 Printed in the United Kingdom

*Variation as accessing 'non-optimal' candidates**

Andries W. Coetzee
University of Michigan

This paper argues that rather than just select the best candidate, EVAL imposes a harmonic rank-ordering on the full candidate set. Language users have access to this enriched information, and it shapes their performance. This paper applies this idea to variation. The claim is that language users can access the full candidate set via the rank-ordering imposed by EVAL. In variation, more than one candidate is well-formed enough to count as grammatical. Consequently, language users will access more than just the best candidate from the rank-ordering. However, the accessibility of a candidate depends on its position on the rank-ordering. The higher the position a candidate occupies, the more likely it is to be selected. In a variable process, variants that appear higher on the rank-ordering (i.e. are more well-formed) will therefore also be the more frequent variants. This model is applied to variation in the phonology of Faialense Portuguese and Ilokano.

Most linguistic theories in the generative tradition are categorical in nature – they map an input onto a single, grammatical output. Natural language, on the other hand, is full of non-categorical phenomena. For instance, a single word can often be pronounced in more than one way. There is a large literature on variation, both in sociolinguistics and in formal phonological theory. One fact on which all the literature agrees is that variation is not random, but is strongly influenced by grammar. This poses a challenge to classical generative grammar, which is a categorical function by design.

Optimality Theory (Prince & Smolensky 1993) has an advantage over most other generative theories in this regard. In other generative theories, the grammar generates only a single form. However, a basic design feature

* I am grateful to Pam Beddor, John Kingston, John McCarthy, Michael Marlo, Barbara Partee, Joe Pater, David Silva, the participants of Linguistics 792 (University of Michigan, Fall 2004) and the members of the UMass phonology group, as well as the audiences at the LSA Annual Meeting (2004, Boston) and MCWOP 10 (2004, Evanston) for discussion, suggestions and encouragement. An associate editor and several anonymous reviewers also contributed greatly to the improvement of this paper. Special thanks to the editors and the associate editor for their patience, encouragement and help. I am ultimately responsible for all views expressed in this paper.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

of OT is that the grammar generates more than one potential output form. In standard OT, only one of these potential output forms gains the status of an actual output. But the other potential outputs still exist. It is therefore not necessary to add generative power to an OT grammar. What needs to be added to the grammar is a mechanism that will allow, in some circumstances, more than one of the already generated possible output forms to become actual outputs.

In this paper, I present a novel approach to variation in OT that does exactly this. In this model, variation does not arise as a result of variation in grammar (ranking) itself. I argue that EVAL imposes a well-formedness rank-ordering on the full candidate set. In speech production, the language user then has access to the full candidate set via this rank-ordering. The likelihood of a candidate being selected as output depends on the position the candidate occupies on the rank-ordering – the higher a candidate appears, the more likely it is to be selected. In most circumstances, only the topmost candidate on the rank-ordering is well-formed enough to be selected as output. In some situations, however, the top two (or more than two) candidates are both well-formed enough and can both surface as grammatical outputs. The output of the grammar is therefore invariant – the same rank-ordered candidate set is output every time. But how the language user uses this output varies. Sometimes s/he accesses the topmost candidate as output, and sometimes s/he accesses a candidate lower down on the rank-ordering. A fundamental difference between the model developed in this paper and other OT models of variation is the locus of variation. In other models (Reynolds 1994, Anttila 1997, Nagy & Reynolds 1997, Boersma 1998, Boersma & Hayes 2001, Jäger 2003, etc.), variation resides in the grammar itself. The ranking between the constraints can vary from one moment to the next, so that a different candidate can be selected as optimal on different occasions.

This paper is structured as follows. In §1, I discuss the theoretical details of the model that I propose. The next two sections are dedicated to illustrating the application of this model to two sets of variable data. In §2, variable deletion of unstressed vowels in Faialense Portuguese (Silva 1997) is discussed, and in §3, variation in the phonology of Ilokano (Hayes & Abad 1989, Boersma & Hayes 2001). §4 considers outstanding issues and domains other than variation in which the model can be applied.

1 The proposal

The basic proposal of this paper is that EVAL does more than simply select the best candidate. It rather imposes a harmonic rank-ordering on the full candidate set. I will therefore refer to this model as the 'rank-ordering model of EVAL', or ROE. This section explains the details of ROE.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

1.1 Imposing a harmonic ordering on the full candidate set

In classical OT, EVAL is assumed to distinguish the winner from the losers, but not to make distinctions within the set of losers between the more and the less well-formed. As soon as a candidate has been eliminated from the race for optimal status, all of its violations in terms of lower-ranked constraints are ignored – indicated by shading the cells for these constraints in OT tableaux. If the information in the shaded cells is not ignored, then it is possible to order even the non-optimal candidates in terms of their well-formedness. To illustrate the point, imagine a language that allows closed syllables but that avoids tautosyllabic consonant clusters by deletion – i.e. a language with the grammar *COMPLEX ⩢ MAX ⩢ NOCODA. The tableau in (1) shows how this language will evaluate output candidates for an input with both a consonant cluster and a coda.

(1)

/prak/	*COMPLEX	MAX	NOCODA
1. pak		*	*
2. pa		**!	
3. prak	*!		*
4. pra	*!	*	

In classical OT, the grammar will say that [park] is better than all of its competitors, but will be quiet about the well-formedness relationship between the losers – i.e. the grammar imposes the ordering [pak] > {[pa], [prak], [pra]} on the candidate set. However, the information to order the losers is contained in this tableau. If the best candidate is removed, then [pa] will be selected as best. If [pa] is also removed, then [prak] will be best. The candidates can therefore be ordered as follows in terms of their well-formedness: [pak] > [pa] > [prak] > [pra] (indicated by numerals in the tableau).

If the information necessary to rank-order even losers is generated by the grammar but is irrelevant and hence ignored, then the grammar generates a large amount of irrelevant information. This overgeneration problem can be (partially) avoided by not allowing every constraint to evaluate every candidate: the full candidate set is evaluated by the topmost constraint. But this constraint hands down only the candidates that are not disfavoured by itself to the next constraint,[1] which then repeats this mode of evaluation. In this way, candidates eliminated by some constraint are not evaluated by lower-ranking constraints. This is equivalent to assuming that no constraint evaluations are indicated in the shaded cells of a

[1] Candidate $cand_1$ is disfavoured by constraint $\mathbb{C}$ if there is some other candidate, $cand_2$, that earns fewer violations in terms of $\mathbb{C}$ than $cand_1$ – i.e. if $\mathbb{C}(cand_2) < \mathbb{C}(cand_1)$, then $cand_1$ is disfavoured by $\mathbb{C}$ (Samek-Lodovici & Prince 1999). A favoured candidate therefore has to be best, not perfect.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

tableau. In terms of the example in (1), since *COMPLEX has already eliminated [prak] and [pra], these two candidates are not evaluated by lower-ranking constraints – i.e. their violations in terms of MAX and NOCODA should not be marked in the tableau. Samek-Lodovici & Prince (1999) and Prince (2002) develop a formal model of EVAL that has this character. However, even if we do this, the tableau still contains information about the well-formedness relationship between non-optimal candidates. Non-optimal candidates eliminated by lower-ranking constraints are more well-formed than ones eliminated by higher-ranking constraints. For instance, the losing candidates from (1) can still be ordered as follows: [pa] > {[prak], [pra]}.

Because of the ranking between constraints, the grammar generates information about the well-formedness relation between losers even under the assumptions in the previous paragraph. Constraint ranking is one of the most central design principles of an OT grammar, so that it is not possible to prevent the grammar from generating this information without making radical changes to the architecture of the grammar. Rather than trying to get rid of this information, another approach is to embrace it as relevant. However, for it to be relevant, it needs to be shown that language users can access this information and that it shapes their linguistic performance. This is the basic tenet of ROE and is the approach taken in this paper. The basic claims of ROE are as follows:

(i) Constraints are functions with the full candidate set as their domain, so that every candidate is evaluated by every constraint. In this, I follow Moreton (1999) and de Lacy (2002: 30), who define constraints in this manner. See also Coetzee (2004: ch. 2). Samek-Lodovici & Prince (1999) and Prince (2002) develop a very different model of EVAL, in which every constraint $\mathbb{C}$ culls the candidate set down to just those candidates that best satisfy $\mathbb{C}$. $\mathbb{C}$ then hands down only these candidates to the next constraint on the hierarchy. In their model, only the topmost constraint gets to evaluate the full candidate set, so that constraints are not functions.

(ii) EVAL uses all the information generated about the candidate set and imposes a well-formedness rank-ordering on the full candidate set. Language users have access to this rank-ordering, and this access can shape their linguistic performance.

(iii) In production, language users can access more than just the best candidate as output, but the likelihood that a candidate will be selected depends on the position it occupies on the rank-ordering. The higher the candidate is on the rank-ordering (the more well-formed it is), the more likely the language user is to select it as output.

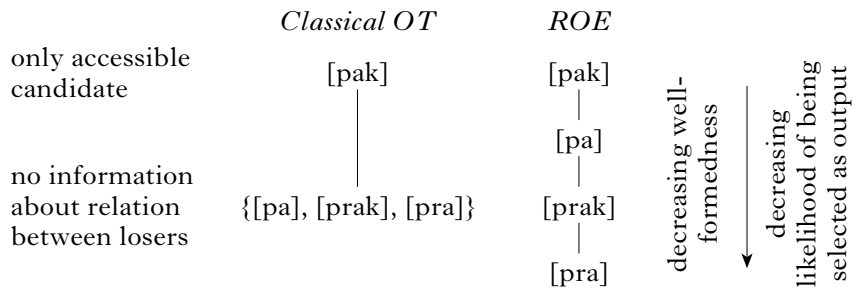
Returning to the example from (1), the best candidate, [pak], is predicted as the most likely candidate to be selected as output for /prak/. But the other candidates are also available and can potentially also be selected. The second-best candidate, [pa], is the second most likely to be selected, and hence the second most frequent variant. [prak], the third candidate on the rank-ordering, is third most likely to be selected, etc. In general,

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

the more well-formed a candidate, the more likely it is to be selected as output and the more frequently it will be observed. It is important to note that no absolute frequency predictions are made – all that is predicted is that a more well-formed candidate will occur more frequently than a less well-formed candidate. In this regard, ROE differs from other theories of variation in OT. For more on this, see the discussion in §4.1.

In (2), a graphic representation of the difference between classical OT and ROE is given. No differentiation is made between the losers in classical OT. The only thing that counts is being the best. In ROE, the losers are differentiated from each other and a candidate can be a better or a worse loser.

(2)



1.2 Limiting variation

Unlike in classical OT, the claim in ROE is that more than just the best candidate is accessible to the language user. Via the rank-ordering imposed by EVAL on the candidate set, the language user has potential access to the full candidate set. Of course, the mere fact that the full candidate set is in principle accessible does not mean that the language user will necessarily always access the candidate set to an arbitrary depth. In fact, it is probably true that for most inputs only the best candidate is ever accessed. And when more than one is accessed, it is usually only the best two or three candidates. What determines when the language user does access the candidate set deeper than just the best candidate? And when the language user does do so, what determines to what depth the candidate set is accessed? Put more generally: what limits variation?

In some theories of variation in OT, there is no principled limit on variation. It is rather the case that (except for harmonically bounded candidates) all candidates are in principle possible outputs, but that for most candidates the probability of being selected is insignificantly small. This is true, for instance, in stochastic OT (Boersma 1998, Boersma & Hayes 2001). The models of Anttila (1997) and Reynolds (1994) are different, in that the non-observed candidates are in principle not possible outputs. In these models, variation arises when a set of

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

constraints are crucially unranked. At every evaluation occasion, one complete ranking between these constraints is selected. If different rankings select different candidates as optimal, then variation is observed. Crucially, only candidates that are selected as optimal under at least one of the rankings are actually possible output candidates. In the stochastic OT model of Boersma & Hayes, there is therefore no concept of absolute ungrammaticality. In this model, ungrammaticality rather becomes equivalent with very low probability. In the Anttila and Reynolds models, the non-observed variants are truly ungrammatical in the sense that they can never be selected as output.

In this section of the paper, I develop a proposal that follows Anttila and Reynolds in that non-observed variants will be classified as candidates that can never be selected as output. The proposal developed here is independent from the general claim of ROE that EVAL imposes a harmonic rank-ordering on the full candidate set. ROE can therefore be accepted without accepting the proposal of this section.

The constraint set is divided into two strata, the one ranked above the other. These two strata are separated by the 'critical cut-off'. The critical cut-off is not a formal object *per se*, but is merely a position on the constraint hierarchy that divides the constraint set into the constraints that rank higher and those that rank lower than the cut-off. The constraints that rank above the cut-off function just like constraints in classical OT. When for some input all candidates are disfavoured by some constraint from above the cut-off, then only the single best candidate (the optimal candidate of classical OT) will be deemed grammatical, and no variation will be observed. However, the constraints that rank below the cut-off at the bottom end of the hierarchy are different. These constraints cannot rule out a candidate as ungrammatical. Put differently: violation of a constraint from below the cut-off is not severe enough to eliminate a candidate from being an output. Variation therefore arises when there is more than one candidate that is disfavoured only by constraints from below the cut-off. All of these candidates are grammatical and are possible outputs. All that the constraints below the cut-off do is impose a harmonic rank-ordering on the candidates, thereby determining the relative frequency with which they will be observed as variant outputs.

There are three logical possibilities with regard to the cut-off, illustrated by the tableaux in (3).[2]

[2] In this and all further tableaux, the position of the cut-off is indicated by a thick vertical line. As in classical OT, pointing hands mark grammatical forms (possible outputs). If there is more than one possible output, the pointing hands are indexed with subscripted numerals. The hand pointing to the most well-formed candidate is indexed with 1. The hand pointing to the second most well-formed candidate is indexed with 2, etc. The indices associated with the pointing hands therefore represent the harmonic rank-ordering that EVAL imposes on the candidate set.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(3) *The influence of the critical cut-off*

a. *Variation:*
more than one candidate disfavoured only by constraints below cut-off

	ℂ1	ℂ2	ℂ3	ℂ4
☞1 i. $cand_1$				*
☞2 ii. $cand_2$			*	
iii. $cand_3$		*!		
iv. $cand_4$	*!			

b. *No variation:*
one candidate disfavoured only by constraints below cut-off

	ℂ1	ℂ2	ℂ3	ℂ4
☞1 i. $cand_1$				*
ii. $cand_2$		*!		
iii. $cand_3$	*!			

c. *No variation:*
no candidate disfavoured only by constraints below cut-off

	ℂ1	ℂ2	ℂ3	ℂ4
☞1 i. $cand_1$		*		
ii. $cand_2$	*!			

In the first scenario, there is more than one candidate that is disfavoured only by constraints that rank below the cut-off. All of these candidates are then possible outputs, while no other candidate is. Tableau (3a) contains a schematic example. In this tableau, *cand₁* and *cand₂* violate $\mathbb{C}_3$ and $\mathbb{C}_4$ respectively. Since these violations are in terms of constraints from below the cut-off, they are not severe enough to eliminate *cand₁* and *cand₂* as possible outputs, so that variation will be observed between these two candidates. The frequency with which each of the variants will be observed is determined by the rank-ordering that EVAL imposes on these candidates. *Cand₁* violates the lower-ranking constraint, and is therefore the more well-formed and the more likely to be accessed as output by the language user. Although both *cand₁* and *cand₂* are possible outputs, *cand₁* is predicted to be observed more frequently. *Cand₃* and *cand₄* are both disfavoured by constraints from above the cut-off. Since there are candidates not disfavoured by constraints from above the cut-off, *cand₃* and *cand₄* are both ungrammatical and will never be accessed by the language user as output. This paper focuses on variation, and many examples of this scenario will therefore be seen in the rest of the paper.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

Under the second possible scenario, there is only one candidate that is disfavoured only by constraints from below the cut-off. This candidate is then the only possible output, so that no variation will be observed. Tableau (3b) gives a schematic example. In this tableau, $cand_2$ and $cand_3$ are both disfavoured by constraints from above the cut-off, and since there is a candidate, $cand_1$, that is not disfavoured by a constraint from above the cut-off, $cand_2$ and $cand_3$ are deemed ungrammatical. $Cand_1$ is hence the only grammatical candidate and the only possible output. This scenario is equivalent to that in (3a), except that there is only one variant.

The last possible scenario is where all candidates are disfavoured by at least one constraint that ranks above the cut-off. In this case, ROE functions exactly like classical OT – all candidates except for the best candidate are ungrammatical. Tableau (3c) gives a schematic example. Both $cand_1$ and $cand_2$ are disfavoured by a constraint from above the cut-off. However, since there is no candidate that is disfavoured only by constraints from below the cut-off, one of $cand_1$ and $cand_2$ has to be chosen as output. The best one, in this case $cand_1$, is selected. In the discussion of Ilokano in §3, examples of the scenarios illustrated in (3b) and (3c) will be given.

The intuition behind the cut-off can also be stated as follows: if the best candidate is relatively very well-formed (if it is disfavoured only by very low-ranking constraints), then it is possible that there are some other candidates that, although not the best, are still very well-formed and can therefore count as grammatical. However, if the best candidate is not that well-formed (it is disfavoured by high-ranking constraints), then all other candidates will be even less well-formed. The non-best candidates are then too ill-formed to count as grammatical.[3]

There is no principled way of deciding whether the cut-off does exist. This is ultimately an empirical question – i.e. can all observed patterns of variable and categorical phenomena be explained in a model of ROE that includes a critical cut-off? In the next two sections, I will show that viable analyses of variation in Faialense Portuguese and in Ilokano can be given within ROE with the critical cut-off. This will give some credence to the existence of the cut-off. However, more research is required before it can be decided definitively either way.

2 Variable vowel reduction and deletion in Faialense Portuguese

In this section, I develop an account of variable vowel reduction and deletion in Faialense Portuguese within ROE. In §2.1, I discuss the basic vocalic phonology of Faialense Portuguese. §2.2 then presents data about

[3] See also §4.2 on the exclusion of harmonically bounded candidates from the set of possible variants.

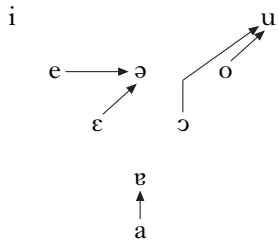
https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

the variation between reduction and delection, and the rest of §2 shows how these data can be explained in ROE.

2.1 The basic vocalic phonology of Faialense Portuguese

Faial is an Azorean island in the mid-Atlantic. Although Faialense Portuguese differs from mainland Portuguese in its consonantal phonology, the vocalic phonologies of these two varieties of Portuguese are virtually identical (Rogers 1949: 48, Mateus & d'Andrade 2000: 2). Both have the vocalic inventory in (4) (cf. Parkinson 1988, Silva 1997: 298). The full inventory occurs only in stressed syllables. A vowel with an arrow pointing away from it does not occur in unstressed syllables, but is replaced with the vowel to which the arrow points (more below).[4]

(4) *Vowel inventory of Faialense Portuguese*



Faialense Portuguese has three levels of vowel height (high = [i u]; mid = [e o ɛ ɔ ə]; low = [ɐ a]) and three degrees of backness (front = [i e ɛ]; central = [ə ɐ a]; back = [u o ɔ]). The members of the pairs [e ɛ], [o ɔ] and [ɐ a] are distinguished by the feature [ATR] ([e o ɐ] = [+ATR] and [ɛ ɔ a] = [−ATR]). [i u] are also treated as [+ATR], and [ə] as [−ATR].

Mainland and Faialense Portuguese also show the same reduction pattern of unstressed vowels (Silva 1997: 298). This reduction pattern is shown by the arrows in (4),[5] with examples given in (5).[6]

[4] Mateus & d'Andrade (2000: 17–18) use [ɨ] rather than [ə] for the non-low, central vowel.

[5] Silva (1997: 18, 1998: 175) and Mateus & d'Andrade (2000: 18) claim that /i/ reduces to [ĭ] in post-tonic position. However, stress placement in Portuguese is such that /i/ never occurs in post-tonic position. Portuguese usually stresses the penultimate syllable, but stress is attracted to a final syllable with /i/ (Thomas 1974: 3, Silva 1997: 299). The result is that /i/ is either stressed or pre-tonic. This is also evidenced in Silva's corpus, in which no word-final unstressed /i/ occurs – see (7) below.

[6] It is unclear whether /ɐ/ has phonemic status in FP or not – the only contexts where it occurs in stressed position is before a palatal segment, as in the example in the text (David Silva, personal communication).

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(5) *Reduction in unstressed syllables*

/e/ → [ə̆]	*peso*	[ˈpe̲su]	‘weight’	*pesado*	[pə̲ˈsadu]	‘heavy’
/ɛ/ → [ə̆]	*selo*	[ˈse̲lu]	‘I stamp’	*selar*	[sə̲ˈlaɾ]	‘to stamp’
/a/ → [ɐ̆]	*paga*	[ˈpa̲gɐ]	‘s/he pays’	*pagar*	[pɐ̲ˈgaɾ]	‘to pay’
/ɔ/ → [ŭ]	*forço*	[ˈfɔ̲ɾsu]	‘I oblige’	*forçar*	[fu̲ɾˈsaɾ]	‘to oblige’
/o/ → [ŭ]	*moro*	[ˈmo̲ɾu]	‘I live’	*morar*	[mu̲ˈɾaɾ]	‘to live’
/i/ → [ĭ]	*livro*	[ˈli̲vɾu]	‘book’	*livraria*	[li̲vɾɐˈɾiɐ]	‘bookshop’
/u/ → [ŭ]	*fuga*	[ˈfu̲gɐ]	‘I escape’	*fugar*	[fu̲ˈgaɾ]	‘to escape’
/ɐ/ → [ɐ̆]	*telha*	[ˈtɐ̲ʎɐ]	‘tile’	*telheiro*	[tɐ̲ˈʎejɾu]	‘tiler’

2.2 The data on variable vowel deletion in Faialense Portuguese

Some varieties of Portuguese variably delete vowels from unstressed syllables – see the examples in (6).

(6) *Variable vowel deletion in Faialense Portuguese*

[ə̆] ~ ∅	*peludo*	‘hairy’	[pə̆ˈludŭ]	~ [pˈludŭ]
	idade	‘idea’	[ĭˈdadə̆]	~ [ĭˈdad]
[ŭ] ~ ∅	*mulher*	‘woman’	[mŭˈʎɛɾ]	~ [mˈʎɛɾ]
	tempo	‘time’	[ˈtempŭ]	~ [ˈtemp]
[ĭ] ~ ∅	*piloto*	‘pilot’	[pĭˈlotŭ]	~ [pˈlotŭ]
[ɐ̆] ~ ∅	*pagar*	‘to pay’	[pɐ̆ˈgaɾ]	~ [pˈgaɾ]
	paga	‘s/he pays’	[ˈpagɐ̆]	~ [ˈpag]

Most literature on Portuguese phonology mentions this phenomenon, but does not report deletion rates or the factors that influence the frequency of deletion (Parkinson 1988: 142, Mateus & d’Andrade 2000: 18). The one exception is Silva (1997), who reports in detail on vowel deletion in Faialense Portuguese.[7] Silva recorded a 22-minute conversation between two native female speakers of Faialense Portuguese, and analysed the utterances of one of the participants. He coded his data according to several linguistic factors, and submitted it to a VARBUL analysis (Sankoff 1988). Of the factors that he considered, three were selected as contributing toward the variation pattern: (i) vowel quality, (ii) position in prosodic word and (iii) stress of the following syllable. The last factor was identified as playing a more marginal role, and I will focus only on the

[7] Silva (1998) also reports on vowel deletion in São Miguel Portuguese. However, São Miguel is an eastern Azorean dialect, and these dialects differ more from mainland Portuguese in their vocalic phonology (Rogers 1949: 48, Silva 1997: 30). I limit myself to Faialense Portuguese in this paper.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

		position in prosodic word		
		final	elsewhere	total
[ə]	deleted	32	36	68
	retained	13	82	95
	% deleted	71%	31%	42%
[ŭ]	deleted	48	16	64
	retained	22	96	118
	% deleted	69%	14%	35%
[ĭ]	deleted	—	5	5
	retained	—[8]	70	70
	% deleted	—	7%	7%
[ĕ]	deleted	2	5	7
	retained	131	193	324
	% deleted	2%	3%	2%
total	deleted	82	62	
	retained	166	441	
	% deleted	33%	14%	

Table I
Vowel-deletion patterns in Faialense Portuguese.

first two in this paper. Table I summarises the relevant aspects of Silva's data.

Silva reports his data based on surface vowel quality, and on the position that the vowel would have occupied in the prosodic word had it not deleted (final syllable or non-final syllable). The data can be characterised as follows in terms of this division: (i) for [ĕ], retention is favoured over deletion in both prosodic contexts. (ii) for [ə] and [ŭ], deletion is favoured over retention when final in the prosodic word. However, elsewhere in the prosodic word, retention is favoured over deletion. (iii) [ĭ] does not occur final in the prosodic word. Elsewhere in the prosodic word, [ĭ] patterns together with [ə] and [ŭ] by showing more retention than deletion. In the rest of this section, I develop and account for this pattern in ROE.

[8] Silva has ten [ĭ]'s in this cell. However, he has explained (personal communication) that these are all instances of the conjunction *e* 'and', and that he incorrectly treated this conjunction as a separate prosodic word in his 1997 paper. Function words are incorporated into the prosodic word headed by the following lexical word. These ten [ĭ]'s should rather be counted as occurring in a syllable that is non-final in a prosodic word. In Table I, the 70 [ĭ]'s in the next cell include the ten that were in this cell in Silva's original.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

2.3 The constraints

In the analysis of Faialense Portuguese vowel reduction, I follow Crosswhite's proposals on how to model vowel reduction in OT (Crosswhite 2000, 2001). Based on an extensive typological review, she shows that vowel reduction is often motivated by a drive to avoid high-sonority vowels in unstressed syllables. Following Crosswhite, as well as de Lacy (2002) and Parker (2002), I assume that the centralised vowels [ɐ ə] are lowest in sonority, and that for [i u e o ɛ ɔ a] sonority decreases as height increases. We can then order the Faialense Portuguese vowels as follows in terms of sonority: [a] > [e o ɛ ɔ] > [i u] > [ɐ ə]. Comparison of this scale with the reduction pattern in (4) shows that all of the reductions involve replacement of a higher-sonority vowel with a lower-sonority vowel. To analyse this kind of vowel reduction, Crosswhite uses a set of markedness constraints against parsing specific vowels into unstressed syllables, and claims that these constraints are universally ranked, with constraints against more sonorous vowels outranking those against less sonorous vowels. The versions of these constraints relevant for Faialense Portuguese vowel reduction are given in (7).

(7) $*\breve{\sigma}/a \gg *\breve{\sigma}/\{e,o,\varepsilon,ɔ\} \gg *\breve{\sigma}/\{i,u\} \gg *\breve{\sigma}/\{ɐ,ə\}$

These markedness constraints make no reference to the position of the vowel in the prosodic word, and yet the data in Table I show that the vowel's position in the prosodic word does play a role in vowel deletion. Specifically, prosodic word-final vowels are more prone to deletion than vowels elsewhere in the prosodic word. Crosswhite argues that vowels are less well licensed in prosodically weak than strong positions. An unstressed syllable in prosodic word-final position is arguably prosodically weaker than an unstressed syllable anywhere else in the prosodic word. Having any vowel in such a syllable is hence more marked than having a vowel in any other unstressed syllable. Based on this, I argue for a markedness constraint, $*\breve{V}]_{\omega}$, which penalises unstressed prosodic word-final vowels.[9] $*\breve{V}]_{\omega}$ is related to FINALC (McCarthy & Prince 1990) and

[9] An anonymous reviewer suggests CONTIGUITY as an explanation for the higher deletion rate in final position. Deletion from word-final position does not violate CONTIGUITY, while deletion from elsewhere in a word does (cf. Kenstowicz 1994, Landman 2002). Although it is likely that CONTIGUITY plays a role, there are reasons why $*\breve{V}]_{\omega}$ would still be needed. Silva also coded prosodic word-initial vowels as 'elsewhere'. Deletion from this position also does not violate CONTIGUITY. Assuming that he was correct in grouping word-initial and medial vowels together, it is clear that a constraint is needed that targets specifically word-final vowels. The reviewer also suggests that syllabic well-formedness constraints such as *COMPLEX could be at play. Silva's results show that this is most likely not the case. He coded his data also for whether deletion results in a licit or illicit syllable. The VARBUL analysis showed that this factor had no influence on the likelihood of deletion (Silva 1997: 304).

A reviewer also points out that it is surprising for final vowels to delete more than non-final vowels – final vowels are often inflectional morphemes and deletion of

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

FREEV (Prince & Smolensky 1993: ch. 7). All of these constraints penalise morphophonological domains that end in vowels.

(8) *V̆]$_{\omega}$
Do not parse a vowel into an unstressed prosodic word-final syllable.

Also following Crosswhite, I use MAX/DEP[αF] rather than IDENT[F] featural faithfulness constraints (Pater 1999, Lombardi 2001), so that there will be a MAX/DEP[+F] and MAX/DEP[−F] constraint for each of the vocalic features [high, low, back, front, ATR]. Lastly, I also use MAX(seg). Segmental deletion will therefore violate MAX(seg) in addition to all the MAX[αF] constraints that apply to the specific vowel.

2.4 The analysis

The central claim of ROE is that the more well-formed a candidate is, the more likely it is to be selected as output. This idea will drive the discussion in the rest of this section. The first thing that needs to be shown is that, for each input, the more frequent variant is more well-formed (occupies a higher position on the rank-ordered candidate set) than the less frequent variant. Next it needs to be shown that the two observed variants are more well-formed than all other candidates. The analysis therefore must impose the following structure on the candidate set for each input: *variant*$_1$ > *variant*$_2$ > the rest. Finally, the position of the critical cut-off on the constraint hierarchy needs to be determined. Specifically, the cut-off needs to be located in such a position that neither of the observed variants is disfavoured by a constraint above the cut-off, while all other candidates are.

2.4.1 *Non-final syllables in the prosodic word.* Table I showed that for all four possible outputs, the non-deletion candidate is more frequent than the deletion candidate when not in prosodic word-final position. In the rank-ordering that EVAL imposes on the candidate set, the non-deletion candidate must therefore occupy a higher position than the deletion candidate. This can be achieved if the highest-ranked constraint that distinguishes the two candidates favours non-deletion over deletion. In Table II, I list the constraints violated by the two observed variants for

these vowels therefore results in the loss of morphemes (i.e. incurs a violation of REALISE-morpheme). Silva does not report that this has an inhibitory influence on the deletion of final vowels, and this most likely means that *V̆]$_{\omega}$ outranks REALISE-morpheme. It is also noteworthy that loss of word-final inflectional affixes is a very common historical process. Modern Arabic dialects have lost nearly all the word-final inflectional vowels of Classical Arabic. Modern Romance languages have lost many of Latin's inflectional vocalic suffixes. English has lost most inflectional affixes of Proto-Germanic, etc.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

/ɔ/	[ŭ]	*$\breve{\sigma}$/{i,u}, MAX[−high], MAX[−ATR], DEP[+high], DEP[+ATR]
	∅	MAX(seg), MAX[−high], MAX[−ATR], MAX[−low], MAX[+back]
/o/	[ŭ]	*$\breve{\sigma}$/{i,u}, MAX[−high], DEP[+high]
	∅	MAX(seg), MAX[−high], MAX[−low], MAX[+back], MAX[+ATR]
/u/	[ŭ]	*$\breve{\sigma}$/{i,u}
	∅	MAX(seg), MAX[−low], MAX[+back], MAX[+high], MAX[+ATR]
/ɛ/	[ə̆]	*$\breve{\sigma}$/{ɐ,ə}, MAX[+front], DEP[−front]
	∅	MAX(seg), MAX[+front], MAX[−high], MAX[−low], MAX[−ATR]
/e/	[ə̆]	*$\breve{\sigma}$/{ɐ,ə}, MAX[+front], MAX[+ATR], DEP[−front], DEP[−ATR]
	∅	MAX(seg), MAX[+front], MAX[+ATR], MAX[−high], MAX[−low]
/i/	[ĭ]	*$\breve{\sigma}$/{i,u}
	∅	MAX(seg), MAX[+front], MAX[+high], MAX[+ATR]
/a/	[ɐ̆]	*$\breve{\sigma}$/{ɐ,ə}, MAX[−ATR], DEP[+ATR]
	∅	MAX(seg), MAX[−ATR], MAX[+low], MAX[−high]
/ɐ/	[ɐ̆]	*$\breve{\sigma}$/{ɐ,ə}
	∅	MAX(seg), MAX[+low], MAX[−high], MAX[−ATR]

Table II

Violation profiles of observed outputs. Violations which are shared between the two variants, and therefore cannot distinguish between them, are underlined.

every input. Violations that are shared between the two variants are underlined, to indicate that these constraints cannot distinguish between the two variants.

As Table II shows, MAX(seg) is the only violation shared by the deletion variant for all of the inputs. If MAX(seg) ranks higher than all the constraints that disfavour a non-deletion variant, non-deletion will be better than deletion for all inputs. The same result can be achieved by using a combination of the MAX[αF] constraints. However, since there is no single MAX[αF] constraint violated by all deletion candidates, several MAX[αF] constraints will have to be used. I therefore opt for MAX(seg) rather than MAX[αF] constraints. In accordance with the principle of ranking conservatism, all faithfulness constraints are ranked at the bottom of the hierarchy in the absence of evidence to the contrary (Tesar & Smolensky 1998, 2000, Itô & Mester 1999, 2003). This gives the ranking represented in (9). (See (7) for the ranking *$\breve{\sigma}$/{i,u} ≫ *$\breve{\sigma}$/{ɐ,ə}.) With this ranking, the deletion candidate will always violate a higher-ranking constraint than the non-deletion candidate. Under the assumption that the more well-formed

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

variant is the more frequent variant, non-deletion will then be observed more frequently than deletion.

(9) *Interim ranking for Faialense Portuguese*

MAX(seg)
|
*$\breve{\sigma}$/{i,u}
|
*$\breve{\sigma}$/{ɐ,ə}
|
DEP[+high], DEP[+ATR], MAX[−low], MAX[+back], DEP[−front], DEP[−ATR], MAX[+high], MAX[+ATR], MAX[+low], MAX[−high], MAX[−ATR], MAX[+front]

Under the ranking in (9), EVAL imposes the ordering 'non-deletion > deletion' on the variants for all inputs. Next, it needs to be ensured that the two observed variants are the best candidates – i.e. all other vowels must be less well-formed than even deletion. This will be the case if each non-observed candidate violates a constraint higher than MAX(seg). Since each of the non-observed candidates violates several constraints, there are many constraints that can be used to achieve this goal. The grammar developed below is therefore only one of the possible grammars. In order to keep the grammar simple, I use as few constraints as possible to rule out the non-observed variants.

The back vowel inputs /u o ɔ/ all map onto [ŭ] or ∅. All vowels except for [ŭ] must therefore violate some constraint ranked higher than MAX (seg). Mapping any of /u o ɔ/ onto [ŏ] or [ɔ̆] violates at least *$\breve{\sigma}$/{e,o,ɛ,ɔ}, so that ranking *$\breve{\sigma}$/{e,o,ɛ,ɔ} above MAX(seg) ensures that [ŏ] and [ɔ̆] are less well-formed than ∅ and [ŭ]. This leaves the non-back vowels. Mapping /u o ɔ/ onto a non-back vowel violates MAX[+back] and DEP [−back]. Since ∅ also violates MAX[+back], we cannot use MAX[+back] to distinguish between ∅ and the non-back vowels. However, ranking DEP[−back] above MAX(seg) will ensure that the non-back vowels are less well-formed than ∅ and [ŭ]. For /u o ɔ/, the following rankings then need to be added to those in (9): {*$\breve{\sigma}$/{e,o,ɛ,ɔ}, DEP[−back]} ≫ MAX(seg). The tableau in (10) shows how the grammar developed thus far evaluates the candidate outputs for the inputs /u o ɔ/. The tableau also shows the position of the critical cut-off on the constraint hierarchy (marked by the thick vertical line). To ensure that both ∅ and [ŭ] are allowed to surface, these candidates can only be disfavoured by constraints below the cut-off, telling us that the cut-off has to rank above MAX(seg). All other candidates must be disfavoured by some constraint above the cut-off, so that both *$\breve{\sigma}$/{e,o,ɛ,ɔ} and DEP[−back] must rank above the cut-off. (This tableau contains only the relevant constraints – i.e. the faithfulness constraints ranked at the bottom are not included.)

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(10) /u o ɔ/ → | ŭ > ∅ > other V̆ |

/u o ɔ/	*σ̆/{e,o,ɛ,ɔ}	Dep[−bk]	Max(seg)	*σ̆/{i,u}
☞1 a. ŭ				*
☞2 b. ∅			*	
c. ŏ ŏ	*!			
d. other V̆		*!		

The vowels /ɐ a/ map onto [ĕ] or ∅. To ensure that these two candidates are the best, all other candidates must violate a constraint above Max(seg). Mapping /ɐ a/ onto [ă] violates *σ̆/a, so that *σ̆/a ≫ Max(seg) can eliminate [ă]. Mapping /ɐ a/ onto a non-low vowel violates at least Max[+low] and Dep[−low]. Max[+low] cannot distinguish the non-low vowels from ∅, since ∅ also violates Max[+low]. But ranking Dep[−low] above Max(seg) ensures that the non-low vowels are less well-formed than ∅ and [ĕ]. Also, as in (10), the cut-off needs to be located just above Max(seg), ensuring that the observed variants are not disfavoured by a constraint from above the cut-off, while all other candidates are. Tableau (11) shows how the output candidates for /ɐ a/ inputs are evaluated.

(11) /ɐ a/ → | ĕ > ∅ > other V̆ |

/ɐ a/	*σ̆/a	Dep[−low]	Max(seg)	*σ̆/{ɐ,ə}
☞1 a. ĕ				*
☞2 b. ∅			*	
c. ă	*!			
d. other V̆		*!		

/i/ maps onto [ĭ] or ∅. All other candidates must violate a constraint above Max(seg). Mapping /i/ onto a non-high vowel violates Max[+high] and Dep[−high]. As before, the Max constraint cannot distinguish between the other vowels and ∅, so that Dep[−high] has to be ranked above Max(seg). The only candidate that remains is [ŭ]. Since [ŭ] is a high vowel, /i/→[ŭ] does not violate Dep[−high]. However, it does violate Max[+front], Dep[−front] and Dep[+back]. Max[+front] cannot be used, since it is also violated by ∅. There is evidence from elsewhere in the system that Dep[−front] cannot rank above Max(seg): the variant [ə̆] for the inputs /e ɛ/ violates Dep[−front], and since [ə̆] is more frequent than ∅ for these inputs, it follows that Dep[−front] must rank below Max(seg). However, there is nothing that prevents the ranking of Dep[+back] above Max(seg). The tableau in (12) shows how the ranking {Dep[−high], Dep[+back]} ≫ Max(seg) ensures that ∅ and [ĭ] are the most well-formed candidates for /i/. Furthermore, as before, with the cut-off located just

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

above MAX(seg), it is ensured that [ĭ] and ∅ will be the only observed variants.

(12) /i/ → | ĭ > ∅ > other V̆ |

/i/	DEP[−hi]	DEP[+bk]	MAX(seg)	*σ̆/{i,u}
☞1 a. ĭ				*
☞2 b. ∅			*	
c. ŭ		*!		*
d. other V̆	*!			

The last vowels to consider are /e ɛ/, which map onto [ə̆] or ∅. All other possible candidates must be disfavoured by a constraint above MAX(seg). Mapping /e ɛ/ onto [ĕ] or [ɛ̆] violates *σ̆/{e,o,ɛ,ɔ}, and mapping them onto any back vowel violates DEP[+back]. It has already been established that these two constraints rank above MAX(seg). Mapping /e ɛ/ onto a low vowel, [ɐ̆] or [ă], violates at least MAX[+front], MAX[−low], DEP[+low] and DEP[−front]. Since ∅ also violates the MAX constraints, they cannot be used to differentiate the low vowels from ∅. Similarly, since [ə̆] also violates DEP[−front], it cannot be used to distinguish between [ə̆] and the low vowels. However, ranking DEP[+low] above MAX(seg) will ensure that [ɐ̆] or [ă] is less well-formed than ∅ and [ə̆]. Thus far the following ranking has been motivated: {*σ̆/{e,o,ɛ,ɔ}, DEP[+back], DEP[+low]} ≫ MAX(seg). This ranking ensures that ∅ and [ə̆] are more well-formed than [ĕ ɛ̆ ŏ ɔ̆ ŭ ă ɐ̆]. This leaves only [ĭ] to be accounted for.

The mapping /e ɛ/→[ĭ] violates *σ̆/{i,u}, DEP[+high] and MAX[−high]. These are the same constraints violated by /o ɔ/→[ŭ]. Since [ŭ] is observed more frequently than ∅ as output for /o ɔ/, the constraints that disfavour /o ɔ/→[ŭ] must rank below MAX(seg). It is therefore not possible to rank any of *σ̆/{i,u}, DEP[+high] or MAX[−high] above MAX(seg). With just the constraints used up to now, it cannot be explained why /e ɛ/→[ĭ] should be less well-formed than /e ɛ/→[ə̆] ~ ∅. What is needed is an additional constraint violated by /e ɛ/→[ĭ], but not by /o ɔ/→[ŭ]. I will use a constraint that targets specifically high, front vowels (*[+high, +front]).[10] If *[+high, +front] ranks above MAX(seg), [ĭ] will be less well-formed than [ə̆] and ∅ as output candidates for /e ɛ/.

[10] Based on a typological review of epenthetic vowels, Lombardi (2003) claims that front vowels are universally more marked than non-front vowels, and that therefore *[+front] ≫ *[−front]. She also claims that languages differ in whether high or non-high vowels are more marked – i.e. *[+high] and *[−high] are not in a fixed ranking. In a language with the ranking *[+high] ≫ *[−high], the low vowel [a] is more likely to be the epenthetic vowel. In a language with the opposite ranking, [i] is more likely as epenthetic vowel. Portuguese does not have productive vowel epenthesis, so that it is not possible to determine with certainty the ranking between *[+high] and *[−high]. However, there is evidence from the acquisition of onset clusters for the ranking *[+high] ≫ *[−high]. Freitas (2003: 35) shows that children sometimes resolve onset clusters via vowel epenthesis. The epenthetic vowel that they use is

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

However, the addition of *[+high, +front] does not solve the problem completely. The input /i/ has [ĭ] and ∅ as variants, with [ĭ] more frequent and hence more well-formed. Consequently, all constraints violated by [ĭ], including *[+high, +front], must rank lower than MAX(seg). The ranking *[+high, +front] ⩾ MAX(seg) is necessary to exclude [ĭ] as an output for /e ɛ/. But to ensure that [ĭ] is more well-formed than ∅ for the input /i/, the opposite ranking, MAX(seg) ⩾ *[+high, +front], is necessary. This is a classic example of a derived environment blocking effect: [ĭ] from a faithful mapping from /i/ is tolerated, but [ĭ] from /e ɛ/ is not. There are several proposals for how to model phenomena such as these in OT. One option is to use conjunction of a markedness and faithfulness constraint (Łubowicz 2002). For instance, *[+high, +front] can be conjoined with MAX[+high] or DEP[−high]. The non-observed mapping /e ɛ/→[ĭ] will then violate this conjoined constraint, while the observed mapping /i/→[ĭ] will not. Another option is to use comparative markedness constraints (McCarthy 2002, 2003). Under this proposal, there are two versions of each markedness constraint. The $*M_{old}$ version is violated by a marked structure that was already present in the input, while the $*M_{new}$ version is violated by a marked structure that was not present in the input. $*[+high, +front]_{old}$ will therefore be violated by the mapping /i/→[ĭ], while /e ɛ/→[ĭ] will violate $*[+high, +front]_{new}$. In the rest of the discussion, I will use comparative markedness, although local conjunction would have worked equally well for the data at hand.[11] By ranking $*[+high, +front]_{new}$ above MAX(seg), we ensure that [ĭ] is less well-formed than ∅ and [ə̆] for the inputs /e ɛ/. $*[+high, +front]_{old}$, on the other hand, has to rank below MAX(seg), so that [ĭ] is still more well-formed than ∅ for the input /i/.

The ranking {$*[+high, +front]_{new}$, DEP[+back], DEP[+low]} ⩾ MAX(seg) ⩾ *σ̆/{ɐ,ə} will therefore impose the correct rank-ordering on the candidate set for the inputs /e ɛ/. All that remains is to determine the placement of the cut-off, and as in all the preceding tableaux in this section, it is clear that it needs to be placed just above MAX(seg). The tableau in (13) illustrates the derivation of inputs with /e ɛ/ that are parsed into unstressed syllables.

---

either [ə] or [ɐ], never [i]. Since front vowels are universally more marked, and since high vowels are probably more marked than low vowels in Portuguese, it follows that the high front vowel [i] is the most marked in Portuguese.

[11] The need for either [M&F]-conjunction or comparative markedness is not unique to ROE. Any analysis needs to account for the fact that derived [ĭ] (from /e ɛ/) is not allowed, while non-derived [ĭ] (from /i/) is. Any analysis will therefore have to use one of the devices developed to deal with this kind of derived environment effect.

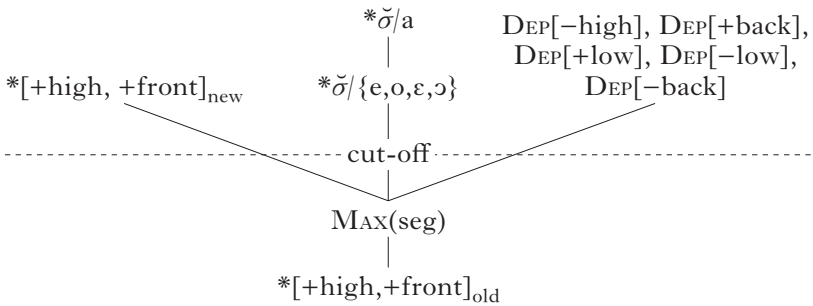
https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(13) /e ɛ/ → | ə > ∅ > other V̆ |

/e ɛ/	*[+hi,+fr]$_{new}$	Dep[+bk]	Dep[+low]	Max(seg)	*σ̆/{ɐ,ə}
☞$_1$ a. ə					*
☞$_2$ b. ∅				*	
c. ɐ̆ ă			*!		
d. ŭ ŏ ŭ		*!			
e. ĭ	*!				

This takes care of all the vowels in non-final position in the prosodic word. In (14), I show the rankings that have been motivated since the summary in (9). In the grammar motivated so far, (i) the two observed variants in each case are rated better than all other candidates, (ii) of the two variants, the non-deletion variant is rated better than the deletion variant and (iii), except for the observed variants, all candidates are disfavoured by some constraint that rank above the cut-off.

(14) *New rankings since* (9)



2.4.2 *Final syllables in the prosodic word.* The same variants occur in prosodic word-final position as elsewhere in the prosodic word. What is different is that the relative frequency of the variants is inverted for all but the low vowels. In this prosodic position, all candidates violate the same constraints as elsewhere in the prosodic word, with one exception: non-deletion candidates violate an additional constraint, *V̆]$_{\omega}$. For the most part, the variation pattern in prosodic word-final position can be explained simply by the addition of *V̆]$_{\omega}$ to the hierarchy.

For the non-low vowels in prosodic word-final position, the grammar has to rate the deletion candidate better than the non-deletion candidate. This can be achieved by ranking *V̆]$_{\omega}$ above Max(seg). Non-deletion now violates a higher-ranking constraint and is consequently rated as less well-formed. Of course, since the deletion variant violates *V̆]$_{\omega}$, this constraint also has to rank below the cut-off. The cut-off can therefore no longer rank just above Max(seg), as in (14). Tableau (15) illustrates

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

the derivation of back vowels parsed into unstressed prosodic word-final position. The only difference between the constraints used in (15) and those used in (10) above is the insertion of *V̆]$_{\omega}$ between Max(seg) and the cut-off. In (15), [ŭ#] and ∅# are still the best candidates, but ∅# is now better than [ŭ#]. The other non-low vowels can be accounted for in the same way, and hence I do not give separate tableaux for these vowels.

(15) /u# o# ɔ#/ → |∅# > ŭ# > other V̆|

/u# o# ɔ#/	*σ̆/{e,o,ɛ,ɔ}	Dep[−bk]	*V̆]$_{\omega}$	Max(seg)	*σ̆/{i,u}
☞$_1$ a. ∅#				*	
☞$_2$ b. ŭ#			*		*
c. ŏ# ŭ#	*!		*		
d. other V̆#		*!	*		

For the low vowels, the non-deletion variant [ĕ#] is more frequent in prosodic word-final position as well. The current ranking, *V̆]$_{\omega}$ ≫ Max (seg), incorrectly rates the deletion candidate ∅# better than the non-deletion candidate [ĕ#]. The solution to this problem is to rank one of the other constraints violated by /ɐ# a#/→∅# higher than *V̆]$_{\omega}$. In addition to Max(seg), /ɐ# a#/→∅# also violates Max[+low], Max[−high], Max[−front] and Max[−back]. Except for Max[+low], each of the other Max[αF] constraints is violated by deletion of some non-low vowel.[12] For these non-low vowels, deletion is more frequent than retention, and the constraints violated by deletion must therefore rank below *V̆]$_{\omega}$. Hence, Max[−high], Max[−front] and Max[−back] cannot outrank *V̆]$_{\omega}$. But since only deletion of a low vowel violates Max[+low], this constraint can be ranked above *V̆]$_{\omega}$ without influencing non-low vowel inputs. The tableau in (16) shows how the ranking Max[+low] ≫ *V̆]$_{\omega}$ ensures that non-deletion is preferred over deletion for low vowels even in prosodic word-final position. Also, the cut-off has to move one spot higher to just above Max[+low] in order to ensure that the deletion candidate violates only constraints below the cut-off.

(16) /ɐ# a#/ → |ĕ# > ∅# > other V̆|

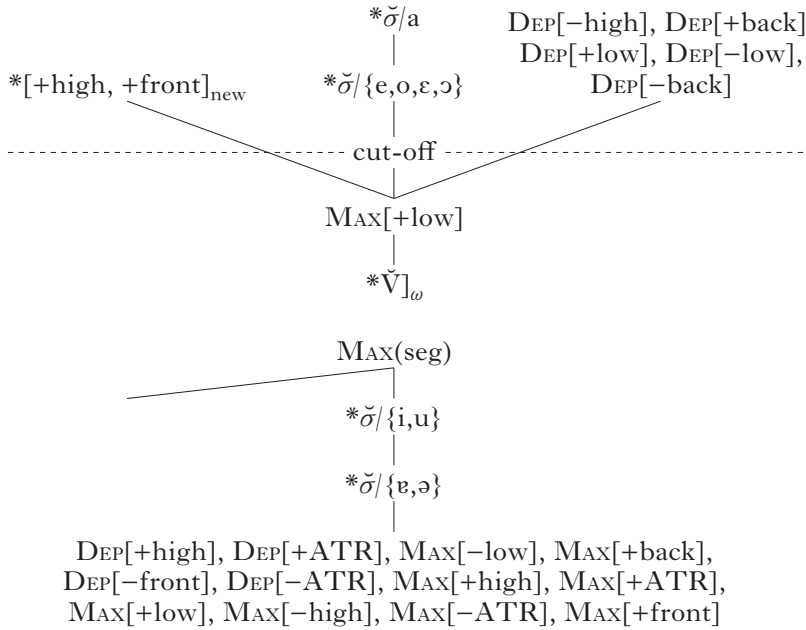
/ɐ# a#/	*σ̆/a	Dep[−low]	Max[+low]	*V̆]$_{\omega}$	Max(seg)	*σ̆/{ɐ,ə}
☞$_1$ a. ĕ#				*		*
☞$_2$ b. ∅#			*		*	
c. ă#	*!			*		
d. other V̆#		*!	*	*		

[12] Max[−high] is violated by /e# o# ɛ# ɔ#/→∅#. Max[−front] is violated by /u# o# ɔ#/→∅#. Max[−back] is violated by /i# e# ɛ#/→∅#.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

2.4.3 *Summary*. All of the rankings necessary to account for the basic variation pattern in Faialense Portuguese are now in place. This ranking is represented graphically in (17).

(17) *Final ranking for Faialense Portuguese*



With this ranking, non-deletion is always better than deletion when not in prosodic word-final position. This follows from the ranking MAX(seg) ⩾ *σ̆/{i,u} ⩾ *σ̆/{ɐ,ə}. In prosodic word-final position, deletion is better than retention for all except the low vowels. This follows from the ranking *V̆]ω ⩾ MAX(seg). For the low vowels, the ranking MAX[+low] ⩾ *V̆]ω ensures that non-deletion is better than deletion even in prosodic word-final position. Finally, the location of the cut-off just above MAX[+low] ensures that all observed variants are only disfavoured by constraints from below the cut-off, while all other candidates are disfavoured by at least one constraint from above the cut-off.

2.5 Evaluation of and possible extensions to the analysis

2.5.1 *Multiple deletion*. The analysis of Faialense Portuguese above is presented in terms of schematic, abstract examples – the inputs and candidate outputs all consist of single vowels. In actual language use, the vowels are of course always embedded in full words. The situation is more complex when full words are considered, especially for words that contain multiple unstressed vowels. In this section, I show the predictions that the

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

analysis makes for such longer words, using the word *peludo* 'hairy' as example. The tableau in (18) shows how this word will be evaluated in the grammar developed above. To simplify, I consider only candidates that do not violate any constraints ranked higher than the cut-off.

(18) *Evaluating longer words*

/peludo/	*V̆]$_{\omega}$	MAX(seg)	*σ̆/{i,u}	*σ̆/{ɐ,ə}
☞$_1$ a. pə.ˈlud		*		*
☞$_2$ b. ˈplud		**		
☞$_3$ c. pə.ˈludŭ	*		*	*
☞$_4$ d. ˈplu.dŭ	*	*	*	

Candidates (a) and (b), from which the final vowel has been deleted, are more well-formed than (c) and (d), in which this vowel has been preserved. This is the result of the ranking *V̆]$_{\omega}$ ≫ MAX(seg). Between candidates (a) and (b), candidate (a), which preserves the remaining unstressed vowel, is more well-formed – this time because MAX(seg) outranks *σ̆/{ɐ,ə}. The same is true of candidates (b) and (c) – candidate (c), which preserves the non-final unstressed vowel, is rated as more well-formed. The analysis therefore orders the four candidates as follows, in terms of decreasing well-formedness: [pə.ˈlud] > [ˈplud] > [pə.ˈlu.dŭ] > [ˈplu.dŭ]. Since all four of these candidates only violate constraints from below the cut-off, it is predicted that all four are possible outputs.

As the analysis now stands, it predicts that there will be multiple variants for words that contain more than one unstressed vowel. Additionally, it also predicts a well-formedness, and hence frequency, relationship between the different variants. Silva's data are based on a spontaneous conversation, so that there are just not enough repetitions of a single word to allow for meaningful generalisations about single words. These predictions can therefore not be tested with the data currently available. However, the analysis makes very specific predictions and is consequently in principle subject to empirical verification. The analysis can serve as a guide for designing future studies to investigate vowel deletion in Faialense Portuguese in more detail. The study by Nagy & Reynolds (1997) of variable deletion in Faetar can serve as an example of how this could be done. Nagy & Reynolds used two kinds of tasks to collect their data. The first was a picture-naming task, and the second a narrative task where subjects were presented with a picture and then asked to tell a story about the situation represented in the picture. In this manner, they were able to elicit multiple repetitions of single words, and hence to state meaningful generalisations about the variant realisations of individual words.

(18) also shows that there are other constraints that could play a role in this variation pattern. The variants that delete the initial vowel all violate *COMPLEX, while those that delete the final vowel violate NOCODA. For

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

unstressed vowels in other phonological environments, still other markedness constraints could be involved. The ranking of these constraints relative to the constraints used in (18) could also have an influence on the well-formedness and frequency relationship between the different variants. However, Silva coded his data for several syllable well-formedness factors such as these before submitting the data to the VARBUL analysis. The VARBUL analysis selected the identity of the vowel and the position of the vowel in the prosodic word as factors that contribute toward the variation. It found no effect of syllable well-formedness on the deletion rate (Silva 1997: 304). This implies that these syllable well-formedness constraints rank lower than the constraints used in (18), so that any effect that they could have are overshadowed by the higher-ranking constraints.

Additionally, there is no *a priori* reason to expect that some vowels are more or less likely than others to occur in specific phonological contexts. Inasmuch as these syllable well-formedness constraints will influence the deletion rate, they will do so equally for all vowels. The data in Table I can be thought of as the deletion rates of the different vowels, assuming that everything is equal except for the identity of the vowel and the position of the vowel in the prosodic word.

2.5.2 *Stressed vowels.* The discussion up to now has only focused on unstressed vowels. Based on the deletion and reduction patterns observed for these vowels, I have argued that MAX(seg) ranks below the cut-off. This could imply that stressed vowels should also be subject to at least variable deletion. To see that this is the case, consider a word like *livro* ['livɾŭ] 'book' (from /livɾo/). The tableau in (19) shows the violation profile of a candidate that preserves the stressed ['i] and a candidate that deletes this vowel. Only the constraint violations not shared by the candidates are included in the tableau.

(19)

/livɾo/	MAX(seg)	*[+hi,+fr]$_{old}$
☞$_1$ a. 'livɾŭ		*
☞$_2$ b. lvɾŭ	*	

Both candidates violate only constraints below the cut-off, and hence both are predicted as possible outputs. Although stressed vowels do not delete in Faialense Portuguese, the analysis currently predicts that they should. The solution is obvious – deletion of a stressed vowel should violate a constraint that ranks higher than the cut-off, so that candidate (19b) will be eliminated as a possible output for the input /livɾo/. Up to now, I have used general MAX(seg) – the constraint violated by deletion of a segment from any position. However, there is also a version of this constraint that is violated specifically by deletion from a stressed syllable, i.e. MAX(seg)/*σ́* (Beckman 1997, 1999). Ranking MAX(seg)/*σ́* above the cut-off correctly predicts that stressed vowels should not delete. This is

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

illustrated in tableau (20). In this tableau, candidate (b) violates a constraint from above the cut-off. Since there is a candidate available that does not violate any constraint from above the cut-off, (b) is deemed ungrammatical, and made inaccessible as possible output.

(20)

/livɾo/	MAX(seg)/σ́	MAX(seg)	*[+hi,+fr]$_{old}$
☞$_1$ a. ˈlivɾŭ			*
b. lvɾŭ	*!	*	

3 Variation in Ilokano

In the version of ROE used above, there was only one cut-off point on the hierarchy. This raises a question: when there is more than one unrelated variable process in a language, is there still only one cut-off, or does each variable process require its own cut-off? If a separate cut-off is required for each variable process, it would undermine the value of the cut-off as a way of limiting variation. In this section, I argue that there is only one cut-off per language, and if more than one variable process is observed it is because the constraints that rank below the cut-off happen to result in variation in more than one context. I discuss Ilokano as an example of a language with multiple variable processes, and show that an analysis of Ilokano variation is possible assuming a single cut-off. The variation in Ilokano phonology was first pointed out by Hayes & Abad (1989), and later analysed by Boersma & Hayes (2001). No information is available about the frequency with which different variants are attested. In the analysis that I present here, I will therefore only show how to account for the fact that variation is observed, and will make no claims about the frequency of different variant outputs.

3.1 The Ilokano data

Like Boersma & Hayes (2001), I analyse optional metathesis in Ilokano, as well as variation between three forms of the so-called heavy reduplicant. This section first presents the data. The analysis of data within ROE follows in §3.2.

3.1.1 *Variation in the form of the reduplicant.* Ilokano has two forms of reduplication: heavy, where the reduplicant is bimoraic, and light, where the reduplicant is monomoraic. Stems that start with CG clusters (where G = glide) have three variant forms for the heavy reduplicant. The first variant copies enough segmental material from the stem to satisfy the bimoraicity requirement – i.e. it copies the initial CGVC sequence from the stem. A second way to achieve a bimoraic reduplicant is to vocalise the glide and resyllabify the first consonant of the stem into the coda position

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

of the reduplicant. Finally, it is also possible to vocalise and lengthen the glide. The examples in (21) are from Boersma & Hayes (2001: 56).

(21)	*stem*	rwa.ŋan 'door'	pja.no 'piano'	bwa.ja 'crocodile'
	reduplicated forms			
	variant 1	rwaŋ.rwa.ŋan	pjan.pja.no	na.ka.bwaj.bwa.ja
	variant 2	rur.wa.ŋan	pip.ja.no	na.ka.bub.wa.ja
	variant 3	ruː.rwa.ŋan	piː.pja.no	na.ka.buː.bwa.ja
		'doors'	'pianos'	'act like a crocodile'

3.1.2 *Variable metathesis of* [ʔw] *sequences.* Ilokano does not tolerate vowel hiatus. Whenever hiatus is created by morphological concatenation, it is resolved either by glottal stop epenthesis or by changing the first vowel into a glide. The rules governing the choice between the two repairs are somewhat involved – see Hayes & Abad (1989) and §3.4.3. However, for now it is sufficient to know that the sequence /oV/ is generally resolved by changing /o/ into the glide [w]. There is a small number of Ilokano roots that end in the sequence /Vʔo/. When a vowel-initial suffix attaches to these roots, root-final /o/ changes to [w] and syllabifies into the onset of the following syllable, while the /ʔ/ forms the coda of the preceding syllable (/-Vʔo+V-/→[-Vʔ.wV-]). All of these forms have an alternative output where [w] and [ʔ] metathesise ([-Vw.ʔV-]). The examples in (22) are from Hayes & Abad (1989: 338) and Boersma & Hayes (2001: 55).

(22)	baʔo 'rat'	/pag+baʔo+an/	→	pag.baʔ.wan ~ pag.baw.ʔan	'place where rats live'
	taʔo 'person'	/taʔo+en/	→	taʔ.wen ~ taw.ʔen	'to repopulate'
	ʔaggaʔo 'to serve rice'	/pag+gaʔo+an/	→	pag.gaʔ.wan ~ pag.gaw.ʔan	'place where rice is served'
	daʔo 'kind of tree'	/pag+daʔo+an/	→	pag.daʔ.wan ~ pag.daw.ʔan	'place where *dao*'s are planted'

3.2 Analysis of variable of reduplication

There are three variants for the heavy reduplicant of a word starting with the sequence CG. There are two things that need to be shown. First, all three observed variants must indeed be possible outputs. This requires that all constraints that disfavour any of the observed variants rank below the cut-off. Secondly, it needs to be ensured that no other candidate reduplicant is predicted as possible output – i.e. all other reduplicants must be disfavoured by some constraint ranked above the cut-off.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

I use the word /rwaŋan/ as an example. The first variant, [rwaŋ.rwa.ŋan], violates *Complex and NoCoda. It also violates Max-BR, because the reduplicant is an incomplete copy of the base. The second variant, [rur.wa.ŋan], avoids violation of *Complex but still violates NoCoda and Max-BR. It also violates two additional constraints: Align(stem-L, σ-L), because the left edge of the stem is not aligned with the left edge of a syllable (Boersma & Hayes 2001: 59), and Ident-BR[syllabic], because it copies the base glide [w] as a vowel [u]. The third variant, [ruː.rwa.ŋan], violates Max-BR, Ident-BR[syllabic], *Complex and NoCoda. Boersma & Hayes (2001) argue that [ruː.rwa.ŋan] also violates Ident-BR[long]. However, I assume that base [w] is not specified for (vowel) length and that realising it as [uː] in the reduplicant therefore does not violate Ident-BR[long]. (It is crucial that [ruː.rwa.ŋan] does not violate Ident-BR[long], since, as I will show just below, Ident-BR[long] ranks above the cut-off, to rule out the unobserved variant [rwaː.rwa.ŋan].)

The variant reduplicants therefore violate Max-BR, Ident-BR[syllabic], Align(stem-L, σ-L), *Complex and NoCoda. For these reduplicants to be possible outputs, all these constraints must rank below the cut-off. Since no information on the frequency of the variants is available, it is not possible to rank these constraints relative to each other. Of course, these reduplicated forms also violate other constraints, such as ordinary segmental markedness constraints. These constraints also have to rank below the cut-off. In §3.4, some of the implications of this are considered.

The last thing that needs to be done is to rule out the non-observed reduplicant forms. Each of these must be disfavoured by a constraint above the cut-off. I discuss a few of the most obvious contenders here. (i) Heavy reduplicants of Ilokano are always exactly bimoraic. Following Boersma & Hayes (2001), I assume that there is a constraint that enforces this shape on the heavy reduplicant. Any reduplicant with fewer or more than two moras can be ruled out by this constraint. (ii) The first segment in the reduplicant is always identical to the first segment of its base – i.e. a form such as [waŋ.rwa.ŋan] is ungrammatical because the reduplicant has skipped the base-initial [r]. This can be enforced by ranking Anchor-L-BR above the cut-off. (iii) The reduplicant never copies a non-contiguous string from the base – i.e. [raŋ.rwa.ŋan] is not observed because the reduplicant has skipped the [w] from the base. This shows that Contiguity-BR ranks above the cut-off. (iv) A bimoraic reduplicant can also be achieved by copying the base up to its vowel, and lengthening the vowel – i.e. a form like [rwaː.rwa.ŋan] also satisfies the bimoraicity requirement. This form, however, copies a short base vowel as long, thereby violating Ident-BR[long], which is responsible for its ungrammaticality. In a similar way, all other potential forms for the reduplicant can be eliminated.

The tableau in (23) shows how a sample of reduplicated forms for [rwa.ŋan] are evaluated. The abstract morpheme /hred/ represents the heavy reduplicant. In this tableau, the three actually observed

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

reduplicants are disfavoured only by constraints below the cut-off, so that they are predicted as possible outputs. All other candidates are disfavoured by at least one constraint ranked above the cut-off, and are therefore excluded from being possible outputs.

(23)

/HRED+rwaŋan/	HRED=*μμ*	ID-BR[long]	ANCHOR-L-BR	CONTIG-BR	MAX-BR	ID-BR[syll]	ALIGN(stem)	*COMPLEX	NOCODA
☞ a. rwaŋ.rwa.ŋan					**			**	**
☞ b. rur.wa.ŋan					****	*	*		**
☞ c. ruː.rwa.ŋan					****	*		*	*
d. rwa.ŋan.rwa.ŋan	*!							**	**
e. rwa.rwa.ŋan	*!				***			**	*
f. rwaː.rwa.ŋan		*!			***			**	*
g. waŋ.rwa.ŋan			*!		***			*	**
h. raŋ.rwa.ŋan				*!	***			*	**

3.3 Analysis of variable metathesis

When a root that ends in /-VʔV/ is followed by a vowel-initial suffix, variation is observed between [-VʔwV-] and [-VwʔV-]. The analysis needs to account for the fact that both of these outputs are observed as variants, and also that these are the only observed variants. I will use the form /taʔo+en/→[taʔ.wen] ~ [taw.ʔen] as an example.

I follow Boersma & Hayes in assuming that both observed variants violate IDENT-IO[syllabic] (because the vowel /o/ is realised as a glide [w]). Additionally, the metathetic variant, [taw.ʔen], also violates LINEARITY-IO. Both of these constraints must therefore rank below the cut-off. The non-metathetic candidate, [taʔ.wen], avoids the LINEARITY-IO violation, but has a glottal stop in coda position, and therefore earns a violation of the constraint *ʔ]$_\sigma$. This shows that *ʔ]$_\sigma$ must also rank below the cut-off. (On derived and underived [ʔ] codas, see §3.4.2.)

Each of the non-observed candidates must be disfavoured by a constraint above the cut-off. The faithful candidate, [ta.ʔo.en], violates ONSET. Ilokano tolerates no onsetless syllables, and ONSET can hence be ranked above the cut-off. A candidate that avoids the ONSET violation by deleting one of the vowels, [ta.ʔen] or [ta.ʔon], can be ruled out by MAX-IO. Again, since Ilokano never tolerates violation of MAX-IO, it can be ranked safely above the cut-off. Another candidate to rule out is one that changes /o/ into a glide and deletes the glottal stop, i.e. [ta.won]. Such a candidate avoids both the ONSET and *ʔ]$_\sigma$ violations, but is ruled out by MAX-IO.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

A last candidate to consider is one that resolves hiatus by epenthesis of a consonant between the two vowels. The epenthetic consonant in Ilokano is the glottal stop, and the candidate to rule out is therefore [ta.ʔo.ʔen]. In this candidate, segmental material intervenes between the stem and the suffix, so that this candidate violates ALIGN(suffix-L, stem-R). Ilokano does not usually tolerate epenthesis between the stem and suffix (Hayes & Abad 1989). ALIGN(suffix-L, stem-R) can therefore be ranked above the cut-off, to rule out [ta.ʔo.ʔen].[13]

Ranking ONSET, MAX-IO and ALIGN(suffix-L, stem-R) above the cut-off rules out the non-observed candidates, while ranking IDENT-IO [syllabic], *ʔ]$_\sigma$ and LINEARITY-IO below the cut-off allows the observed variants, as shown in the tableau in (24). In this tableau, all candidates except for the first two violate some constraint above the cut-off. Only the first two candidates are therefore predicted to be observed as output variants.

(24)

/taʔo+en/	ONSET	MAX-IO	ALIGN(sfx)	ID-IO[syll]	*ʔ]$_\sigma$	LINEARITY-IO
☞ a. taʔ.wen				*	*	
☞ b. taw.ʔen				*		*
c. ta.ʔo.en	*!					
d. ta.ʔo.ʔen			*!			
e. ta.ʔen		*!				
f. ta.wen		*!		*		

3.4 Testing the implications of the analysis

In (25), I give a summary of the analysis developed above. Since the frequency of the different variants is not known, the constraints below the cut-off cannot be ranked.

(25) *Interim ranking for Ilokano*

HRED=$\mu\mu$, ANCHOR-L-BR, CONTIGUITY-BR, ONSET, MAX-IO, ALIGN(suffix-L, stem-R), IDENT-BR[long]

---------- cut-off ----------

MAX-BR, IDENT-BR[syllabic], ALIGN(stem-L, σ-L), *COMPLEX, NOCODA, IDENT-IO[syllabic], *ʔ]$_\sigma$, LINEARITY-IO

With so many constraints below the cut-off, it seems possible that this analysis might predict variation in contexts that have not been considered up to now. Put differently: it seems possible that there is some previously

[13] Boersma & Hayes use DEP-IO(ʔ) to rule out this candidate. However, Ilokano does have a process of variable [ʔ]-epenthesis, so that DEP-IO(ʔ) must rank below the cut-off.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

unconsidered input that will have more than one candidate disfavoured only by constraints from below the cut-off. If this were the case, then the analysis would predict variation also for that input. In this section, I consider a few such cases and show that for some of them this prediction is borne out. I also show how to rule out variation in some of the contexts for which the current analysis incorrectly predicts variation.

3.4.1 *Metathesis in contexts other than* /-ʔo+V/. To explain variable metathesis in [taʔ.wen] ~ [taw.ʔen], I argued that LINEARITY-IO ranks below the cut-off. Since some of the observed variants violate *COMPLEX and NOCODA, I argued that these constraints also rank below the cut-off. With all three of these constraints below the cut-off, we should expect variable metathesis, to avoid violation of *COMPLEX and NOCODA. The tableau in (26) illustrates this with made-up examples. (A leftward pointing hand marks a candidate incorrectly predicted as possible output.)

(26) *Variable metathesis to avoid violation of* *COMPLEX *and* NOCODA

	/traka/	*COMPLEX	NOCODA	LINEARITY-IO
a.	☞ i. tra.ka	*		
	☜ ii. tar.ka		*	*
b.	/palta/			
	☞ i. pal.ta		*	
	☜ ii. pla.ta	*		*

(26) shows that, for both kinds of inputs, there are two candidates disfavoured only by constraints below the cut-off. The analysis then incorrectly predicts variable metathesis for these inputs. Boersma & Hayes (2001 : 61) explicitly state that metathesis is not observed in this context. This presents a problem to the analysis developed above. LINEARITY-IO was ranked below the cut-off for one reason, and *COMPLEX and NOCODA for different reasons. However, once all of them are below the cut-off, there is nothing preventing them from interacting. In this context, this interaction leads to incorrect predictions.

Up to now, only one of the constraints violated by metathesis, LINEARITY-IO, has been considered. But metathesis also violates CONTIGUITY-IO. When an input /abcd/ maps onto [acbd], [a] and [b], which were contiguous in the input, are separated by [c] in the output. When two consonants metathesise, the segment causing the CONTIGUITY-IO violation will be a consonant. However, when a consonant and a vowel metathesise, it is a vowel that causes the CONTIGUITY-IO violation. This is illustrated in (27). In (27a), /a/ and /p/ are separated by the consonant [k] in the output, and /k/ and /i/ by the consonant [p]. In (27b), /a/ and /k/ are separated by the consonant [r], but /p/ and /r/ are separated by the vowel [a].

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(27) a. CC-*metathesis* /tapki/ → [takpi]
b. CV-*metathesis* /praki/ → [parki]

In her discussion of metathesis in Leti, Hume (1998) argues for separate CONTIGUITY constraints for these two situations. Leti uses metathesis to avoid *COMPLEX violations (cf. /ulit+prai/→[ul.tip.ra.i], *[u.lit.pra.i]).[14] The metathetic candidate violates not only CONTIGUITY, but also ONSET. The problem is that there is another metathetic candidate that avoids violation of *COMPLEX and ONSET, [u.lit.pa.ri]. Hume claims that the crucial difference between these two candidates is that it is a consonant, [t], that causes the CONTIGUITY violation in the observed [ul.tip.ra.i], while it is a vowel, [a], that causes the violation in the ungrammatical [u.lit.pa.ri]. Leti tolerates a consonant between two underlyingly contiguous segments, but not a vowel. Based on this, Hume defines two separate CONTIGUITY-IO constraints, which are violated by intervening consonants and vowels respectively. The definition in (28) is based on Hume (1998: 167).

(28) CONTIGUITY-IO-X
A contiguous string in the input may not be separated by X in the output (where X is a variable that ranges over the set {consonant, vowel}).

In all contexts in which metathesis is observed in Ilokano, it is always a consonant which intervenes between contiguous underlying segments – in /taʔo+en/→[taw.ʔen], [w] separates /a/ and /ʔ/, and [ʔ] separates /o/ and /e/. However, in (26), it is a vowel that is responsible for the CONTIGUITY-IO violation. Like Leti, Ilokano allows a consonant to intervene between underlying segments, but not a vowel. Tableau (29) shows how ranking CONTIGUITY-IO-V above the cut-off rules out the non-observed cases of metathesis.

(29) *Ruling out non-observed metathesis in Ilokano*

	/traka/	CONTIGUITY-IO-V	*COMPLEX	NOCODA	LINEARITY-IO
a.	☞ i. tra.ka		*		
	ii. tar.ka	*!		*	*
b.	/palta/				
	☞ i. pal.ta			*	
	ii. pla.ta	*!	*		*

Another context where the current analysis predicts metathesis is in a morpheme with the shape /… $V_1C_2C_3V_4$ …/. Metathesis of /C_2C_3/ will

[14] CONTIGUITY-IO applies only intramorphemically in Leti. In order to make the different morphemes more easily recognisable, I underline the first morpheme in all Leti examples.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

violate CONTIGUITY-IO-C, like the observed metathesis in /taʔo+en/ →[taw.ʔen]. This kind of metathesis is observed for at least some words in the dialect of Ilokano described by Vanoverbergh (1955). He gives the following examples: /takpil/→[takpil] ~ [tapkil] and /nakaldeman/→ [nakaldeman] ~ [nakadleman]. Hayes & Abad (1989) do not refer to metathesis in this context. If their dialect is like that described by Vanoverbergh, then the analysis developed here correctly predicts that this kind of variable metathesis should be observed. However, if their dialect does not have metathesis in this context, it can easily be ruled out. The metathesis that is observed in the Hayes & Abad dialect always occurs at the right edge of a root. Hume (2001) and Hume & Mielke (2001) show that by far the most examples of metathesis are observed only at the right edge of a morpheme. They ascribe this to the fact that metathesis earlier in the root inhibits word recognition more than metathesis at the end of the root. It is therefore likely that there are separate positional versions of LINEARITY-IO. Ilokano would then tolerate violation only of the version of LINEARITY-IO that applies to the right edge of the root morpheme.

3.4.2 *Derived and underived* [ʔ] *codas.* The discussion in §3.3 showed that Ilokano allows [ʔ] codas in forms such as [taʔ.wen]. This [ʔ] coda is 'derived' in the sense that the /ʔ/ does not occur before a consonant or word-boundary in the input, /taʔo+en/. Except in derived environments such as these, Ilokano does not tolerate [ʔ] codas. Put differently: the Ilokano lexicon contains no morphemes with /ʔ/ followed directly by a consonant (/… ʔC …/) or a word boundary (/… ʔ#/). In pre-OT phonology, it could simply be stated that such forms are absent from the lexicon (see Hayes & Abad 1989). However, under richness of the base this kind of stipulation is not tolerated in OT. In an OT grammar, inputs such as /… ʔC …/ and /… ʔ#/ must be mapped onto some grammatical output. The grammar that is developed for Ilokano must therefore tolerate violation of $*ʔ]_{\sigma}$ in derived but not in underived contexts. Several ways of handling this kind of phenomenon have been proposed in OT. I opt to use comparative markedness (McCarthy 2002, 2003) here simply because I have already made use of comparative markedness in this paper (see §2.4.1). However, there are other ways, such as OO-correspondence, which would have worked equally well (see Boersma & Hayes 2001 for an example of the use of OO-correspondence).

Under comparative markedness, there are two versions of $*ʔ]_{\sigma}$: $*ʔ]_{\sigma\ \text{old}}$ and $*ʔ]_{\sigma\ \text{new}}$. If for some input the fully faithful candidate contains a [ʔ] coda, then all candidates in which this corresponding /ʔ/ also appears in coda position violate $*ʔ]_{\sigma\ \text{old}}$. All other [ʔ] codas violate $*ʔ]_{\sigma\ \text{new}}$. The difference between derived and underived [ʔ] codas is then that underived codas violate $*ʔ]_{\sigma\ \text{old}}$, while derived codas violate $*ʔ]_{\sigma\ \text{new}}$. It is therefore $*ʔ]_{\sigma\ \text{new}}$ that is violated in the mapping /taʔo+en/→[taʔ.wen], so that $*ʔ]_{\sigma}$ used in §3.3 above needs to be replaced with $*ʔ]_{\sigma\ \text{new}}$. Since $*ʔ]_{\sigma\ \text{new}}$ ranks below the cut-off, a candidate that violates this constraint can take part in a variable process.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

Underived [ʔ] codas violate *ʔ]$_{\sigma\ \text{old}}$. To avoid candidates that violate this constraint from surfacing as outputs, *ʔ]$_{\sigma\ \text{old}}$ needs to rank above the cut-off. In addition, it also has to rank above at least one of the faithfulness constraints that can be violated, to avoid violation of *ʔ]$_{\sigma\ \text{old}}$. It has already been established that LINEARITY-IO ranks below the cut-off, and hence below *ʔ]$_{\sigma\ \text{old}}$. With this ranking, inputs such as /… VʔCV …/ and /… Vʔ/ will always surface with metathesis. This is illustrated in (30). This tableau also contains examples of derived [ʔ] codas, to show how the grammar handles the two kinds of [ʔ] codas differently. The first example is a real Ilokano word. The last two are made up, since no morphemes with these sequences actually exist in Ilokano.

(30) *Derived and underived* [ʔ] *codas*

		*ʔ]$_{\sigma\ \text{old}}$	LINEARITY-IO	*ʔ]$_{\sigma\ \text{new}}$
a.	/taʔo+en/			
	☞ i. taʔ.wen			*
	☞ ii. taw.ʔen		*	
b.	/taʔlak/			
	i. taʔ.lak	*!		
	☞ ii. tal.ʔak		*	
c.	/bataʔ/			
	i. ba.taʔ	*!		
	☞ ii. bat.ʔa		*	

Tableau (30) shows how a derived [ʔ] coda only violates constraints below the cut-off, and is therefore tolerated in the output. An underived [ʔ] coda, however, violates a constraint above the cut-off, and since there is a candidate that violates no constraint above the cut-off, the faithful candidate will not be selected as output. This serves as an example of one of the ways in which a categorical process is modelled in ROE (see §1.2). When there is only one candidate that is not disfavoured by any constraints above the cut-off, this single candidate is the only grammatical candidate, and hence the only observed output. The next section provides an example of the other way in which a categorical process can arise in ROE.

3.4.3 *Compulsory* [ʔ]*-epenthesis*. In the analysis above, I argued that IDENT-IO[syllabic] ranks below the cut-off, while ALIGN (suffix-L, stem-R) ranks above the cut-off. The former ranking is needed to explain why [taʔ.wen] is a variant output for /taʔo-en/. The latter ranking is necessary to rule out *[ta.ʔo.ʔen] (with epenthetic [ʔ] between stem and suffix) as an output for this input – see (24). The ranking ALIGN(suffix-L, stem-R) ≫ cut-off ≫ IDENT-IO[syllabic] implies that vowel hiatus between a stem and suffix will always be resolved by changing the stem-final vowel into a glide and never by consonantal epenthesis. Inserting a consonant

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

between stem and suffix will result in violation of ALIGN(suffix-L, stem-R) from the above the cut-off, while glide formation will violate IDENT-IO[syllabic] from below the cut-off. This prediction, however, is not correct.

Vowel hiatus created by morphological concatenation is usually resolved by changing the first of the two vowels into a glide (/i e/ change to [j], and /u o/ to [w]), as in the example /taʔo+en/→[taʔ.wen] discussed in §3.3. However, when the first vowel is /a/, glottal stop epenthesis is used to resolve the hiatus. The examples in (31) are from Hayes & Abad (1989).

(31) a. *Glide formation*

/babawi+en/	→ [ba.baw.jen]	'regret-GOAL-FOCUS'
/masahe+en/	→ [ma.sah.jen]	'massage-GOAL-FOCUS'
/saŋo+en/	→ [saŋ.wen]	'to cause to face forwards'

b. *Glottal stop epenthesis*

/pag+pja+en/	→ [pag.pja.ʔen]	'to make healthy'
/basa+en/	→ [ba.sa.ʔen]	'read-GOAL-FOCUS'
/pag+saka+an/	→ [pag.sa.ka.ʔan]	'place to walk barefoot'

As discussed in §1.2, there are two ways in which a categorical output can arise in ROE. One is when there is only one candidate that is disfavoured only by constraints below the cut-off. This single candidate is then observed as the only variant. Examples of this were seen in (29) and (30) above. The second way is when all candidates are disfavoured by constraints from above the cut-off. The best candidate is then observed as the only output. Glottal stop epenthesis after a root-final /a/ presents an example of this.

I follow Boersma & Hayes (2001) in assuming that there is a markedness constraint against low glides (*a̯). For an input like /basa+en/, the observed output, [ba.sa.ʔen], violates ALIGN(suffix-L, stem-R) from above the cut-off, while a candidate that changes the /a/ into a low glide, [bas.a̯en], violates IDENT-IO[syllabic] from below the cut-off. However, it also violates the markedness constraint *a̯. If *a̯ is ranked above ALIGN(suffix-L, stem-R), then both the actual output [ba.sa.ʔen] and the ungrammatical [bas.a̯en] violate constraints above the cut-off, and because [ba.sa.ʔen] violates the lower-ranking constraint, it is selected as the output. To ensure that [ba.sa.ʔen] is the best candidate, a few additional rankings between the constraints above the cut-off are required. To rule out the faithful candidate, [ba.sa.en], ONSET has to dominate ALIGN(suffix-L, stem-R). To rule out a candidate that deletes one of the vowels, MAX-IO must rank higher than ALIGN(suffix-L, stem-R). This is illustrated in (32). In this tableau, all candidates violate at least one constraint above the cut-off. The candidate that violates the lowest-ranking of these constraints, [ba.saʔen], is therefore selected as the only output.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(32)

/basa+en/	MAX-IO	ONSET	*a̱	ALIGN(sfx)	ID-IO[syll]
☞ a. ba.sa.ʔen				*	
b. bas.a̱en			*!		*
c. ba.sa.en		*!			
d. ba.san	*!				

3.4.4 *Intervocalic consonant clusters.* In the grammar developed above, *COMPLEX and NOCODA both rank below the cut-off – see (25). These rankings are necessary, since several observed variants do violate these constraints – see, for instance, the discussion of variable reduplication in §3.2. However, with both *COMPLEX and NOCODA below the cut-off, any intervocalic biconsonantal cluster should allow two syllabifications – one that violates *COMPLEX, and another that violates NOCODA.

An input with the form /… VCCV …/ should variably map onto […VC.CV …] and […V.CCV …]. This is, however, not true. See Hayes & Abad (1989: 335): 'sequences of the form VC_1C_2V are syllabified $VC_1.C_2V$, even where C_1C_2 can appear word-initially: compare *pur.wák* 'to scatter' with *rwár* 'outside'.'

(33)

/purwak/	*COMPLEX	NOCODA
☞ a. pur.wak		*
☜ b. pu.rwak	*	

The tableau in (33) illustrates the problem for the input /purwak/. In this tableau, both candidates violate only constraints below the cut-off, so that [pu.rwak] is incorrectly predicted as a variant output. What is necessary is a constraint that ranks above the cut-off that penalises ungrammatical [pu.rwak]. The natural choice seems to be *COMPLEX. However, ranking *COMPLEX above the cut-off will prevent any form with a tautosyllabic consonant cluster from taking part in a variable process, implying, for instance, that no variation should be observed in the heavy reduplication of roots that start with /CGV/. This gives evidence for the need to split the old and new versions of *COMPLEX, and for ranking $*\text{COMPLEX}_{\text{new}}$ above the cut-off and $*\text{COMPLEX}_{\text{old}}$ below the cut-off. The ungrammatical mapping /purwak/→[pu.rwak] then violates $*\text{COMPLEX}_{\text{new}}$, and is correctly ruled out as a possible output. This is shown in (34). It should be noted that this analysis depends on the assumption that syllabification is not specified in the input. If syllabification could be present underlyingly, an input like /pu.rwak/ would have been allowed under richness of the base, and for such an input the candidate [pu.rwak] would have violated $*\text{COMPLEX}_{\text{old}}$. However, since syllabification is never contrastive (Clements 1986: 318, Hayes 1989: 260, Blevins 1995: 221, McCarthy 2003: 10), it is standard to assume that syllabification is absent from the input.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(34) /purwak/ *with* *COMPLEX$_{old}$ *and* *COMPLEX$_{new}$ (FFC: [pur.wak])

/purwak/	*COMPLEX$_{new}$	*COMPLEX$_{old}$	NOCODA
☞ a. pur.wak			*
b. pu.rwak	*!		

In comparative markedness theory, markedness constraints compare output candidates to the fully faithful candidate (FFC). Since syllabification is assumed not be specified in the input, there are always multiple FFC's that differ only in terms of syllabification. It then becomes necessary to decide which of the different FFC's should be used as *the* candidate against which output candidates are compared. In his discussion of comparative markedness theory, McCarthy dedicates a long section to this question (2002: §6.2). His argument comes down to the following: every FFC is assigned markedness violations, assuming that the candidate being evaluated is the comparison FFC – the effect of this is that only the old versions of the markedness constraints are active at this stage. The FFC that comes out most harmonic in this competition is then selected as *the* comparison FFC that is used in ordinary production-oriented use of EVAL. For an input like /purwak/ from (34) above, there are two FFC's to consider, [pur.wak] and [pu.rwak]. As shown in (34), the FFC that should be selected as the comparison candidate is [pur.wak]. [pur.wak] will be selected if it is rated as more harmonic than [pu.rwak] in terms of the M_{old} constraints, and this will be the result under the ranking *COMPLEX$_{old}$ ⩾ NOCODA$_{old}$.

In the grammar developed above, I have not distinguished between NOCODA$_{old}$ and NOCODA$_{new}$. However, an undifferentiated markedness constraint is just shorthand for the situation in which no constraints rank between the old and new versions of the markedness constraint. The slot occupied by NOCODA in the hierarchy is therefore really occupied by NOCODA$_{old}$ and NOCODA$_{new}$ together, with no ranking specified between them. Since undifferentiated NOCODA ranks below the cut-off, it follows that NOCODA$_{old}$ does too. Tableau (34) shows that *COMPLEX$_{old}$ also ranks below the cut-off. The ranking between the constraints below the cut-off influences only the relative frequency of the different variants, and since no information is available about the frequency of the variants in Ilokano, the ranking *COMPLEX$_{old}$ ⩾ NOCODA$_{old}$ is consistent with the available data.

It is important to note that splitting the old and new versions of *COMPLEX does not have an influence on word-initial consonant clusters, on triconsonantal intervocalic clusters or on the variation observed in heavy reduplication. In all of these instances, there is no FFC candidate that does not violate *COMPLEX. The comparison FFC will hence violate *COMPLEX, and any candidate that shares this violation with the FFC will violate *COMPLEX$_{old}$ (below the cut-off), and not *COMPLEX$_{new}$ (above the cut-off). Consider the heavy reduplicated form of /rwaŋan/. Following

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

McCarthy (2002), I assume that the input actually contains a full copy of the base for the reduplicant. The full segmental specification of /HRED+rwaŋan/ is therefore /rwaŋan+rwaŋan/ (the reduplicant is underlined). The comparison FFC for this input will be [rwa.ŋan.rwa.ŋan]. There is no faithful candidate that does not have a consonant cluster in word-initial position. Similarly, there is no FFC that does not have a tautosyllabic cluster where the two morphemes meet. The tableau in (35) repeats the relevant part of tableau (23), where heavy reduplication was first discussed, with *COMPLEX split into its old and new versions. As is clear, not one of the observed variants has a tautosyllabic consonant cluster that is not shared by the FFC. All of the observed variants violate only $*\text{COMPLEX}_{\text{old}}$, which ranks below the cut-off. They are therefore still correctly predicted as possible outputs.

(35) FFC: [rwa.ŋan.rwa.ŋan]

/rwaŋan+rwaŋan/	HRED=*μμ*	$*\text{COM-PLEX}_{\text{new}}$	MAX-BR	ID-BR [syll]	ALIGN (stem)	$*\text{COM-PLEX}_{\text{old}}$	NO CODA
☞ a. rwaŋ.rwa.ŋan			**			**	**
☞ b. rur.wa.ŋan			****	*	*		**
☞ c. ruː.rwa.ŋan			****	*		*	*
d. rwa.ŋan.rwa.ŋan	*!					**	**

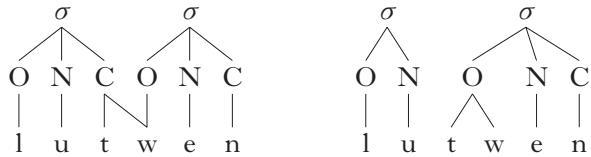
One complication remains – there is one context in which a form with a derived onset cluster does take part in a variable process. Glide formation (discussed above in §3.3 in relation to metathesis) optionally induces gemination of a preceding consonant (Hayes 1989: 269–276, Hayes & Abad 1989: 338). The examples below are from Hayes (1989).

(36) /luto+en/ → [lut.wen] ~ [lut.twen] 'cook-GOAL-FOCUS'
/pag+ʔaso+an/ → [pag.ʔas.wan] ~ [pag.ʔas.swan] 'where dogs are raised'
/damo+en/ → [dam.wen] ~ [dam.mwen] 'to be new to something'

The geminate variants have derived clusters, and hence violate $*\text{COMPLEX}_{\text{new}}$. Since these forms take part in a variable process, they cannot violate any constraint from above the cut-off, and $*\text{COMPLEX}_{\text{new}}$ must rank below the cut-off, counter to the claims above. A full explanation for these forms is left for future research. However, it is noteworthy that the only derived onset clusters observed in Ilokano are those where the first member of the cluster is one half of a geminate, and hence is not affiliated exclusively with onset position. Ilokano therefore allows structures like those in (37a), but not like those in (37b).

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(37) a. *Allowed:* [lut.twen] b. *Not allowed:* *[lu.twen]



The only instances of derived consonant clusters in Ilokano are those in which the first member of the cluster is licensed by its affiliation to the coda of the preceding syllable. If *COMPLEX were reformulated such that it bans syllabic margins that license more than one consonantal slot, it would be violated by the structure in (37b), but not by that in (37a). The geminate variants from (36) would then not violate *$\text{COMPLEX}_{\text{new}}$, so that this constraint can be ranked above the cut-off without ruling out the geminate variants from (36). However, such a reformulation of *COMPLEX is a departure from the way that it is traditionally viewed, and it will be necessary to consider the broader implications of such a move before this explanation of the geminate variants can be accepted. There is some evidence from English that also suggests this kind of interpretation of word-medial onset clusters. English has a process that devoices liquids after tautosyllabic voiceless obstruents – cf. *app*[ɹ̥]*ise vs. up* [ɹ]*ise* (Hoard 1971). In the pronunciation of a word like *atlas*, the /l/ is usually voiced, implying that /t/ is not tautosyllabic with /l/. However, Hoard (1971: 138) notes that the /l/ in *atlas* can variably be pronounced voiceless, and Anderson & Ewen (1987: 57) claim that there are dialects of English in which the /l/ is always voiceless, suggesting that it is sometimes possible to parse the /t/ into the onset of the second syllable. This is problematic, since English does not generally tolerate onset clusters of the form [tl]. The solution is to treat the instances where /l/ devoices as forms in which the /t/ is ambisyllabic and hence also affiliated with the coda of the preceding syllable. The two pronuciations of *atlas* are then [æt.ləs] and [æt.t̬l̥əs]. English can then be said to disallow onsets of the form [tl] when both of the consonants are solely licensed by the onset. However, when the [t] is (co-)licensed by a preceding coda, such a cluster is allowed. This correctly excludes word-initial [tl] clusters, since word-initially there is no preceding coda that can license the [t].

3.5 Summary and evaluation of the Ilokano analysis

In (38), I give a summary of the Ilokano grammar developed above. This summary is identical to that given in (25), with the following exceptions: (i) *ʔ]$_{\sigma}$ and *COMPLEX have been split into their respective old and new versions, (ii) *ʔ]$_{\sigma\ \text{old}}$, *$\text{COMPLEX}_{\text{new}}$, *a̰ and CONTIGUITY-IO-V have been added to the constraints above the cut-off, (iii) *ʔ]$_{\sigma\ \text{new}}$, *$\text{COMPLEX}_{\text{old}}$ and CONTIGUITY-IO-C have been added to the constraints below the cut-off

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

and (iv) the ranking {MAX-IO, ONSET, *a̱} ≫ ALIGN(suffix-L, stem-R) has been added to constraints above the cut-off.

(38) *Final ranking for Ilokano*

MAX-IO *a̱ ONSET

ALIGN(suffix-L, stem-R)

HRED=*μμ*, ANCHOR-L-BR, CONTIGUITY-BR, CONTIGUITY-IO-V, *ʔ]$_{\sigma\,old}$, *COMPLEX$_{new}$

cut-off

MAX-BR, IDENT-BR[long], ALIGN(stem-L, *σ*-L), *COMPLEX$_{old}$, NOCODA, IDENT-IO[syllabic], *ʔ]$_{\sigma\,new}$, LINEARITY-IO, CONTIGUITY-IO-C

In order to test the accuracy of this grammar, I used OTSoft (Version 2.1) (Hayes *et al*. 2004) to calculate its typology. I used the 'constrained factorial typology' option offered by the software, as this option allows specification of rankings that are not allowed to change, so that it is possible to limit reranking only to the rankings that are not specified as part of the grammar in (38). All the constraints above the cut-off therefore always dominate all the constraints below the cut-off. Additionally, ALIGN(suffix-L, stem-R) is not allowed to outrank MAX-IO, ONSET or *a̱. As inputs I used a representative example from each of the patterns discussed in this paper (/HRED+bwaja, taʔo+en, traka, palta, taʔlak, basa+en, purwak/). As output candidates, I used all acceptable and unacceptable outputs listed by Boersma & Hayes (2001: 65), and added to these the additional candidates that were used in the discussion above. When the typology is calculated in this manner, only the actually observed outputs are generated. This serves as confirmation that the analysis as developed above does indeed work for the part of the Ilokano grammar that is discussed in this paper.[15]

Unlike the Faialense Portuguese example discussed in §2, Ilokano has several unrelated variable processes. As a consequence, it is necessary to rank several unrelated constraints below the cut-off. This could create problems – as more constraints are ranked below the cut-off, these constraints can interact in unexpected ways and can predict variation in contexts where it should not be observed. Stated in more general terms: if some constraint ℂ is violated by a winning candidate in any derivation that involves variation, then ℂ must rank below the cut-off. And since ℂ ranks below the cut-off it can never rule out as ungrammatical any candidate in any other derivation.

A possible response to this is to introduce several cut-offs on the hierarchy, and to limit the activity of a cut-off to a specific process or class of morphemes. The presence of a heavy reduplicant in the input could

[15] The files that were used as input to OTSoft, as well the output files generated by OTSoft, are available (September 2006) from www.umich.edu/~coetzee/Ilokano/.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

trigger activation of a cut-off that could be used to explain the variation observed with heavy reduplication. Similarly, there could be a cut-off that is activated when /ʔ/ is parsed into coda position – i.e. by violation of *ʔ]$_{\sigma\ \text{new}}$. This could be used to explain variable metathesis in mappings such as /taʔo+en/ → [taʔ.wen] ~ [taw.ʔen]. In this way, it would be possible to limit the contexts in which variation is observed. However, this would result in a much less restrictive theory. The theory would then allow for process-specific grammars – something that OT avoids. In this section, I have shown that ROE can account for the variation patterns of Ilokano without resorting to these kinds of stipulations. It remains to be shown that ROE can successfully account for other similarly complicated variable processes.

4 Outstanding questions and extensions to ROE

4.1 Relative and absolute frequency

ROE predicts the relative frequency of variants. If *variant*$_1$ is more well-formed than *variant*$_2$, *variant*$_1$ is predicted to occur more frequently. However, ROE is silent about how much more frequent *variant*$_1$ will be. In this regard, ROE differs from other theories of variation in OT. Anttila (1997) proposes a theory in which constraints are allowed to be crucially unranked. On every evaluation occasion, one of the possible complete rankings between the unranked constraints is selected. If the candidate that is optimal differs under the different rankings, variation is observed. But this theory also makes predictions about the absolute frequency of the different variants. If there are n unranked constraints, then there are $n!$ possible rankings. If m out of these select some variant as optimal, then this variant should occur $m/n!$ of the time. In the Boersma & Hayes (2001) model, ranking is along a continuous scale. Constraints are not fixed at a specific place on this scale, but rather have a normally distributed ranking range on the scale. At every evaluation occasion, the constraint is ranked at a specific point in its ranking range. Constraints with overlapping ranges can then rank in opposite orders on different evaluation occasions, so that different candidates can be optimal. By moving the ranking ranges of two constraints closer or further apart, it is possible to control the frequency with which different rankings, and hence different variants, are observed.

A theory that predicts absolute frequencies accounts better for the data than a theory that predicts relative frequencies – provided that the absolute predictions are accurate. The other OT theories therefore seem to be better than ROE. However, the choice is not necessarily straightforward. In this section, I discuss some of the problems with theories that make absolute frequency predictions.

A theory that accounts for absolute frequencies must decide which frequencies to account for. Anttila (1997) shows that his model can account for the frequencies of the different allomorphs of the Finnish

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

genitive plural. Boersma & Hayes (2001) show that their model can also account for these frequencies. However, the percentages with which Anttila works are based on averages calculated over a corpus that contains texts written by several different authors. Anttila does not refer to the frequencies with which different individuals use different variants. However, the long tradition of variation studies in sociolinguistics has shown that the individual members of a speech community do not generally have exactly the same variation pattern. If absolute frequencies arise directly from the grammar, then individuals with different frequency patterns should have different grammars. But if the different members of a community have different grammars, then showing that there is a grammar that can account for the average community-wide frequency patterns has less meaning. What would be required, then, is to show that a grammar exists that can account for each individual variation pattern. Especially in a theory such as Anttila's, this can very quickly become practically impossible. If individuals in a community can have different grammars, there also needs to be a theory of how much difference is allowed between the individuals. Even if it has been shown that absolute frequency theories can be reasonably successful at accounting for community averages, these theories are still faced with conceptual questions about the relation between individual and community grammars.

Some of these problems can be sidestepped in ROE. Since ROE models only relative frequency, two individuals with the same relative frequency difference between two variants can have the same grammar even if they differ in their absolute frequency patterns. This suggests the following as a hypothesis for how individual grammars relate to community grammars. The community grammar specifies the relative well-formedness (and hence frequency) of variants. All members of the community share this grammar, and will therefore show the same relative frequency relationship between variants. Differences in absolute frequencies between individuals then do not originate directly in the grammar, but rather in the different ways in which individuals use the information supplied by grammar.[16] There is some evidence that the individuals in a community are related

[16] This is in general agreement with the way in which variation is traditionally treated in the 'variable rule' approach used in sociolinguistics. A variable rule expresses the probability that some variable grammatical rule will apply. The probability p of application is calculated with the logistic function $p/(1-p) = p_0/(1-p_0) \times p_1/(1-p_1) \times \ldots \times p_n/(1-p_n)$ (Rousseau & Sankoff 1978). In this expression, $p_1 \ldots p_n$ represent factors that either inhibit or promote application of the rule. p_0 is the probability that the rule will apply independently from any of these factors. The standard assumption is that individuals in the same speech community (that share the same grammar) can differ only in terms of p_0 and not in terms $p_1 \ldots p_n$ – see Guy (2003: 381): 'variation within the community will be confined to differences in the p_0 values …, while constraint effects, along with other features of the phonology, will be consistent for all community members'. In this approach, the differences between members of the same community do not reside in the grammar, but rather in p_0, which is not a grammatical factor.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

like this. Kariņš (1995) studied the deletion of unstressed vowels in Latvian (/pele/→['pe.le] ~ ['pel] 'mouse'). He found an average deletion rate of 86%. However, the eight individuals on whom his corpus is based have deletion rates ranging from 67% to 97%. Although the individuals differed in their absolute deletion rates, they all had the same relative frequency between the variants.[17]

A related problem for the absolute frequency theories is the well-known fact that a large part of variation is conditioned by non-grammatical factors such as age, gender, etc. If grammar accounts perfectly for the observed frequencies, it actually accounts for more than its fair share of the variation. The close fit that is sometimes found between observed and predicted frequencies in these models can then be a liability rather than an asset. Again, since ROE makes only relative frequency predictions, it does not face this problem to the same extent. The grammar is but one of the things that determine the frequency of variants, and all that it contributes is a general preference relationship between variants. It is the complex interplay between this preference provided by grammar and the preferences provided by other factors that co-determine the actually observed frequencies. Of course, to be an adequate theory of variation, ROE must ultimately contain a model of how all these different factors interact to produce actual frequencies. This is something that is still absent in the literature on formal phonology, although there are some well-developed ideas about this in the sociolinguistics literature.

In the end, the choice between an absolute frequency and a relative frequency theory might be a philosophical choice. However, I hope to have shown that absolute frequency theories do not necessarily account better for the data and are not *per se* superior.

4.2 Gratuitous repairs

In ROE, all candidates that are disfavoured only by constraints from below the cut-off are possible outputs. An anonymous reviewer points out a potential problem for this aspect of ROE. To illustrate, consider a language that allows no codas and that variably avoids them by vowel epenthesis or consonantal deletion – i.e. /pikat/→[pi.ka] ~ [pi.ka.ti]. In ROE, this language will have the grammar NoCoda ≫ cut-off ≫ {Max, Dep}. However, since Dep ranks below the cut-off, nothing can prevent a candidate with gratuitous epenthesis from surfacing as a variant. Tableau (39) illustrates the point.

[17] This begs the question of what do with individuals that subvert the relative frequencies. If the argument above is correct, then such individuals must have different grammars. One has a grammar that prefers $variant_1$ over $variant_2$, and the other that prefers $variant_2$ over $variant_1$.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

(39)

/pikat/	NoCoda	Max	Dep
a. pi.kat	*!		
☞ b. pi.ka		*	
☞ c. pi.ka.ti			*
☞ d. pi.ka.ti.ti			***
☞ e. etc.			******

As (39) shows, not only the candidate with minimal violation is predicted as possible, but also candidates with multiple gratuitous violations of Dep. This is an example of a more general problem – once a constraint like Max or Dep is below the cut-off, what will prevent variable epenthesis and deletion from occurring all over the place? There are two separate solutions to this problem. First consider the examples in (39). One of the properties of OT that has been subject to the most severe criticism is the infinity of the candidate set. Language users must compute the output in finite time and with finite computational resources. If the optimal candidate is identified by comparison with each competitor, then the infinite candidate set implies that an infinite number of comparisons are necessary. Samek-Lodovici & Prince (1999) have shown that the source of the infinity problem is located in the harmonically bounded candidates – they prove that for any input the number of non-harmonically bounded candidates (candidates that can be optimal under some ranking of the constraints) is actually finite. The infinity of the candidate set is the result of there being an infinite number of harmonically bounded candidates – specifically, there are infinitely many epenthetic candidates like the problematic candidates in (39). The computational problem can therefore be solved if there is a way to limit Gen so that it generates only the finite set of possible outputs. Riggle (2004) develops a finite state model of Gen that uses the constraint set to do exactly this. If Gen is understood in Riggle's terms, something that is independently necessary to solve the computational problem, then the problematic candidates from (39) disappear. These problematic candidates are all harmonically bounded and would therefore not even be generated under this conceptualisation of Gen.

What about non-harmonically bounded candidates? Imagine a situation where variable metathesis is observed in some context, while other contexts allow no metathesis. Because variable metathesis is observed, Linearity-IO must rank below the cut-off. But what then prevents variable metathesis from occurring in all contexts? The Ilokano example discussed earlier in §3 gives a real-life example, and also illustrates that this is only an apparent problem. Ilokano has variable metathesis to avoid violation of $*ʔ]_{\sigma new}$ (/taʔo+en/→[taʔ.wen] ~ [taw.ʔen]). But metathesis is not allowed as a way of avoiding violation of NoCoda (/palta/→[pal.ta], *[pla.ta]). To allow variable metathesis in the former context, I argued that Linearity-IO ranks below the cut-off. But the

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

mapping /palta/→[pla.ta] also violates LINEARITY-IO. What then prevents [pla.ta] from surfacing as a variant output for /palta/? The answer is that the non-observed /palta/→[pla.ta] violates more constraints than just LINEARITY-IO, and specifically it violates CONTIGUITY-V-IO. This constraint ranks above the cut-off and rules out metathesis as a possible repair for a NOCODA violation (cf. §3.4.1). Ranking a faithfulness constraint below the cut-off does not necessarily mean that rampant variable unfaithfulness will follow.

4.3 Variation only of comparatively well-formed candidates

In ROE, only constraints that rank above the cut-off can rule out a candidate as ungrammatical. On the other hand, a variable output candidate can only be disfavoured by constraints that rank below the cut-off. Consequently, a constraint that is used to rule out a candidate as ungrammatical for one input can never be violated by a variant for another input. This is one of the fundamental differences between ROE and other theories of variation in OT. In the theories of Anttila and Boersma & Hayes, it is possible that a constraint violated by a variant for some input is used to rule out a candidate as ungrammatical for a different input. The tableau in (40) gives a schematic example of this kind of situation in the non-ROE models of OT.

(40) a.

/input$_1$/	ℂ1	ℂ2	ℂ3
☞ i. cand$_1$	*		
☞ ii. cand$_2$		*	

b.

/input$_2$/			
i. cand$_1$	*!		
ii. cand$_2$		*!	
☞ iii. cand$_3$			*

The non-ROE theories explain variation as the consequence of rerankable constraints. In (40), {ℂ1, ℂ2} are rerankable relative to each other, but both outrank ℂ3. For /input$_1$/, the two best candidates violate ℂ1 and ℂ2, and because ℂ1 and ℂ2 are rerankable both of these candidates are grammatical. For /input$_2$/, however, there is a candidate that violates only ℂ3, so that candidates violating ℂ1 and ℂ2 are now ungrammatical. In ROE, this is not possible. Because ℂ1 and ℂ2 are violated by observed variants, these constraints must rank below the cut-off, and once they rank below the cut-off they cannot rule out a candidate as ungrammatical.

Boersma & Hayes's (2001) analysis of Ilokano provides a real-life example. A [ʔw] sequence variably undergoes metathesis, so that /taʔo+en/ can be pronounced as [taʔ.wen] or [taw.ʔen]. Boersma & Hayes explain

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

this by allowing *ʔ]$_\sigma$ and LINEARITY-IO to rerank freely. However, tautosyllabic consonant clusters are never repaired by metathesis – i.e. /pjano/ always surfaces as [pja.no] and never as [paj.no]. To account for this, they rank *COMPLEX below LINEARITY-IO. This is equivalent to the situation in (40) above, and is illustrated in (41). LINEARITY-IO is violated by a variant of the input /taʔo+en/, but it rules out a non-observed candidate for the input /pjano/.

(41)

	/taʔo+en/	LINEARITY-IO	*ʔ]$_\sigma$	*COMPLEX
a.	☞ i. taʔ.wen	*		
	☞ ii. taw.ʔen		*	
b.	/pjano/			
	☞ i. pja.no			*
	ii. paj.no	*!		

If examples such as these do exist, they would be a problem for ROE. In the discussion above (§3.4.1), I showed that an alternative analysis is possible for Ilokano. The ungrammaticality of [paj.no] is not due to LINEARITY-IO, but to a different constraint. In a ROE analysis, LINEARITY-IO ranks below the cut-off, since it disfavours the observed variant [taw.ʔen] (from /taʔo+en/). It cannot then be used to rule out the ungrammatical [paj.no] (from /pjano/). I used the constraint CONTIGUITY-IO-V (Hume 1998) for this purpose. The ROE account of these data is shown in (42).

(42)

	/taʔo+en/	CONTIGUITY-IO-V	LINEARITY-IO	*ʔ]$_\sigma$	*COMPLEX
a.	☞ i. taʔ.wen		*		
	☞ ii. taw.ʔen			*	
b.	/pjano/				
	☞ i. pja.no				*
	ii. paj.no	*!	*		

Examples such as these are what will provide the true test cases for ROE. For ROE to be accepted as a possible model of variation, it must be shown that a viable alternative analysis exists for all examples that are currently analysed along the lines of (40). In this paper, I have shown that this is possible for Ilokano, and I claim that it should be possible for other similar examples. However, ultimately only more research will determine whether this is truly possible.

This difference between ROE and the other theories of variation runs deeper than just the practicalities of the analysis. In ROE, variation is only

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

possible between candidates that are all relatively well-formed (candidates that are disfavoured only by constraints that rank at the bottom end of the hierarchy). This implies the following as a motivation for variation. For a certain input the second (and third, etc.) best candidate is disfavoured only by very low-ranking constraints. Although it is not the best candidate, its violations are not serious enough to disqualify it as a possible output. The alternative theories of variation do not share this assumption. In these theories, variation is possible even between two candidates that are both rather ill-formed (that are both disfavoured by very high-ranking constraints, as long as these constraints are rerankable with regard to each other). There is therefore no implicit connection between relative well-formedness and the possibility that variation will be observed.

4.4 Other domains of application for ROE

This paper has shown how variation can be accounted for by assuming that language users have access to the richer information structure imposed by EVAL on the candidate set. There is also evidence from other aspects of phonological grammar and language use that this richer information structure exists and that language users have access to this information. I briefly mention two of these here.

Language users have the ability to judge different forms for their well-formedness. These judgements are usually more fine-grained than just a two-level distinction between the best form and the non-best forms. Silva (2004) collected judgements on the prosodic phrasing of Korean SOV sentences. He presented Korean speakers with four phrasing options for each sentence, [SOV], [S]-[OV], [SO]-[V] and [S]-[O]-[V]. The subjects in his experiment did not simply make a two-way distinction between the best phrasing and the three non-best phrasings. Rather, they showed the following gradient preference: [S]-[OV] > [SOV] > [S]-[O]-[V] > [SO]-[V]. Silva explains this with a model that is practically identical to ROE. He says: 'it is assumed that the relative preferences can be formally captured by allowing constraint ranking … to derive a corresponding set of candidate rankings' (2004: 27). In the OT grammar that he develops, the [S]-[OV] phrasing is selected as the best, [SOV] as second best, etc. See also Berent *et al*. (2001a) and Berent *et al*. (2001b) for similar examples from Hebrew, and Coetzee (to appear) for examples from English.

The idea that the grammar can impose a multi-level well-formedness hierarchy on a set of candidates can also be used to account for the structure observed in the lexicon of some languages. The constraints on the co-occurrence of homorganic consonants in Arabic words are well known (Greenberg 1950, Frisch *et al*. 2004, etc.). For instance, Arabic allows words with contiguous coronals. But words with coronals that agree in sonorancy and continuancy are much scarcer than words with coronals

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

agreeing only in sonorancy – i.e. fewer [TaDaKa] words than [TaSaKa] words.[18] Words such as [TaLaKa], with coronals that differ in sonorancy and continuancy, are even more frequent. Coetzee & Pater (2006) analyse this kind of pattern with OCP-constraints: OCP-Cor[αson][βcont] is violated by [TaDaKa], OCP-Cor[αson] by [TaSaKa] and general OCP-Cor by [TaLaKa]. The proposal is then that when new words are added to the lexicon, forms that are more well-formed are more likely to be added than forms that are less well-formed. With the ranking OCP-Cor[αson][βcont] $\gg$ OCP-Cor[αson] $\gg$ OCP-Cor, the grammar imposes the well-formedness ranking [TaLaKa] > [TaSaKa] > [TaDaKa] on the three word-types, corresponding to the frequency with which the different word-types are observed in the lexicon.

Examples like these serve as additional motivation that the information about the well-formedness relationship between non-optimal candidates is relevant, and that language users can access and use this information. Models like ROE can increase our understanding of these kinds of phenomena.

5 Conclusion

An OT grammar generates more information than simply what the best candidate for some input is. It also generates information about the well-formedness relationship between the non-best candidates. This constitutes a fundamental difference between OT and most other generative theories of grammar. In other generative theories, the grammar generates a single grammatical output form for every input, but is completely silent about other possible outputs and how they are related to each other in terms of their well-formedness. In classical OT, this difference between OT and other generative theories does not receive any attention. In fact, in classical OT the additional information generated by the grammar is considered to be irrelevant and is hence ignored. In this paper, I have argued that this additional information is not irrelevant, but rather that language users have access to this information and that it can shape their linguistic performance. I have shown that this information can be used to explain phonological variation.

If the additional information generated by OT grammars were truly irrelevant, then these grammars would be more powerful than necessary. The generation of this information could then be used as an argument against the general design of OT grammars. However, as shown in this paper, this additional information is not irrelevant. Rather than a problem for OT, it can be interpreted as confirmation that the general architecture

[18] Capital letters here represent root consonants. [T] is a voiceless coronal stop, [D] a voiced coronal stop, [K] voiceless dorsal stop, [L] a coronal sonorant and [S] a voiceless coronal fricative.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

of an OT grammar is on the right track. OT grammars generate more information than other generative theories of grammar, and hence can more easily account for some aspects of linguistic performance than other generative theories.

REFERENCES

Anderson, John M. & Colin J. Ewen (1987). *Principles of dependency phonology*. Cambridge: Cambridge University Press.

Anttila, Arto (1997). Deriving variation from grammar. In Hinskens *et al*. (1997). 35–68.

Beckman, Jill N. (1997). Positional faithfulness, positional neutralisation and Shona vowel harmony. *Phonology* **14**. 1–46.

Beckman, Jill N. (1999). *Positional faithfulness : an Optimality Theoretic treatment of phonological asymmetries*. New York: Garland.

Berent, Iris, Daniel L. Everett & Joseph Shimron (2001a). Do phonological representations specify variables? Evidence from the Obligatory Contour Principle. *Cognitive Psychology* **42**. 1–60.

Berent, Iris, Joseph Shimron & Vered Vaknin (2001b). Phonological constraints on reading: evidence from the Obligatory Contour Principle. *Journal of Memory and Language* **44**. 644–665.

Blevins, Juliette (1995). The syllable in phonological theory. In John A. Goldsmith (eds.) *The handbook of phonological theory*. Cambridge, Mass. & Oxford: Blackwell. 206–244.

Boersma, Paul (1998). *Functional phonology : formalizing the interactions between articulatory and perceptual drives*. PhD dissertation, University of Amsterdam.

Boersma, Paul & Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *LI* **32**. 45–86.

Carpenter, Angela C., Andries W. Coetzee & Paul de Lacy (eds.) (2002). *Papers in Optimality Theory II*. Amherst: GLSA.

Clements, George N. (1986). Syllabification and epenthesis in the Barra dialect of Gaelic. In Koen Bogers, Harry van der Hulst & Martin Mous (eds.) *The phonological representation of suprasegmentals*. Dordrecht: Foris. 317–336.

Coetzee, Andries W. (2004). *What it means to be a loser : non-optimal candidates in Optimality Theory*. PhD dissertation, University of Massachusetts, Amherst.

Coetzee, Andries W. (to appear). Grammar is both categorical and gradient. In Steven Parker (ed.) *Phonological argumentation*. London: Equinox.

Coetzee, Andries W. & Joe Pater (2006). Lexically ranked OCP-Place constraints in Muna. Ms, University of Michigan & University of Massachusetts, Amherst. Available as ROA-842 from the Rutgers Optimality Archive.

Crosswhite, Katherine (2000). T*he analysis of extreme vowel reduction*. *UCLA Working Papers in Linguistics* **4** *: Papers in Phonology*. 1–12.

Crosswhite, Katherine (2001). *Vowel reduction in Optimality Theory*. New York: Routledge.

de Lacy, Paul (2002). *The formal expression of markedness*. PhD dissertation, University of Massachusetts, Amherst.

Freitas, M. João (2003). The acquisition of Onset clusters in European Portuguese. *Probus* **15**. 27–46.

Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe (2004). Similarity avoidance and the OCP. *NLLT* **22**. 179–228.

Greenberg, Joseph H. (1950). The patterning of root morphemes in Semitic. *Word* **6**. 162–181.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

Guy, G. (2003). Variationist approaches to phonological change. In Brian D. Joseph & Richard D. Janda (eds.) *The handbook of historical linguistics*. Malden, Mass. & Oxford: Blackwell. 369–400.
Hayes, Bruce (1989). Compensatory lengthening in moraic phonology. *LI* **20**. 253–306.
Hayes, Bruce & May Abad (1989). Reduplication and syllabification in Ilokano. *Lingua* **77**. 331–374.
Hayes, Bruce, Bruce Tesar & Kie Zuraw (2004). *OTSoft 2.1*. Software package, http://www.linguistics.ucla.edu/people/hayes/otsoft/.
Hinskens, Frans, Roeland van Hout & W. Leo Wetzels (eds.) (1997). *Variation, change and phonological theory*. Amsterdam & Philadelphia: Benjamins.
Hoard, James E. (1971). Aspiration, tenseness, and syllabication in English. *Lg* **47**. 133–140.
Hume, Elizabeth (1998). Metathesis in phonological theory: the case of Leti. *Lingua* **104**. 147–186.
Hume, Elizabeth (2001). Metathesis: formal and functional considerations. In Hume *et al*. (2001). 1–25.
Hume, Elizabeth & Jeff Mielke (2001). Consequences of word recognition for metathesis. In Hume *et al*. (2001). 135–158.
Hume, Elizabeth, Norval Smith & Jeroen van de Weijer (eds.) (2001). *Surface syllable structure and segment sequencing*. Leiden: Holland Institute of Generative Linguistics.
Itô, Junko & Armin Mester (1999). The phonological lexicon. In Natsuko Tsujimura (ed.) *The handbook of Japanese linguistics*. Malden, Mass. & Oxford: Blackwell. 62–100.
Itô, Junko & Armin Mester (2003). On the sources of opacity in OT: coda processes in German. In Caroline Féry & Ruben van de Vijver (eds.) *The syllable in Optimality Theory*. Cambridge: Cambridge University Press. 271–303.
Jäger, Gerhard (2003). *Maximum entropy models and Stochastic Optimality Theory*. Available as ROA-625 from the Rutgers Optimality Archive.
Kariņš, A. Krišjānis (1995). Vowel deletion in Latvian. *Language Variation and Change* **7**. 15–34.
Kenstowicz, Michael (1994). Syllabification in Chukchee: a constraints-based analysis. Available as ROA-30 from the Rutgers Optimality Archive.
Landman, Meredith (2002). Morphological contiguity. In Carpenter *et al*. (2002). 139–169.
Lombardi, Linda (2001). Why Place and Voice are different: constraint-specific alternations in Optimality Theory. In Linda Lombardi (ed.) *Segmental phonology in Optimality Theory: constraints and representations*. Cambridge: Cambridge University Press. 13–45.
Lombardi, Linda (2003). Markedness and the typology of epenthetic vowels. Available as ROA-578 from the Rutgers Optimality Archive.
Łubowicz, Anna (2002). Derived environment effects in Optimality Theory. *Lingua* **112**. 243–280.
McCarthy, John J. (2002). Comparative markedness. In Carpenter *et al*. (2002). 171–246.
McCarthy, John J. (2003). Comparative markedness. *Theoretical Linguistics* **29**. 1–51.
McCarthy, John J. & Alan Prince (1990). Prosodic morphology and templatic morphology. In Mushira Eid & John McCarthy (eds.) *Perspectives on Arabic Linguistics II*. Amsterdam & Philadelphia: Benjamins. 1–54.
Mateus, Maria Helena & Ernesto d'Andrade (2000). *The phonology of Portuguese*. Oxford: Oxford University Press.
Moreton, Elliott (1999). Non-computable functions in Optimality Theory. Ms, University of Massachusetts, Amherst. Available as ROA-364 from the Rutgers Optimality Archive.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press

Nagy, Naomi & Bill Reynolds (1997). Optimality Theory and variable word-final deletion in Faetar. *Language Variation and Change* **9**. 37–55.
Parker, Steve (2002). *Quantifying the sonority hierarchy*. PhD dissertation, University of Massachusetts, Amherst.
Parkinson, Stephen (1988). Portuguese. In Martin Harris & Nigel Vincent (eds.) *The Romance languages*. London & Sydney: Croom Helm. 131–169.
Pater, Joe (1999). Austronesian nasal substitution and other NÇ effects. In René Kager, Harry van der Hulst & Wim Zonneveld (eds.) *The prosody–morphology interface*. Cambridge: Cambridge University Press. 310–343.
Prince, Alan (2002). Arguing optimality. In Carpenter *et al*. (2002). 269–304.
Prince, Alan & Paul Smolensky (1993). *Optimality Theory : constraint interaction in generative grammar*. Ms, Rutgers University & University of Colorado, Boulder. Published 2004, Malden, Mass. & Oxford: Blackwell.
Reynolds, William (1994). *Variation and phonological theory*. PhD dissertation, University of Pennsylvania.
Riggle, Jason (2004). Contenders and learning. *WCCFL* **23**. 649–662.
Rogers, Francis Millet (1949). Insular Portuguese pronunciation: central and western Azores. *Hispanic Review* **17**. 47–70.
Rousseau, Pascale & David Sankoff (1978). Advances in variable rule methodology. In David Sankoff (ed.) *Linguistic variation : models and methods*. New York: Academic Press. 57–69.
Samek-Lodovici, Vieri & Alan Prince (1999). Optima. Ms, University College London & Rutgers University. Availabe as ROA-363 from the Rutgers Optimality Archive.
Sankoff, David (1988). Variable rules. In Ulrich Ammon, Norbert Dittmar & Klaus J. Mattheier (eds.) *Sociolinguistics : an international handbook of the science of language and society*. Vol. 2. Berlin & New York: de Gruyter. 984–997.
Silva, David James (1997). The variable deletion of unstressed vowels in Faialense Portuguese. *Language Variation and Change* **9**. 295–308.
Silva, David James (1998). Vowel lenition in São Miguel Portuguese. *Hispania* **81**. 166–178.
Silva, David James (2004). Quantified optimality and the phonological parsing of Korean SOV sentences. *Korean Linguistics* **12**. 25–54.
Tesar, Bruce & Paul Smolensky (1998). Learnability in Optimality Theory. *LI* **29**. 229–268.
Tesar, Bruce & Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, Mass.: MIT Press.
Thomas, Earl W. (1974). *A grammar of Spoken Brazilian Portuguese*. Nashville: Vanderbilt University Press.
Vanoverbergh, Morice (1955). *Iloko grammar*. Bauko, Philippines: Catholic School Press.

https://doi.org/10.1017/S0952675706000984 Published online by Cambridge University Press