Phonology **18** (2001) 147–197. Printed in the United Kingdom

*Consonant cluster neutralisation and targeted constraints**

Colin Wilson
University of California, Los Angeles

In this paper, I propose an optimality-theoretic account of the generalisation that deletion processes that apply to intervocalic biconsonantal clusters canonically delete the first consonant (schematically, $VC_1C_2V \rightarrow VC_2V$). The approach to contextual neutralisation proposed here has two main components. First, I follow the licensing-by-cue framework (e.g. Steriade 1997) in identifying 'weak' elements as those without strong perceptual cues. Second, I argue that the constraints responsible for contextual neutralisation 'target' weak elements. This approach captures the deletion generalisation above, because the relevant targeted constraint prefers only the correct output VC_2V (from which the weak consonant C_1 has been removed), not the incorrect output VC_1V. Intuitively, the representation containing a weak element (VC_1C_2V) is compelled to neutralise to a representation that is perceptually very similar (VC_2V). The targeted-constraint approach is formalised by replacing the standard violation-based definition of OT optimisation with a new definition – which is equivalent except when 'targeted' constraints are involved – based on harmonic orderings. The approach is shown to extend to certain cases of (i) contextually determined feature neutralisation and (ii) phonological opacity.

1 Introduction

Recent work within Optimality Theory (OT; Prince & Smolensky 1993) has focused intensely on CONTEXTUAL NEUTRALISATION (e.g. Alderete 1999, Beckman 1998, Itô & Mester 1993, Kirchner 1998, Lombardi 1997, 1999, Padgett 1995, Steriade 1997, 2000; see also Steriade 1995 for a review of the empirical domain and non-OT analyses). In the present paper, I extend this line of research by proposing an OT account of the following

* For very helpful comments on portions of the work presented here, I would like to thank Mark Allen, Eric Baković, Laura Benua, Luigi Burzio, Lisa Davidson, Bob Frank, Matt Goldrick, Carlos Gussenhoven, Bruce Hayes, René Kager, John McCarthy, Paul Smolensky, Donca Steriade, the participants of the GLOW Workshop and the audience at the LSA annual meeting (Chicago, January 2000). Special thanks go to René Kager and Carlos Gussenhoven for organising the conference and acting as editors of this special issue, to Luigi Burzio, whose guidance was invaluable during the writing of the paper, and to the two anonymous reviewers and the associate editor, for their insightful comments and questions.

typological generalisation about contextually determined consonant deletion.

(1) *First consonant deletion*
Across languages, deletion processes that apply to intervocalic biconsonantal clusters consistently delete the *first* consonant (schematically, $VC_1C_2V \rightarrow VC_2V$).

This generalisation can be illustrated with the two well-known cases of consonant deletion in (2). Additional processes that conform to the generalisation are found in Basque (Hualde 1987), Carib (Gildea 1995) and Tunica (Haas 1946); the Carib and Tunica cases are introduced later in this section.

(2) *Two cases of first consonant deletion*
a. *Diola-Fogny* (Sapir 1965, Kiparsky 1973, Steriade 1982, Itô 1986, Lombardi 1997)

/le**t**+**k**u+jaw/	le**k**ujaw	'they won't go'
/kutɛ**b** **s**inaŋas/	kutɛ**s**inaŋas	'they carried the food'
/ɛkɛ**t** **b**o/	ɛkɛ**b**o	'death there'

b. *West Greenlandic* (Rischel 1974, Fortescue 1980, Jenkins 1984)

/qani**k**+**l**erpoq/	qani**l**erpoq	'begins to approach'
/ukiju**q**+**t**uqaq/	ukiju**t**uqaq	'old year'
/angu**t**i+**k**ulak/	angu**k**ulak	'he goat'

In each of these input–output mappings, the consonant that is deleted is the one that would occupy the first position of an illegal cluster if the input were realised completely faithfully. For example, the Diola input /let+ku+jaw/ is realised as [lekujaw], with deletion of the input consonant /t/. The fully faithful output would be *[letkujaw], but intervocalic obstruent–obstruent clusters (here [tk]) are not legal in this language.[1] Similarly, the Greenlandic input /qanik+lerpoq/ is realised as [qanilerpoq], with deletion of the input consonant /k/. Fully faithful morpheme combination would yield *[qaniklerpoq], but obstruent-initial intervocalic clusters (here [kl]) are illegal.[2]

[1] The only intervocalic consonant sequences permitted in Diola are those containing a nasal or liquid that is homorganic to a following obstruent stop. The language also permits intervocalic geminates, both nasal and non-nasal. Of the three examples cited in (2a), the first illustrates deletion within a word and the second two illustrate deletion at a word boundary. Sapir (1965: 19) also mentions that epenthesis occurs instead of deletion in 'slow and rather deliberate speech'; see §3 below for discussion of the relationship between deletion and epenthesis in the present approach.

[2] West Greenlandic allows certain liquid-initial clusters, as well as homorganic nasal–stop clusters and geminates, in intervocalic position. Consonant-initial suffixes in this language fall into two classes: suffixes in one class delete a stem-final consonant; suffixes in the other class fully assimilate the stem-final consonant to the suffix-initial consonant. Only the first class, illustrated in (2b), is considered here.

With respect to intervocalic obstruent–obstruent clusters, generalisation (1) has only one, very restricted, type of exception. It appears that some languages resolve illegal biconsonantal clusters that arise at root+suffix boundaries by deletion of the second (i.e. suffixal) consonant. The theory proposed here accounts for this exception (see §3.3.1) as well as the general pattern of first consonant deletion. (Note that some languages delete the first consonant even at root+suffix boundaries; witness the West Greenlandic examples above.)

With respect to other intervocalic biconsonantal clusters (i.e. those that contain a sonorant in first, second or both positions), the cross-linguistic evidence is less clear. There are certainly languages (such as West Greenlandic) in which generalisation (1) holds regardless of whether the input consonants are obstruents or sonorants. But there also seem to be languages in which illegal sonorant–obstruent and obstruent–sonorant clusters are all resolved by deletion of the sonorant; for example, see the deletion (or total progressive assimilation) process that applies to certain intervocalic clusters in Pali (Middle-Indic; Hankamer & Aissen 1974, Zec 1994, 1995, among others). Such cases could plausibly be analysed in terms of a preference for less sonorous syllable onsets (Hankamer & Aissen 1974, Zec 1994, 1995; see also Clements 1990, Prince & Smolensky 1993).

In this paper, however, I put aside the possible effects of sonority, focusing instead on the many other factors that logically could, but empirically do not, influence the decision about which of two consonants in an intervocalic cluster deletes. According to generalisation (1), a consonant is deleted or preserved based solely on the *position* that it would occupy in the cluster. But previous OT approaches to consonant deletion predict that the decision about which consonant deletes will instead be made – either universally or as a typological option – on grounds of *markedness*. Put simply, the problem addressed here is that previous OT approaches predict a pattern in which the more marked consonant deletes regardless of the position that it would occupy in the cluster. An abstract characterisation of this problem is given in the tableau below, where the pointing finger marks the correct output and the thumbs-down symbol marks the incorrectly optimal candidate.

(3) *Markedness constraint* **M** *incorrectly causes deletion of second* C

VC_1C_2V	ClusterCond	Max	**M**
a. VC_1C_2V	*!		
☞ b. VC_2V		*	*!
👎 c. VC_1V		*	

In this tableau, ClusterCond stands for any contextual markedness constraint that is violated by outputs that contain the consonant cluster C_1C_2. Max is the faithfulness constraint that is violated by segment deletion.[3] And **M** stands for any markedness constraint that is violated by the output [VC_2V], but not (or to a lesser degree) by the output [VC_1V]. (Whether or not [VC_1C_2V] violates **M** is irrelevant for the point made here, as the dark shading indicates. Note that there may in principle be other markedness constraints that favour [VC_2V] over [VC_1V], but we are considering here rankings in which those constraints, should they exist, are dominated by **M**.) The ranking ClusterCond ≫ Max forces one of the input consonants to be deleted, thereby avoiding the marked cluster [C_1C_2] in the output. (Other ways of avoiding the cluster, such as vowel epenthesis, are ignored for the moment, but are treated in the body of the paper; see especially §3.3.)

The important point made by this tableau is that neither ClusterCond nor Max prefers candidate [VC_2V], which is the correct output according to generalisation (1), over candidate [VC_1V]. Both of these candidates satisfy ClusterCond, because neither of them contains a consonant cluster; and both of them violate Max exactly once, because each of them lacks an output correspondent for one of the input consonants. Therefore, ClusterCond and Max *inevitably* pass the decision about which of the two candidates is optimal (i.e. about which of the two consonants is deleted) to other constraints.

Here the decision is passed to lowest-ranked **M**, which by hypothesis selects the incorrect output [VC_2V]. As should be obvious, the same output would be optimal if the input were /VC_2C_1V/ instead. Thus this hierarchy generates an unattested type of deletion process, according to which markedness (specifically, **M**-violation) rather than position governs which consonant deletes.

A concrete illustration of the problem is provided in the following tableau, which is based on the first Diola example in (2a) and in which the place-markedness constraint *Pl(lab,dor) instantiates **M**. (Only the constraint violations that differentiate the candidates under consideration are recorded in the tableau, a practice I observe throughout the paper.)

[3] Max is part of the Correspondence Theory of faithfulness (McCarthy & Prince 1995), which I adopt throughout the paper. The following constraints in particular will be employed (my definitions).

(i) Max: A segment in the input must have a correspondent in the output (no deletion).

Dep: A segment in the output must have a correspondent in the input (no epenthesis).

Ident: Corresponding segments must be featurally identical (no feature change).

Contiguity: If two segments are contiguous, then their correspondents (if any) must also be contiguous.

Linearity. If segment α precedes segment β, then the correspondent (if any) of α must also precede the correspondent (if any) of β (no metathesis).

Note that Ident can be specified for particular features (Ident(voice), Ident(place), etc.). When correspondence relations are not obvious from the context, they are indicated with subscripts. Unless explicitly mentioned, every faithfulness constraint in the paper regulates input–output correspondence.

(4) *Incorrect deletion of non-coronal* C *in Diola*

let+ku+jaw	ClusterCond	Max	*Pl(lab,dor)	*Pl(cor)
a. letkujaw	*!		*	*
☞ b. lekujaw		*	*!	
☹ c. letujaw		*		*

The ranking *Pl(lab,dor) ≫ *Pl(cor), which is part of the universal place-markedness subhierarchy (Prince & Smolensky 1993, Lombardi 1997), states that Labial and Dorsal are intrinsically more marked place features than Coronal. This markedness relationship is supported by evidence from segmental inventories (Prince & Smolensky 1993: §§9.1.2, 9.2), epenthetic segments[4] and default place specifications (McCarthy & Prince 1993, Lombardi 1997), reduplication (Alderete *et al.* 1999) and contextual restrictions on labial and dorsal consonants (Prince & Smolensky 1993, Zoll 1998); see also the contributions to Paradis & Prunet (1991).

As shown in the tableau, the place-markedness subhierarchy incorrectly selects *[letujaw] as the optimal candidate, because that candidate contains a segment (namely, [t]) with a less marked place specification than the corresponding segment (namely, [k]) in [lekujaw]. In general, the deletion process that is generated by this hierarchy resolves clusters that contain a coronal and a non-coronal consonant by deleting the *non-coronal* consonant, regardless of its position in the cluster. (When the members of a cluster are both coronal consonants, or both non-coronal consonants, then the hierarchy predicts that either one can be deleted.) This predicted deletion process is not the one observed in Diola, and moreover violates the typological generalisation in (1).

This is just one example of a very general problem. Notice in particular that nothing essential changes if ClusterCond is replaced with a constraint against certain types of coda consonants (call it CodaCond; see Itô 1986, Lombardi 1997, Prince & Smolensky 1993). With the syllabification [let.ku.jaw], the obstruent [t] in some sense 'causes' the violation of CodaCond. And indeed the processes at work in examples such as those in (2) could be pre-theoretically described as coda deletion. But satisfaction of CodaCond, like satisfaction of ClusterCond, is actually compatible

[4] An anonymous reviewer asserts that claims about markedness based on the quality of epenthetic segments are mitigated by the fact that (for example) glottal stop and schwa are canonical epenthetic segments, but are excluded from many phoneme inventories. While correct at an empirical level, this observation seems to show only that there is not one sense, or dimension, of markedness, but many. For example, the very perceptual weakness of glottal stop and schwa that makes them well-suited for epenthesis (Steriade 2000) may also make them ill-suited for membership in an inventory of contrastive segments. The points made in the text depend only on the existence of markedness constraints, not on the apparently incorrect claim that a segment that is unmarked on one dimension (e.g. as an epenthetic segment) must be unmarked on all dimensions (e.g. as a phoneme).

with deletion of either consonant. The two candidates [le.ku.jaw] and [le.tu.jaw] contain exactly the same set of coda-affiliated segments, and therefore receive exactly the same violations on any CodaCond constraint. Consequently, the choice between these two candidates is passed – again incorrectly – to the place-markedness subhierarchy.

Notice also that the decision to use the place-markedness subhierarchy to illustrate the problem was arbitrary. Any non-contextual markedness constraint that prefers certain consonants over others can be used to generate a deletion process that violates (1). To give one more example, substituting *[−son, +voice] for the place-markedness subhierarchy in the hierarchy above yields a process that deletes a *voiced obstruent* (if any) regardless of its position in the cluster. Applied to the second two Diola inputs in (2a), this process gives /kutɛb sinaŋas/ → [kutɛsinaŋas] (correct output; C_1 deleted) and /ɛkɛt bo/ → *[ɛkɛto] (incorrect output; C_2 deleted). Evidence for the markedness constraint *[−son, +voice] (i.e. for the markedness of voiced obstruents relative to voiceless obstruents) is comparable in scope to the evidence for the place-markedness subhierarchy (Keating 1984, Lombardi 1997, 1999).[5]

Before turning to my solution, I note an additional type of data that will be relevant for the discussion of possible alternative solutions. All of the examples in (2) have the property that the undeleted consonant is the only member of the potential cluster that is *prevocalic* in the input. Looking only at examples of this kind, one might conclude that consonants that are prevocalic in the input are subject to a special type of faithfulness, and that this accounts for the fact that only the first consonant can be deleted. That conclusion would be incorrect, however, given that there are cases of consonant deletion in which *both* of the consonants are prevocalic in the input (i.e. in which vowel deletion descriptively 'feeds' consonant deletion). Importantly, these cases also obey generalisation (1), as illustrated by the following examples from Carib (Northern Carib; Hoff 1968, Gildea 1995) and Tunica (Haas 1941, 1946).

[5] Furthermore, the problem would not be solved by eliminating all non-contextual markedness constraints. In addition to raising the question of how to reanalyse the evidence that motivated these constraints (see references cited in the text), this move underestimates the scope of the problem in two ways. First, simply removing the non-contextual markedness constraints from the two hierarchies above leaves us with the incorrect prediction that *either* consonant can be deleted to produce a grammatical output. Without some further constraint that prefers deletion of the first consonant, the analysis is still incomplete. Second, *contextual* markedness constraints can also give rise to unattested deletion processes. For example, consider the contextual markedness constraint that is responsible for intervocalic voicing of (certain) obstruents in Korean, Mohawk and other languages. The existence of this constraint – whatever its precise formulation – predicts a process that preferably deletes *voiceless* obstruents, because the consonant that remains after deletion is by hypothesis intervocalic. This predicted process is the mirror image of the one that *[−son, +voice] gives rise to, and is equally unnatural. Although for reasons of space I concentrate only on the problem as it is caused by non-contextual markedness constraints, the solution I propose appears to carry over straightforwardly to contextual markedness constraints as well.

(5) *Vowel deletion feeds first consonant deletion*

a. *Carib* (data from Gildea 1995: 91)

/s+enaa**pɨ**+**s**a/	senaa**s**a	'I eat it'
/s+enee**pɨ**+**s**a/	senee**s**a	'I bring it'

b. *Tunica* (data from Haas 1946: 343–345)

/tiˈtih**ki**+**t**ʔɛ/	tiˈtih**t**ʔɛ	'a river'
/tiˈtih**ki** **p**iˈrʔutakʔahča/	tiˈtih**p**iˈrʔutakʔahča	'it will turn into a bayou'

These examples are discussed in detail in §4.2. Even without detailed discussion, however, they rule out the hypothesis that deletion processes that apply to intervocalic clusters only delete consonants that are not prevocalic in the input. Indeed, I will argue further below that no analysis based solely on input properties can give a satisfactory account of the generalisation in (1).

I propose to solve the general problem identified above with a new OT approach to contextual neutralisation. Working within the licensing-by-cue framework of Steriade (1997) and others, I characterise 'weak' phonological elements as those that lack strong perceptual cues. Most relevantly here, consonants that are not released into a vowel are relatively 'weak', because they lack the cues in the forceful burst and formant transitions that a following vowel provides. I then argue that the constraints responsible for contextual neutralisation *target* such weak elements, in this specific sense: given a candidate containing a weak element α, these constraints prefer *only* the candidate that is identical except that α has been removed. More intuitively, a representation containing a weak element is only compelled to neutralise to a representation that is auditorily/perceptually very similar. For example, consider again the Diola input /let+ku+jaw/ from (2a) above. A certain contextual markedness constraint states that the fully faithful candidate [letkujaw], which contains a weak consonant ([t]), is worse (formally, less harmonic) than the alternative candidate [lekujaw], which is identical except that the weak consonant has been removed. In contrast, no contextual markedness constraint states that [letkujaw] is worse (less harmonic) than the alternative candidate [letujaw], from which the strong consonant ([k]) has been removed. In other words, only deletion of the first consonant is preferred by the set of contextual markedness constraints, once these are properly formulated. Because the faithfulness constraint MAX states that [letkujaw] is better (more harmonic) than both [lekujaw] and [letujaw], it follows that a consonant deletion process cannot remove the second consonant (as in [letujaw]), but only the first one (as in [lekujaw]). This analysis is formalised by replacing the standard violation-based definition of OT optimisation with a new definition – which is equivalent to the standard definition except when 'targeted' constraints are involved – that is based on harmonic orderings. The resulting theory yields a more restrictive typology of consonant cluster neutralisation in general, and in particular captures the generalisation about consonant deletion in (1).

1.1 Outline of the paper

In the next section (§2), I develop the idea of 'targeted' markedness constraints and provide an introduction to the licensing-by-cue approach to contextual neutralisation. I also give an informal analysis of deletion processes like those in (2). In §3, I formalise that analysis, and most importantly show that the targeted-constraint approach captures the generalisation in (1). This result is first established with respect to relatively restricted candidate sets (those limited to candidates created by segment deletion and/or feature change) and is then shown to be maintained when a much wider range of candidate outputs is considered. Another important result of §3 is that the targeted constraint proposed in §1, which only directly prefers consonant deletion, nevertheless plays a crucial role in the analysis of other consonant cluster 'repairs' such as vowel epenthesis.

In §4, I discuss possible alternative accounts of generalisation (1), including a number of approaches based on contextual (or 'positional') faithfulness constraints. All of the alternatives are shown to be theoretically problematic and empirically inadequate. In §5, I extend the targeted-constraint approach to two additional empirical domains. First, targeted constraints are used to capture the pattern of voice neutralisation in Lithuanian (see also §2.1 below). Second, targeted constraints are shown to provide a novel account of cases of phonological opacity, and are deployed in this capacity to account for one case of opacity observed in Nancowry. In §6, I summarise the main point of the paper and state the conclusions.

2 Targeted constraints and consonant deletion

As mentioned above, instead of searching for some *input* property that can predict which consonant deletes in examples such as those in (2) and (5), I propose that we examine the relationships among the relevant candidate *outputs*. More specifically, I claim that standard OT formulations of contextual markedness constraints like ClusterCond lack a notion of inter-candidate similarity that is crucial for an understanding of generalisation (1).

To illustrate the idea, consider again the Diola input /let+ku+jaw/. The fully faithful (but ungrammatical) candidate output for this input, [letkujaw], violates the postulated contextual markedness constraint ClusterCond. In standard OT, this violation entails that [letkujaw] is less harmonic, according to ClusterCond, than *any* alternative candidate that does not violate the constraint. The alternative candidates preferred by ClusterCond can differ in many ways from [letkujaw], and may even be completely distinct from it, as the following diagram indicates.

(6) *Harmonic orderings asserted by standard contextual markedness*

	[letujaw] (C_2 *deleted*) ↑	
[ba], etc. ← (*other changes*)	[letkujaw] ↓	→ [lekujaw] (C_1 *deleted*)
	[letikujaw] (V *epenthesised*)	

The arrows in this diagram point in directions of increasing harmony according to ClusterCond. They lead away from the fully faithful candidate [letkujaw] to candidates that 'repair' the intervocalic consonant cluster in various ways: by deletion of one of the consonants; by vowel epenthesis; and even by radical reduction to a default representation such as [ba]. Indeed, Gen provides an infinite set of 'repair' candidates that ClusterCond prefers over [letkujaw]. Thus the standard contextual markedness constraint itself does not distinguish the grammatical 'repair' (i.e. first consonant deletion, yielding [lekujaw]) from any of a myriad of logically possible alternatives.

In contrast to the standard picture of contextual markedness, I propose that constraints such as ClusterCond are *targeted*, in the sense that they prefer only certain possible 'repairs' of the structures that violate them. The following diagram, which is a revision of (6), introduces both the general notion of a targeted constraint and the particular way in which the contextual markedness constraints that drive consonant deletion must be targeted.

(7) *Harmonic orderings asserted by targeted contextual markedness*

	[letujaw] (C_2 *deleted*)	
[ba], etc. (*other changes*)	[letkujaw]	→ [lekujaw] (C_1 *deleted*)
	[letikujaw] (V *epenthesised*)	

As in the previous diagram, the arrow points in a direction of increasing harmony according to the contextual markedness constraint. Here, however, there is only one such direction: with respect to the candidate [letkujaw], which contains an offending consonant cluster, the targeted contextual markedness constraint prefers *only* the candidate [lekujaw], from which the first member of the cluster has been deleted. Of course, [lekujaw] is the grammatical output for the input /let+ku+jaw/ in Diola;

moreover, deletion of the first member of the cluster in [letkujaw] is the 'repair' that conforms to the typological generalisation in (1). The targeted contextual markedness constraint that distinguishes this 'repair' from all the others is preliminarily defined as follows (see §2.1 for the precise definition).

(8) NoWeakConsonant (NoWeakCons)
(*targeted contextual markedness constraint – to be revised*)
Let x be any candidate and α be any preconsonantal (or 'weak') consonant in x. If candidate y is exactly like x except that α has been removed, then y is more harmonic than x (i.e. $y \succ x$).

Candidate [lekujaw] is exactly like candidate [letkujaw] except that the preconsonantal consonant (namely, [t]) in the latter has been removed. Therefore, NoWeakCons states that [lekujaw] is more harmonic than [letkujaw] (i.e. [lekujaw] $\succ$ [letkujaw]). In contrast, candidate [letujaw] does *not* differ from [letkujaw] in the prescribed manner. Therefore, NoWeakCons is silent about the relative harmonic ordering of these two candidates (i.e. the constraint states neither that [letujaw] is more harmonic than [letkujaw] nor that [letkujaw] is more harmonic than [letujaw]); and the same holds true, as indicated by the shading, for all the other 'repair' candidates that are shown (or implied) in diagram (7).

The introduction of targeted constraints such as NoWeakCons into the OT framework raises a number of formal and substantive issues. On the formal side, the standard violation-based definition of harmonic ordering (Prince & Smolensky 1993: ch. 5) is not sufficiently expressive to accommodate targeted constraints. On the substantive side, the question immediately arises of why targeted constraints such as NoWeakCons should prefer particular 'repairs' over others. Leaving the formal issues for later in the paper, I now turn to the substantive ones.

2.1 Perceptual cues and the weak element principle

As mentioned above, the main component that is missing from previous formulations of contextual markedness is a notion of *similarity* among candidate outputs. In particular, I propose that contextual markedness constraints such as NoWeakCons only compare candidates that are sufficiently auditorily/perceptually similar. This proposal builds on the licensing-by-cue approach to contextual neutralisation of Jun (1995), Kirchner (1998), Silverman (1995) and Steriade (1997); see also the closely related proposals of Flemming (1995).

Work within the licensing-by-cue approach has argued convincingly that the phonological contexts in which a given contrast is neutralised (as opposed to 'licensed') are defined primarily by the absence of strong perceptual cues for that contrast. For example, the strongest perceptual cues for a voice contrast in obstruents (e.g. voice onset time) are provided by release into a following modally voiced sonorant (Keating 1984, Kingston 1990). Steriade (1997) makes use of this fact in her analysis of a pattern of voice neutralisation that is illustrated with the Lithuanian data in (9).

(9) *Contextual voice neutralisation in Lithuanian* (Steriade 1997: 17–18, and references cited there)[6]

a. *Voice contrast maintained before sonorants*

/au**k**le/	au**k**le	'governess'
/au**g**lingas/	au**g**lingas	'fruitful'

b. *Voice contrast neutralised word-finally*

/dau**g**/	dau**k**	'much'
/ka**d**/	ka**t**	'that'

In Lithuanian, as in many languages, voiced and voiceless obstruents contrast before sonorants (9a), but not before obstruents or in word-final position (9b). In other words, the contrast is maintained in contexts where strong perceptual cues are present, and neutralised elsewhere. Steriade (1997: Part I) analyses this pattern with a set of contextual markedness constraints, each of which bans contrastive voice specifications in a particular context. In accordance with the basic insight of the licensing-by-cue approach – that better cued contrasts are less marked – the constraint that bans contrastive voice specifications in the word-final context (*[α voice]/V__ #) universally dominates the constraint that bans contrastive voice specifications in the pre-sonorant context (*[α voice]/V__ [+son]). Languages like Lithuanian rank the faithfulness constraint on voice specifications (here, PRESERVE[voice]) between these two contextual markedness constraints. The following tableau, reproduced from Steriade's paper, shows how this hierarchy generates word-final voice neutralisation.

(10) *Contextual markedness neutralises voice word-finally in Lithuanian*

daug	*[α vce]/V__#	PRES[vce]	*[α vce]/V__[+son]
a. daug	*!		
☞ b. dauK		*	
c. dauk	*!	*	

Following Steriade's discussion, the outcome of word-final neutralisation is identified here as a segment that has no phonetic specification for voice (transcribed [K]).[7] Although I do not provide a tableau to show this, it should be clear that the same constraint hierarchy that neutralises

[6] Diacritics indicating rhyme length and pitch accent are omitted here.

[7] Steriade (1997) asserts that the outcome of neutralisation is a segment without a phonetic voice specification, based on phonetic studies in various languages (see the

the voice contrast word-finally (by the ranking (*[α voice]/V__ # ≫ PRESERVE[voice]) preserves the contrast in presonorant contexts (by the ranking PRESERVE[voice] ≫ *[α voice]/V__ [+son]).[8]

The licensing-by-cue framework also sheds light on cases of contextually determined consonant deletion such as those in (2). Focusing on the core cases, in which deletion is used to avoid a surface obstruent cluster (e.g. Diola /let+ku+jaw/ → [lekujaw]), the following facts are relevant. Experimental work has shown that the perceptual cues provided by the release of a consonant are particularly important for signalling the existence and the features of that consonant (Ohala 1990, Kingston 1990, 1994). Converging evidence comes from recent phonological work, which shows that release plays a key role in restricting the inventory and distribution of contour segments (Steriade 1993) and in accounting for neutralisation and assimilation of place and voice features (Beckman 1998, Jun 1995, Lombardi 1997, 1999, Padgett 1995, Steriade 1997). Furthermore, both lines of research have identified consonants that are released into a sonorant – and in the best case into a vowel – as the 'strongest' ones, where 'strength' is measured (experimentally) in terms of perceptibility and (phonologically) in terms of the licensing of phonological contrasts. (Many languages prohibit obstruents from being released immediately before other obstruents, and release in word-final position is cross-linguistically variable; see Padgett 1995, Steriade 1993, 1997).[9]

In light of these facts, it is apparent that consonant deletion as in /let+ku+jaw/ → [lekujaw] serves the same basic purpose as voice neutralisation in (9a). Both processes eliminate phonological elements (segments or features) that would be poorly cued (or 'weak') in a more faithful output. With respect to consonant deletion, this is a restatement of generalisation (1) in terms of the licensing-by-cue approach to contextual neutralisation.

Within standard OT, however, the licensing-by-cue approach is unable to account for the generalisation about consonant deletion. This point is underscored by the following tableau, in which the stand-in contextual markedness constraint CLUSTERCOND of tableau (4) is replaced by the cue-based constraint *[−son]/ __ C ('no preconsonantal obstruent'; Zoll 1998).

references cited in Steriade 1997). The issue of whether neutralisation results in a phonetically unspecified segment or a [−voice] segment is orthogonal to the present discussion; I have transcribed the segment as unspecified for voice in order to maintain consistency with Steriade's work.

8 Steriade (1997: Part II) proposes to replace the markedness constraints that refer to segmental contexts (e.g. *[α voice]/V __ #) with constraints that refer directly to the absence of perceptual cues (roughly, *[α voice] in contexts where voicing lacks transitional cues). This change in formalisation does not substantially affect the present discussion.

9 The more general notion that is relevant here is that of the transitional cues provided by the context of a given segment. Thus the cues from an immediately following vowel are equally important for the licensing of continuant consonants, which are not formally 'released' in the aperture theory of Steriade (1993).

(11) *Incorrect deletion of non-coronal* C *in Diola*

let+ku+jaw	*[−son]/__C	Max	*Pl(lab,dor)	*Pl(cor)
a. letkujaw	*!		*	*
☞ b. lekujaw		*	*!	
☹ c. letujaw		*		*

The medial obstruent in both [lekujaw] and [letujaw] is immediately followed by a vowel (and is furthermore assumed to be released into that vowel, thus receiving strong perceptual cues). Consequently, *[−son]/ __ C is satisfied by both of these candidates, just as ClusterCond and CodaCond are. The decision about which input consonant deletes therefore falls once again to the place-markedness constraint *Pl(lab,dor), which selects the incorrect output [letujaw].

The standard OT conception of markedness must therefore be modified to correctly express the insights of the licensing-by-cue framework. Consider a surface representation such as [letkujaw], which contains a poorly cued consonant (i.e. a consonant that is not released into a following vowel of other sonorant). Suppose that such a representation is marked, not in any absolute sense, but only relative to another representation. In particular, suppose that [letkujaw] is marked only relative to [lekujaw], because (i) the two representations have very similar auditory/perceptual components – that is, they differ only with respect to a segment (namely, [t]) that has few auditory/perceptual cues – and (ii) [lekujaw] contains less structure (specifically, one fewer consonant) than [letkujaw]. The general principle that establishes markedness relations of this kind is given below.

(12) *Weak element principle*

A representation x that contains a poorly cued (or 'weak') element α is marked relative to the representation y that is identical to x except that α has been removed.

Informally, this principle states that, given two surface representations that sound basically the same (and which could therefore be easily confused by the hearer), the more complex representation is marked relative to the less complex representation.

I propose the weak element principle as the substantive basis for targeted constraints such as NoWeakCons (8). Indeed, all of the targeted constraints discussed in this paper should be thought of as being mechanically derived from the weak element principle in the following way. If (12) states that representation x is marked relative to representation y, then the corresponding targeted constraint asserts that y (as a candidate output) is more harmonic than x (as a candidate output). Individual targeted constraints are defined by selecting a value for α (e.g. consonant or [voice] feature) and setting a criterion for being poorly cued. To simplify the discussion, I will assume throughout the paper that a consonant is poorly cued/'weak' unless it is released into a vowel (but see the references cited above for more fine-grained theories of cueing). Thus the final version of the targeted constraint NoWeakCons is as follows.

(13) NoWeakCons
(*targeted contextual markedness constraint – final version*)
Let x be any candidate and α be any consonant in x that is not released by a vowel. If candidate y is exactly like x except that α has been removed, then y is more harmonic than x (i.e. $y \succ x$).[10]

2.2 Informal presentation of the account

In the preceding discussion, I introduced both the particular targeted constraint NoWeakCons (13) and the general principle that underlies such constraints (12). Here, I give a preliminary and informal presentation of how this theory accounts for the generalisation in (1), returning to a more formal presentation in §3.

Consider first the abstract input /VC_1C_2V/. The fully faithful candidate, [VC_1C_2V], contains a consonant that is not released into a following vowel (i.e. that is 'weak'), namely [C_1]. Therefore, by the weak element principle, NoWeakCons asserts that [VC_1C_2V] is less harmonic than the alternative candidate that is exactly the same except that [C_1] has been removed, namely [VC_2V]. The standard faithfulness constraint Max asserts the opposite harmonic ordering: it prefers candidate [VC_1C_2V], in which both of the input consonants have output correspondents, over candidate [VC_2V], in which only the second input consonant has an output correspondent. But if NoWeakCons dominates Max, the targeted markedness constraint overrides the faithfulness constraint and fixes the harmonic ordering between the two candidates as [VC_2V] $\succ$ [VC_1C_2V].

Of course, these are not the only two members of the candidate set. Another candidate that must be considered is the one that was incorrectly ruled optimal in tableau (3), namely [VC_1V]. This candidate is not identical to [VC_1C_2V] except that the preconsonantal/weak consonant [C_1] has been removed (i.e. removing [C_1] from [VC_1C_2V] and making no other change yields [VC_2V], not [VC_1V]). Therefore, NoWeakCons makes no assertion about their relative harmony. In terms of the earlier reasoning

[10] An anonymous reviewer suggests that the theory of targeted constraints could be made more restrictive by imposing the condition that a targeted constraint can designate a target, but cannot specify the change that applies to the target. Revising NoWeakCons according to this suggestion would yield a constraint that still targets weak consonants, but which prefers *any* change that eliminates a target while leaving the segments in the context of the target unaffected (i.e. not just deletion, but also vowel epenthesis, etc.). This suggested revision of the theory of targeted constraints does not appear to me to have any obvious flaws, and is therefore worthy of further investigation, but unfortunately I must leave this project for future research.

about the motivation for targeted constraints, [VC_1C_2V] and [VC_1V] are too auditorily/perceptually dissimilar to be compared by NoWeakCons.

In contrast, Max is a standard (i.e. untargeted) OT constraint, and is therefore perfectly capable of asserting a harmonic ordering between [VC_1V] and [VC_1C_2V]. In particular, Max asserts that [VC_1C_2V] is more harmonic than [VC_1V]. Combining the harmonic ordering established by Max with the one established by NoWeakCons yields a total ordering of the three candidates, with [VC_2V] at the top: [VC_2V] $\succ$ [VC_1C_2V] $\succ$ [VC_1V]. According to this ordering, the only optimal candidate is [VC_2V] (because it is the only candidate that is not less harmonic than some other candidate). And this declaration of optimality cannot be refuted by any lower-ranked markedness constraint **M** that prefers [VC_1V] over [VC_2V]. Therefore, NoWeakCons and Max together ensure that the first consonant of the cluster is deleted, in accordance with generalisation (1), even if the second consonant is more marked on some dimension.

This informal optimisation is summarised in the following diagram, which contrasts with the unsuccessful tableau (3) above.

(14) *Optimality of first consonant deletion*

VC_1C_2V

① NoWeakCons ↙ ↖ ② Max

☞ VC_2V (③ **M** →) VC_1V

In diagrams such as this one, which will be used to summarise optimisations throughout the paper, each arrow points in a direction of increasing harmony according to the constraint with which the arrow is labelled. The number next to a constraint indicates its position in the hierarchy (i.e. as top-ranked or first in the hierarchy, ①; second in the hierarchy, ②, and so on). In general, only the harmonic orderings that are crucial for a given optimisation are shown in the corresponding diagram. Parentheses around an arrow indicate constraint violation. Here, the arrow labelled by **M** is parenthesised, because, as mentioned above, this constraint is overridden by the two higher-ranked constraints (NoWeakCons and Max).

A concrete example of this type of optimisation is summarised in the diagram below, which should be compared to tableau (4) above.

(15) *Optimality of first consonant deletion in Diola*

letkujaw

① NoWeakCons ↙ ↖ ② Max

☞ lekujaw (③ *Pl(lab,dor) →) letujaw

As in tableau (4), the candidate outputs shown in the diagram are possible realisations of the Diola input /let+ku+jaw/. But in contrast to the tableau, the correct output ([lekujaw]) is optimal here, despite the fact that the alternative candidate *[letujaw] better satisfies the place-markedness constraint *PL(lab,dor). This result, like the one summarised in the preceding diagram, is directly attributable to the targeting of NOWEAKCONS. Because the targeted contextual markedness constraint asserts only that [lekujaw] is more harmonic than [letkujaw], it allows the faithfulness constraint MAX to render [letujaw] non-optimal by placing it below [letkujaw] (and thus below [lekujaw]) in the harmonic ordering.

3 The typology of consonant deletion

This section is organised around two main themes. In the first part of the section (§§3.1 and 3.2), I focus on the typological *restrictiveness* of the targeted-constraint approach to consonant deletion introduced above. I begin by formalising the account of consonant deletion in Diola that was given informally in (15). I then show that the present approach, unlike the ones discussed in §1 (and in §4), captures the typological generalisation about consonant deletion in (1). In the second part of the section (§3.3), I focus on the typological *sufficiency* of the targeted-constraint approach. I demonstrate that, through its interaction with other constraints, NOWEAKCONS can account for non-deletion 'repairs' of illegal consonant clusters (e.g. vowel epenthesis). The second part of the section also supplements the first part by considering a larger set of candidate outputs.

3.1 Relative harmony and order-based optimisation

As discussed in §2.1, the introduction of targeted constraints such as NOWEAKCONS depends on a notion of RELATIVE HARMONY (which is in turn derived from a notion of RELATIVE MARKEDNESS). Given a candidate x, a targeted constraint does not assign a list of violations (or 'marks') to x. Instead, the constraint asserts that each member of a (possibly empty) set of alternative candidates is more harmonic than x. In other words, the constraint asserts a (possibly empty) set of pairwise harmonic orderings of the form 'candidate y is more harmonic than candidate x (i.e. $y \succ x$)'.

The standard OT definition of harmonic ordering by a constraint hierarchy (Prince & Smolensky 1993: ch. 5) is based on violations (or 'marks'), and therefore does not accommodate targeted constraints.[11] I

[11] As an indication of the inadequacy of violation-based harmonic ordering, consider that neither of the following two tableaux correctly expresses the harmonic orderings that NOWEAKCONS asserts.

(i)

let+ku+jaw	*not* NOWEAKCONS
a. letkujaw	*
b. lekujaw	
c. letujaw	

let+ku+jaw	*not* NOWEAKCONS
a. letkujaw	*
b. lekujaw	
c. letujaw	*

begin this subsection by presenting an alternative, *order-based* definition of harmonic ordering by a hierarchy. This definition is equivalent to the violation-based one when only untargeted constraints are at play, but it allows targeted constraints to participate in selecting the optimal candidate as well. The new definition is illustrated with the optimisation shown in the tableau below, which repeats part of (10) above.

(16) *Word-final voice neutralisation in Lithuanian*

daug	*[αvce]/V__#	*PRES[vce]
a. daug	*!	
☞ b. dauK		*
c. dauk	*!	*

The constraint violations recorded in this tableau, like those in any tableau, entail a set of pairwise harmonic orderings; in fact, the violations themselves have no significance apart from the harmonic orderings that can be derived from them. For example, *[α voice]/V __ # assigns one violation apiece to [daug] and [dauk], but assigns no violations to [dauK]. This is just one way of expressing the two harmonic orderings that the constraint asserts, namely: [dauK] ≻ [daug] and [dauK] ≻ [dauk]. All of the violations in the tableau can be 'unpacked' into harmonic orderings in this fashion, yielding a different representation of exactly the same optimisation:

(17) *Word-final voice neutralisation in Lithuanian*

daug	*[αvce]/V__#	*PRES[vce]
a. daug	dauK ≻ daug !	
☞ b. dauK		(daug ≻ dauK)
c. dauk	dauK ≻ dauk !	daug ≻ dauk
cumulative ordering	dauK ≻ daug, dauk	dauK ≻ daug ≻ dauk

In this alternative formalism, the course of the optimisation is tracked in the final row (labelled 'cumulative ordering'). The higher-ranked con-

In the tableau on the left, the fact that the constraint is violated only by [letkujaw] entails that it prefers [letujaw] over this candidate. In the tableau on the right, the fact that the constraint is violated by [letujaw] entails that it prefers [lekujaw] over this candidate. But neither the harmonic ordering [letujaw] ≻ [letkujaw] nor the harmonic ordering [letujaw] ≻ [letkujaw] is actually asserted by NoWeakCons.

The notion of relative harmony appealed to in the text also plays a role in recent papers by Samek-Lodovici & Prince (1999) and Prince (2000). But there is one quite crucial difference between the proposals: while Samek-Lodovici & Prince employ relative harmony to elucidate the interaction of constraints whose preferences are defined in terms of violation assignment, the central claim of this paper is that there are certain constraints – the targeted ones – whose preferences cannot be defined in the standard way, as illustrated in the tableaux in (i).

straint (*[α voice]/V __ #) establishes that candidate [dauK] is more harmonic than both [daug] and [dauk] (i.e. [dauK] $\succ$ [daug], [dauk]). The lower-ranked constraint (Preserve[voice]) *refines* this cumulative harmonic ordering, adding another ordering to it. Specifically, Preserve[voice] places the faithful candidate [daug] higher in the cumulative ordering than the unfaithful candidate [dauk]. Thus the final cumulative ordering – the one that incorporates the contributions of both constraints in the hierarchy – is [dauK] $\succ$ [daug] $\succ$ [dauk]. Candidate [dauK] is the only optimal candidate according to this final ordering, just as it is the only optimal candidate in the violation-based tableau (16), because it is the only one that is not less harmonic than any other candidate.

Of course, the faithfulness constraint Preserve[voice] would prefer the optimal candidate to be [daug]. But the harmonic ordering that would make this candidate optimal, namely [daug] $\succ$ [dauK], directly contradicts one of the harmonic orderings that is established by the higher-ranked markedness constraint, namely [dauK] $\succ$ [daug]. Consequently, violation of Preserve[voice] is compelled by *[α voice]/V __ #, exactly as in tableau (16). The order-based definition of harmonic ordering that resolves this and other cases of constraint conflict is given in (18); a more formal version appears in the Appendix.

(18) *Order-based optimisation by a constraint hierarchy*

a. Starting with the highest-ranked constraint and descending the hierarchy, if the current constraint asserts the harmonic ordering $x \succ y$, then add $x \succ y$ to the cumulative harmonic ordering (O), except when the opposite ordering (i.e. $y \succ x$) is in O.

b. A candidate is optimal iff it is not less harmonic than any other candidate according to the final cumulative harmonic ordering.

The only crucial aspect of order-based optimisation that does not appear in (18) is the following. For any candidates x, y and z, if both $x \succ y$ and $y \succ z$ are in the cumulative harmonic ordering, then $x \succ z$ is also automatically in the cumulative harmonic ordering (i.e. $x \succ y$ & $y \succ z \Rightarrow x \succ z$). This automatic deduction of additional harmonic orderings, which is referred to as transitive closure, applies after each constraint adds its (consistent) harmonic orderings to the cumulative ordering.

Although (18) is a necessary prerequisite for the targeted-constraint approach, I will seldom refer to it (let alone the more formal definition given in the Appendix) in the rest of the paper. Instead, I will trace through the other order-based tableaux below as I did the one in (17) above, and I will highlight the most important optimisations with summary diagrams like (15) in §2.2.

As noted above, the standard violation-based definition of harmonic ordering by a hierarchy and the order-based definition in (18) give equivalent results when only untargeted constraints are considered. But only the order-based definition is compatible with the crucial role of targeted constraints. The following tableau illustrates this with the familiar Diola example. (Henceforth, the symbol '⇒' is used to identify targeted constraints in tableaux. Every constraint is assumed to be untargeted in the absence of an argument to the contrary.)

(19) *First consonant deletion in Diola*

let+ku+jaw	⇒NoWeakCons	Max	*Pl(lab,dor)	*Pl(cor)
a. letkujaw	leku ≻ letku !		(letu ≻ letku)	leku ≻ letku
☞ b. lekujaw		(letku ≻ leku)	(letu ≻ leku)	
c. letujaw		letku ≻ letu !		leku ≻ letu
cumulative ordering	leku ≻ letku	leku ≻ letku ≻ letu		

The constraint hierarchy in this tableau maps the Diola input /let+ku+jaw/ to the grammatical output [lekujaw]. To understand how it does so, consider first the harmonic ordering that is asserted by the targeted constraint NoWeakCons. The fully faithful candidate [letkujaw] contains the poorly cued (or 'weak') consonant [t]. Therefore, in line with the weak element principle (12), NoWeakCons asserts only that [letkujaw] is less harmonic than the candidate that is identical except that [t] has been removed: namely, [lekujaw]. Thus the cumulative harmonic ordering below the highest-ranked constraint is just an ordering of these two candidates: [lekujaw] ≻ [letkujaw]. (To save space, I have omitted the final substring [jaw] from all of the harmonic orderings in the tableau.)

The next constraint, Max, asserts the opposite harmonic ordering of [lekujaw] and [letkujaw] (i.e. [letkujaw] ≻ [lekujaw]). But this constraint, like all constraints, cannot change harmonic orderings that have been fixed by higher-ranked constraints. What Max can do, however, is decide the relative ordering of [letkujaw] and [letujaw], which are left unordered by NoWeakCons. The constraint places the faithful candidate above the unfaithful one: [letkujaw] ≻ [letujaw]. The cumulative harmonic ordering below Max is formed by stringing together the two pairwise orderings just established: [lekujaw] ≻ [letkujaw] ≻ [letujaw]. This is a *total* ordering of the candidate set considered here, and at the top of the ordering lies the optimal candidate [lekujaw]. This candidate beats [letkujaw] straightforwardly: the highest-ranked constraint declares that [lekujaw] (and only [lekujaw]) is more harmonic than [letkujaw]. Notice, however, that [lekujaw] beats the other candidate, [letujaw], in a more interesting fashion: namely, by *transitivity* of harmonic ordering. Neither of the two highest-ranked constraints states that [lekujaw] is more harmonic than [letujaw]. But NoWeakCons asserts [lekujaw] ≻ [letkujaw], Max asserts [letkujaw] ≻ [letujaw], and these two orderings together yield, by transitivity, the ordering [lekujaw] ≻ [letujaw]. The possibility of establishing harmonic orderings purely by transitivity, of which this is a concrete and central example, is precisely the way in which targeted constraints enrich the standard OT conception of optimisation.

Whenever the higher-ranked constraints in a hierarchy establish that one candidate is more harmonic than all the others, the harmonic orderings asserted by the lower-ranked constraints are irrelevant. Thus, in this tableau, the fact that the place-markedness constraint *PL(lab,dor) prefers [letujaw] over both [lekujaw] and [letkujaw] has no effect on the selection of the optimal candidate. In other words, the present account successfully prevents this constraint from determining which of the two input consonants is deleted. Note that *PL(lab,dor) must indeed be ranked below MAX, which in turn is ranked below NOWEAKCONS – otherwise, labial and dorsal consonants would be incorrectly excluded from the segmental inventory of Diola. (The inactivity of *PL(lab,dor) and lowest-ranked *PL(cor) is indicated by shading the cumulative ordering cells below them.)

More generally, the hierarchy above accounts for the fact that it is consistently the *first* consonant of an illegal intervocalic cluster that is deleted in Diola (2a). And it does so, as just discussed, despite the existence of markedness constraints that sometimes favour deletion of the second consonant (e.g. *PL(lab,dor) prefers [letujaw] over [lekujaw] in the example above). Thus the targeted-constraint approach to consonant deletion solves the problem that was identified in §1, at least for this particular language.

In the following subsections, I demonstrate that the solution is in fact a much more general one. The targeted-constraint approach to consonant deletion accounts for the typological generalisation in (1) both with respect to restricted candidate sets of the type considered so far (see §3.2) and with respect to significantly larger candidate sets (see §3.3).

3.2 Basic factorial typology of consonant deletion

In OT, typological predictions are verified by computing a factorial typology: that is, by computing the input–output mappings that result from every possible ranking of the constraints of interest. This subsection begins the study of the typological predictions of the targeted-constraint approach to consonant deletion by computing a *basic* factorial typology. The typology is basic in two senses: (i) the only candidates considered are those that can be derived from the input by segment deletion or feature change; and (ii) the constraint set is restricted to NOWEAKCONS, the two faithfulness constraints MAX and IDENT, and non-contextual markedness constraints such as *PL(lab,dor). The extended factorial typology computed in the next subsection (§3.3) considers additional candidates derived by epenthesis and metathesis and the faithfulness constraints that those additional candidates violate.

The basic factorial typology observes the cross-linguistic generalisation about consonant deletion in (1). In different terms, the typology obeys the following implication, which relates consonant deletion processes to segmental inventories.

(20) *Inventory-restricted first consonant deletion*

Let α and β be any two consonants in the segmental inventory of language L. If L resolves intervocalic $\alpha\beta$ and $\beta\alpha$ clusters by deletion, then it does so by consistently deleting the *first* member of the cluster (i.e. /V$\alpha\beta$V/ $\rightarrow$ [VβV] and /V$\beta\alpha$V/ $\rightarrow$ [VαV]).

This implication is empirically equivalent to generalisation (1), because the data that supports the generalisation is data about the resolution of clusters containing consonants that actually exist in the inventory of a given language. The implication is more useful from the present theoretical perspective, however, because it evokes certain rankings that must hold in the hierarchies that are relevant for testing the analysis (see immediately below).

In light of the problem identified in §1, the most important aspect of (20) is that it makes no reference to the relative markedness (in language L or universally) of α and β. As long as α and β are both in the segmental inventory of L, the implication precludes a logically possible deletion process that removes the more marked consonant (say β) regardless of its position in the cluster (i.e. /V$\alpha\beta$V/ $\rightarrow$ [VαV] and /V$\beta\alpha$V/ $\rightarrow$ [VαV]). The actual selection of the deleted consonant is insulated from any (non-contextual) markedness constraint that prefers α over β.

To see why this implication holds, consider the rankings that are necessary to (a) place both α and β in the inventory and (b) force deletion in order to avoid intervocalic consonant clusters.

(21) *Rankings for consonant inventory and deletion*

a. *Inventory*: Max *and* Ident $\gg$ *β
b. *Deletion*: NoWeakCons $\gg$ Max

I have arbitrarily used β to represent the more marked consonant. The set of non-contextual markedness constraints that are violated by β, but not by α, are designated *β. (The corresponding set of constraints against α are lower-ranked, by hypothesis, and can therefore be safely ignored.) In order for β to be included in the inventory, the two faithfulness constraints Max and Ident must outrank *β, as in (21a). Otherwise, *every* instance of β in the input would be deleted or mapped to some less marked segment (e.g. α) in the output.[12] In order for deletion to be used to avoid intervocalic consonant clusters, the targeted contextual markedness constraint NoWeakCons must dominate Max, as in (21b). There are exactly

[12] To simplify the exposition, I consider only feature-changing mappings between the two abstract consonants α and β. Consideration of a full range of mappings between actual consonants would not affect the results about deletion presented here.

three hierarchies, listed immediately below, that are compatible with all of the rankings in (21).

(22) *Hierarchies compatible with consonant inventory* (21a) *and deletion* (21b)

a. NoWeakCons $\gg$ Max $\gg$ Ident $\gg$ $*\beta$
b. NoWeakCons $\gg$ Ident $\gg$ Max $\gg$ $*\beta$
c. Ident $\gg$ NoWeakCons $\gg$ Max $\gg$ $*\beta$

These three hierarchies differ only with respect to the position of Ident, which can be ranked anywhere above $*\beta$. All three of them map input /V$\alpha\beta$V/ to output [VβV] and input /V$\beta\alpha$V/ to [VαV], as desired. I have arbitrarily selected hierarchy (22b) to illustrate this; the following tableau shows how the hierarchy generates the crucial mapping /V$\alpha\beta$V/ $\rightarrow$ [VβV], in which the deleted first consonant is also the *less marked* one.

(23) *Deletion of the first* (*and less marked*) *consonant*

V$\alpha_1\beta_2$V	$\Rightarrow$NoWeakCons	Ident	Max	$*\beta$
a. V$\alpha_1\beta_2$V	$\beta_2 \succ \alpha_1\beta_2$!			$(\alpha_1,\alpha_2 \succ \alpha_1\beta_2)$ $(\alpha_1\alpha_2 \succ \alpha_1\beta_2)$
☞ b. Vβ_2V			$(\alpha_1\beta_2 \succ \beta_2)$ $(\alpha_1\alpha_2 \succ \beta_2)$	$(\alpha_1,\alpha_2 \succ \beta_2)$ $(\alpha_1\alpha_2 \succ \beta_2)$
c. Vα_1V			$\alpha_1\beta_2 \succ \alpha_1$! $(\alpha_1\alpha_2 \succ \alpha_1)$	
d. Vα_2V		$\alpha_1\beta_2 \succ \alpha_2$! $\beta_2,\alpha_1 \succ \alpha_2$!	$\alpha_1\beta_2 \succ \alpha_2$ $(\alpha_1\alpha_2 \succ \alpha_2)$	
e. V$\alpha_1\alpha_2$V	$\alpha_1,\alpha_2 \succ \alpha_1\alpha_2$!	$\alpha_1\beta_2 \succ \alpha_1\alpha_2$ $\beta_2,\alpha_1 \succ \alpha_1\alpha_2$		
cumulative ordering	$\beta_2 \succ \alpha_1\beta_2$ $\alpha_1,\alpha_2 \succ \alpha_1\alpha_2$	$\beta_2 \succ \alpha_1\beta_2 \succ \alpha_2 \succ \alpha_1\alpha_2$ $\alpha_1 \succ \alpha_2 \succ \alpha_1\alpha_2$	$\beta_2 \succ \alpha_1\beta_2 \succ \alpha_1 \succ \alpha_2 \succ \alpha_1\alpha_2$	

In this tableau, input–output correspondence relations are indicated by subscripts. For reasons of space, the harmonic orderings do not show the flanking vowels that appear in every candidate (e.g. the harmonic ordering [Vβ_1V] $\succ$ [V$\alpha_1\beta_2$V] is abbreviated by [β_1] $\succ$ [$\alpha_1\beta_2$]).

The most important point to notice about this tableau is that the three highest-ranked constraints together establish a *total* ordering of the candidate set, with the candidate [Vβ_2V] correctly at the top. This prevents lowest-ranked $*\beta$ from affecting the optimisation. The following remarks describe how the total ordering is constructed (see also the summary diagram in (24) below).

– NoWeakCons asserts that candidate [Vβ_2V] is more harmonic than candidate [V$\alpha_1\beta_2$V], and that candidates [Vα_1V] and [Vα_2V] are both more harmonic than candidate [V$\alpha_1\alpha_2$V]. These harmonic orderings follow directly from the definition of the targeted constraint: given any candidate containing a consonant that is not released into a vowel (i.e. that is 'weak'), the constraint prefers any alternative candidate that is identical except that the weak consonant has been removed. (Recall that these preferences are grounded in auditory/perceptual similarity. Therefore, input–output correspondence relations, which have no auditory/perceptual effect, are ignored. For purposes of evaluation by NoWeakCons, [Vα_1V] and [Vα_2V] are identical, and are both more harmonic than [V$\alpha_1\alpha_2$V].)

Because NoWeakCons is the highest-ranked constraint, the harmonic orderings that it asserts immediately enter the cumulative harmonic ordering. This narrows the field of potentially optimal candidates to those that do not contain a consonant cluster (namely [Vβ_2V], [Vα_1V] and [Vα_2V]).

– Ident further narrows the field by removing [Vα_2V], in which an input /β/ has been mapped to an output [α], from the set of potentially optimal candidates. In particular, Ident places both [Vβ_2V] and [Vα_1V] above [Vα_2V] in the cumulative harmonic ordering. Ident also places [Vβ_2V], [Vα_1V] and [V$\alpha_1\beta_2$V] above [V$\alpha_1\alpha_2$V], which has already been ruled non-optimal by NoWeakCons.

– The only remaining decision for the grammar to make is which one of [Vβ_2V] and [Vα_1V] is more harmonic – and Max decides in favour of [Vβ_2V] in an indirect fashion, establishing a harmonic ordering by transitivity. To see how this works, notice that [Vβ_2V] and [Vα_1V] do not stand on equal footing with respect to the two candidates that satisfy Max, namely [V$\alpha_1\beta_2$V] and [V$\alpha_1\alpha_2$V]. Candidate [Vβ_2V] lies above both [V$\alpha_1\beta_2$V] (by NoWeakCons) and [V$\alpha_1\alpha_2$V] (by Ident) in the cumulative ordering. Candidate [Vα_1V] also lies above [V$\alpha_1\alpha_2$V] (by NoWeakCons and Ident). But the relative ordering of [Vα_1V] and [V$\alpha_1\beta_2$V] has not yet been determined. Therefore, Max is able to assert its preference for candidates in which all of the input segments have output correspondents, placing [V$\alpha_1\beta_2$V] above [Vα_1V]. By transitivity, it follows from this ordering that [Vβ_2V] is more harmonic than [Vα_1V]. And thus [Vβ_2V] becomes optimal at this point, because it has been ruled more harmonic than *every* other candidate.

The following diagram summarises this optimisation.

(24) *Optimisation summary:* /V$\alpha_1\beta_2$V/ $\rightarrow$ [Vβ_2V]

$$
\begin{tikzcd}
V\alpha_1\beta_2V \arrow[d, "\text{① NoWeakCons}"'] & & V\alpha_1\alpha_2V \arrow[ll, "\text{② Ident}"'] \arrow[dl, "\text{① NoWeakCons}"] \arrow[dr] & \\
☞\ V\beta_2V \arrow[r, "\text{(④ *}\beta\text{)}"] & V\alpha_1V \arrow[ul, "\text{③ Max}"'] & & V\alpha_2V \arrow[ll, "\text{② Ident}"']
\end{tikzcd}
$$

Note in particular that the harmonic ordering [Vβ_2V] $\succ$ [Vα_1V], which is asserted by *β, contradicts a chain of harmonic orderings established by higher-ranked constraints: [Vβ_2V] $\succ$ [V$\alpha_1\beta_2$V] (by NoWeakCons) and [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] (by Max). Consequently, and as desired, the non-contextual markedness constraint *β cannot force deletion of the more marked consonant. (Recall that *β must be ranked below Max, which is in turn ranked below NoWeakCons, in order for β to be a member of the inventory at all.)

As just demonstrated, the hierarchy NoWeakCons $\gg$ Ident $\gg$ Max $\gg$ $*\beta$ maps input /$V\alpha_1\beta_2V$/ to the output [$V\beta_2V$], deleting the first consonant of the cluster in accordance with implication (20). The same hierarchy also maps input /$V\beta_1\alpha_2V$/ to output [$V\alpha_2V$], again deleting the first consonant (which in this case is the more marked one). The relevant optimisation is very similar to the one described above, therefore I provide only a summary diagram of it here.

(25) *Optimisation summary:* /$V\beta_1\alpha_2V$/ $\rightarrow$ [$V\alpha_2V$]

$$\begin{tikzcd}
V\alpha_1V \arrow[d, "\text{② }\textsc{Ident}"'] & V\alpha_1\alpha_2V \arrow[l, "\text{① }\textsc{NoWeakCons}"'] \arrow[d, "\text{② }\textsc{Ident}"'] & \\
\text{☞ } V\alpha_2V & V\beta_1\alpha_2V \arrow[l, "\text{① }\textsc{NoWeakCons}"'] & V\beta_1V \arrow[l, "\text{③ }\textsc{Max}"']
\end{tikzcd}$$

The non-contextual markedness constraint $*\beta$ is satisfied by the output in this case. However, as the diagram indicates, the outcome of the optimisation is fully determined by the three higher-ranked constraints. As in the previous optimisation, the decision about which consonant to delete is made on the basis of position in the cluster, not on the basis of intrinsic markedness.

Although there is not space to show this here, the other two hierarchies in (22) generate the same input–output mappings. Therefore, as claimed, the basic factorial typology obeys the implication in (20) (and hence also obeys the empirically equivalent generalisation in (1)). If consonants α and β are both in the segmental inventory of a language, then intervocalic clusters that contain them can be resolved by consistently deleting the first consonant, but *not* by consistently deleting the more marked consonant. In other words, the existence of the more marked consonant β in the segmental inventory forces the non-contextual markedness constraint $*\beta$ to be ranked so low that it cannot affect the decision about which consonant deletes. In the following subsection, I strengthen this result by showing that the implication also holds in an extended factorial typology.[13]

[13] At this point, it is important to emphasise that the targeted constraint NoWeakCons is being proposed as a *replacement* for the untargeted version that appears earlier in the paper. If both the targeted and the untargeted version of the constraint were allowed to coexist, the problem addressed in the paper would not be solved, because there would still be rankings under which deletion is determined by markedness, not by position. More generally, the claim being made is that all the contextual markedness constraints that can force consonant deletion – including NoCoda (Prince & Smolensky 1993) – must be replaced by targeted NoWeakCons or else reinterpreted as targeted themselves. To put the claim in a more positive way, the proposal is that all such constraints must be projected, in the sense described for the targeted NoWeakCons, from the weak element principle (12).

As the associate editor observes, this proposal naturally raises the question of whether *all* contextual markedness constraints (including, for example, those responsible for word-final devoicing and voicing assimilation in obstruent clusters)

3.3 Extended factorial typology

The previous subsection focused exclusively on the *restrictiveness* of the targeted-constraint approach to consonant cluster resolution. One of the goals of this subsection, as just mentioned, is to continue the demonstration of restrictiveness. Another equally important goal is to show that the approach is typologically *sufficient*: that it can account for other 'repairs' that languages employ to eliminate consonant clusters. Here, I illustrate the typological sufficiency of the approach with respect to vowel epenthesis.[14]

(26) *Partial typology of attested repairs*

Repair	*Faithfulness constraint violated by repair*	*Example language*
First consonant deletion ($VC_1C_2V \rightarrow VC_2V$)	Max	Diola-Fogny (see §§1, 3.1)
Vowel epenthesis ($VC_1C_2V \rightarrow VC_1VC_2V$)	Dep	Ponapean (see below)

The existence of epenthesis (and other non-deletion repairs) might appear to pose a problem for the present theory. Recall that the targeted constraint NoWeakCons only directly prefers deletion (i.e. it only asserts that candidates of the form [VC_2V] are more harmonic than candidates of the form [VC_1C_2V]). Therefore, it might appear that an entirely different type of contextual markedness constraint is needed to account for epenthesis, etc. Contrary to appearances, however, the observed typological variation in fact follows from the present theory in standard OT fashion. The targeted constraint NoWeakCons is sufficient to trigger

are targeted. Although no answer to such a broad question could be defended in a paper of this size, I believe that it is the case that all contextual markedness constraints, and perhaps all non-contextual markedness constraints as well, are in fact targeted; see §5.1 and especially Wilson (2000) for further discussion.

The associate editor also brings up the issue of whether the present proposal can accommodate a deletion pattern that was alluded to in the first section of the paper, namely one in which intervocalic obstruent+sonorant (O+R) and sonorant +obstruent (R+O) are consistently resolved by deletion of the sonorant. Although a fully worked-out account must be left for future research, the idea to be pursued is that a candidate of the form [VO_1R_2V] is equally perceptually similar to both [VR_2V] and [VO_1V], with the similarity of [VO_1R_2V] and [VO_1V] following from the fact that the transition from an obstruent to a vowel has many of the same acoustic properties as the transition from an obstruent to a sonorant (see Fleischhacker 2000 on obstruent+liquid clusters in particular). The targeted constraint against intervocalic clusters would then not distinguish between [VR_2V] and [VO_1V], and would therefore allow [VO_1V] to be selected as optimal by constraints that prefer less sonorous onsets (Clements 1990, Prince & Smolensky 1993, among others).

[14] As indicated, this typology is not intended to be exhaustive; due to space limitations, I cannot present analyses of other attested 'repairs' such as metathesis or feature change (but see the following note).

repairs other than deletion, and the particular repair that is selected by a given language is determined by the language-particular ranking of faithfulness constraints such as those listed in the centre column of (26).

For example, consider the epenthesis process that breaks up obstruent clusters in Ponapean (Austronesian).

(27) *Vowel epenthesis in Ponapean* (Itô 1986: §4.1; see also Rehg & Sohl 1981)[15]

a. /a**k** + **p**ʷuŋ/ a**kup**ʷuŋ 'petty'
b. /a**k** + **s**uwei/ a**kus**uwei 'demonstrating boorishness'

The hierarchy that generates this epenthesis process has NoWeakCons at the top, Dep at the bottom and all the other faithfulness constraints in the middle. The optimality of epenthesis according to this hierarchy is shown in the following tableau, which for reasons of space includes only a subset of the relevant constraints and candidates (adding the others would not affect the result).

(28) *Optimality of vowel epenthesis in Ponapean*

ak+pʷuŋ	⇒NoWeakCons	Max	Dep
a. akpʷuŋ	apʷuŋ ≻ akpʷuŋ !		
b. apʷuŋ		(akpʷuŋ ≻ apʷuŋ) akupʷuŋ ≻ apʷuŋ !	
c. akuŋ		akpʷuŋ ≻ akuŋ ! akupʷuŋ ≻ akuŋ !	
☞ d. akupʷuŋ			(akpʷuŋ ≻ akupʷuŋ) (apʷuŋ ≻ akupʷuŋ) (akuŋ ≻ akupʷuŋ)
cumulative ordering	apʷuŋ ≻ akpʷuŋ	akupʷuŋ ≻ apʷuŋ ≻ akpʷuŋ ≻ akuŋ	

In this tableau, there are two candidates that satisfy Max: [akpʷuŋ] (the fully faithful candidate, in which the two obstruents are adjacent) and [akupʷuŋ] (an unfaithful candidate, in which an epenthetic [u] separates the two obstruents). The fully faithful candidate [akpʷuŋ] is rendered non-optimal by the highest-ranked constraint: as always, NoWeakCons prefers the minimally different candidate from which the first member of the intervocalic cluster has been removed ([apʷuŋ]). But the epenthesis candidate [akupʷuŋ] is placed above all of the other candidates by Max,

[15] I do not account for the quality of the epenthetic vowel here. Note also that Ponapean has a process of 'nasal substitution' that applies to clusters that are underlyingly homorganic under certain circumstances, as illustrated in /ak + keelail/ → [aŋkeelail] 'demonstrate strength' (Itô 1986: 137). Although an analysis of this nasal substitution process cannot be given here, note that the targeted-constraint approach is in general able to account for feature-changing 'repairs' by ranking Ident constraints (e.g. Ident(son) and Ident(nas)) below the other faithfulness constraints.

and is therefore optimal. Lowest-ranked DEP prefers every other candidate over [akup^{w}uŋ], but these preferences are overridden by the dominant constraints.[16]

This optimisation illustrates the following important point: the effects of a targeted constraint are *not* limited to guaranteeing the optimality of the candidates that it directly prefers. The only candidate that is directly preferred by NOWEAKCONS in the tableau above is [ap^{w}uŋ]. This candidate is not the optimal one. But the fact that NOWEAKCONS places [ap^{w}uŋ] above [akp^{w}uŋ] in the cumulative harmonic ordering is nevertheless crucial for establishing the optimality of [akup^{w}uŋ]. As shown in the following diagram, [akup^{w}uŋ] beats the fully faithful candidate [akp^{w}uŋ] only *by transitivity*: [akup^{w}uŋ] ≻ [ap^{w}uŋ] and [ap^{w}uŋ] ≻ [akp^{w}uŋ] implies [akup^{w}uŋ] ≻ [akp^{w}uŋ].

(29) *Optimisation summary: vowel epenthesis*

$$\begin{tikzcd} \text{ap}^{w}\text{uŋ} \arrow[d, "\text{② Max}"'] & \text{akp}^{w}\text{uŋ} \arrow[l, "\text{① NoWeakCons}"'] \\ \text{☞ akup}^{w}\text{uŋ} & \text{akuŋ} \arrow[l, "\text{② Max}"] \arrow[u, "\text{② Max}"'] \end{tikzcd}$$

As this case demonstrates, a targeted constraint can have widespread effects through its interaction with other constraints. This is arguably the most desirable property of *any* OT constraint. And it clearly distinguishes targeted constraints from the rules that they superficially resemble. For example, the formal statement of NOWEAKCONS (13) looks much like the following deletion rule: $C \rightarrow \emptyset / \_\_ C$. But this rule, unlike the constraint, could never play a crucial role in accounting for vowel epenthesis. Thus the main argument for OT – that it allows us to capture the common surface motivation for a variety of descriptive repairs (e.g. deletion *and* epenthesis) – is not jeopardised by the switch to targeted contextual markedness constraints.

The remainder of this subsection is divided into two parts. In §3.3.1, I present the extended factorial typology that results when root-faithfulness constraints are included in the constraint set. In §3.3.2, I introduce a principle that eliminates certain circular harmonic orderings that would otherwise be created, under the order-based definition of optimisation in (18), when candidates that have undergone metathesis are considered.

3.3.1 *Root dominance in consonant deletion.* Many researchers have argued that certain morphological categories – in particular *roots* – are subject to special faithfulness requirements (Beckman 1998, Casali 1997, McCarthy & Prince 1995, Zoll 1998). The implementation of this idea that I consider here is the one proposed by Beckman (1998): a set of general faithfulness constraints (FAITH) apply to all morphological

[16] In the optimisations prior to this one, it was implicitly assumed that DEP dominated MAX, and therefore that vowel epenthesis was less harmonic than consonant deletion.

categories, roots included; in addition, a special set of faithfulness constraints (FAITH$_{\text{Root}}$) apply to roots only.

When added to the constraint set of the basic factorial typology, root-faithfulness constraints expand the predicted typology in a very restricted way. Without these constraints, it is impossible to generate a pattern in which the second member of an intervocalic cluster is deleted ([VC_1C_2V] → [VC_1V]). With the addition of root-faithfulness constraints, it becomes possible to generate this pattern in one particular environment: namely, root+suffix junctures. In this environment, deletion of the second (i.e. suffixal) consonant allows the first (i.e. root) consonant to be preserved, thus satisfying MAX$_{\text{Root}}$. The following tableau displays the predicted root-dominance pattern with abstract candidates of the type used in the basic factorial typology.

(30) *Root-dominated consonant deletion*

Vβ+αV (*root+suffix*)	⇒NOWEAK CONS	MAX$_{\text{Rt}}$	MAX	*β
a. V$\beta\alpha$V	VαV ≻ V$\beta\alpha$V !			(VαV ≻ V$\beta\alpha$V)
b. VαV		(V$\beta\alpha$V ≻ VαV) VβV ≻ VαV !	(V$\beta\alpha$V ≻ VαV)	
☞ c. VβV			(V$\beta\alpha$V ≻ VβV)	(VαV ≻ VβV)
cumulative ordering	VαV ≻ V$\beta\alpha$V	VβV ≻ VαV ≻ V$\beta\alpha$V		

MAX$_{\text{Root}}$
If ζ is a root segment in the input, then ζ must have a correspondent in the input.

As far as I can determine, the predicted root-dominated pattern of consonant deletion occurs with the negative suffix of the language Ibibio (Benue Congo; Akinlabi & Urea 1993, Beckman 1998). Even if that particular case does not instantiate the pattern, however, it is likely that an example will be found in a broader typological survey; the same pattern of preservation of root material at the expense of affixal material is the norm in vowel harmony, and is also attested – though apparently less frequently – for voice assimilation (e.g. in Tangale; Kidda 1995), place assimilation (e.g. in Gidabal; Anderson 1974, Geytenbeek & Geytenbeek 1971) and vowel deletion (in various languages; Casali 1997).

At other morphological junctures (e.g. prefix+root, affix+affix and root+root) and morpheme-internally, the prediction of the theory remains that only the first member of an illegal consonant cluster is susceptible to deletion (absent any interfering inventory restrictions). This prediction is as far as I know correct. Most importantly, the introduction of root-faithfulness constraints does not affect the central result of the basic factorial typology: namely, that a language cannot decide which consonant to delete on a cluster-by-cluster basis, dropping the more marked consonant in each case. (Note also that the addition of root-faithfulness constraints does not nullify the prediction that consonant clusters arising at root + suffix junctures are resolved by deletion of the first/root consonant in some languages. This pattern, which is observed for example in West Greenlandic (2a), is gotten by simply re-ranking $\textsc{Max}_{\text{Root}}$ below Max in the hierarchy above.)

3.3.2 *Metathesis candidates and priority of the more harmonic.* The basic factorial typology of § 3.2 considered only candidates that are derived from the input by segment deletion and/or feature change. Earlier in this subsection (§ 3.3), the candidate set was expanded to include candidates derived by vowel epenthesis (see tableau (28) above). Complete candidate sets are standardly assumed to include candidates derived by metathesis as well. In the Correspondence Theory of faithfulness, a metathesis candidate is one in which either (i) output segments are reordered relative to their input correspondents (e.g. $/VC_1C_2V/ \rightarrow [VC_2C_1V]$) or (ii) output segments affiliated with different morphemes are intercalated (e.g. $/VC_1+C_2V/ \rightarrow [VC_2C_1V]$). (I assume that the exponents of different morphemes are not ordered in the input. Therefore, mappings such as $/VC_1+C_2V/ \rightarrow [VC_2C_1V]$ violate only Contiguity, while morpheme-internal metathesis such as $/VC_1C_2V/ \rightarrow [VC_2C_1V]$ violates both Contiguity and Linearity. The distinction is not relevant for the following discussion, therefore I henceforth use NoMetathesis to refer to both Contiguity and Linearity.)

The inclusion of metathesis candidates in the candidate set creates a potential threat to the restrictive predictions of the basic factorial typology. To see the potential problem, consider again the abstract input $/V\alpha_1\beta_2V/$. As in the basic factorial typology, we hypothesise that β is more marked than α according to certain non-contextual markedness constraints, which are represented by $*\beta$. The following diagram indicates how this markedness relation could potentially force deletion of β – that is, deletion of the *second* consonant – when the metathesis candidate $[V\beta_2\alpha_1V]$ is allowed to compete against $[V\alpha_1\beta_2V]$, $[V\beta_2V]$ and $[V\alpha_1V]$. The relevant hierarchy is NoWeakCons ≫ Max ≫ $*\beta$ ≫ NoMetathesis.

(31) *Incorrect deletion of the second consonant* (*but see below*)

$$\begin{tikzcd}
& \text{☞ } V\alpha_1V & & V\beta_2\alpha_1V \arrow[ll, "①\ \textsc{NoWeakCons}"'] \arrow[dl, "④\ \textsc{NoMetathesis}"] \\
V\beta_2V \arrow[ur, "③\ {*}\beta"] & & V\alpha_1\beta_2V \arrow[ll, "①\ \textsc{NoWeakCons}"'] &
\end{tikzcd}$$

Notice, however, that this diagram is incomplete in one crucial respect: it does not take into account the harmonic orderings asserted by Max, which must be ranked higher than $*\beta$ in order for β to be included in the segmental inventory. Max asserts that the fully faithful candidate $[V\alpha_1\beta_2V]$ and the metathesis candidate $[V\beta_2\alpha_1V]$, which both contain an output correspondent for each of the input consonants, are more harmonic than both [Vβ_2V] and [Vα_1V]. Two of these harmonic orderings (namely, [V$\alpha_1\beta_2$V] $\succ$ [Vβ_2V] and [V$\beta_2\alpha_1$V] $\succ$ [Vα_1V]) contradict higher-ranked NoWeakCons, and are therefore cancelled out immediately. The two remaining orderings asserted by Max are [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] and [V$\beta_2\alpha_1$V] $\succ$ [Vβ_2V]. The first of these asserts that the fully faithful candidate ([V$\alpha_1\beta_2$V]) is more harmonic than the candidate from which the second consonant has been deleted ([Vα_1V]). If this ordering is added to the diagram, then the optimal candidate becomes – as desired – the one from which the first consonant has been deleted (i.e. [Vβ_2V]). As in the basic factorial typology, [Vβ_2V] beats its competitor [Vα_1V] by transitivity: [Vβ_2V] $\succ$ [V$\alpha_1\beta_2$V] and [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] together imply [Vβ_2V] $\succ$ [Vα_1V].

(32) *Correct deletion of the first consonant*

$$\begin{tikzcd}
 & \mathrm{V}\alpha_1\mathrm{V} \arrow[dr, "② \textsc{Max}"] & & \mathrm{V}\beta_2\alpha_1\mathrm{V} \arrow[ll, "① \textsc{NoWeakCons}"'] \arrow[dl, "④ \textsc{NoMetathesis}"] \\
☞\ \mathrm{V}\beta_2\mathrm{V} \arrow[ur, "(③ *\beta)"] & & \mathrm{V}\alpha_1\beta_2\mathrm{V} \arrow[ll, "① \textsc{NoWeakCons}"'] &
\end{tikzcd}$$

The question to ask now is what happens when the final harmonic ordering that is asserted by Max, namely [V$\beta_2\alpha_1$V] $\succ$ [Vβ_2V], is also added to the diagram in (32). The answer is that adding *both* [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] *and* [V$\beta_2\alpha_1$V] $\succ$ [Vβ_2V] gives rise to a circular harmonic ordering, as shown in the diagram below. The circularity is established by the two highest-ranked constraints, NoWeakCons and Max, therefore the harmonic orderings asserted by the other constraints are irrelevant.

(33) *A circular harmonic ordering*

$$\begin{tikzcd}
\mathrm{V}\beta_2\alpha_1\mathrm{V} \arrow[r, "① \textsc{NoWeakCons}"] & \mathrm{V}\alpha_1\mathrm{V} \arrow[d, "② \textsc{Max}"] \\
\mathrm{V}\beta_2\mathrm{V} \arrow[u, "② \textsc{Max}"] & \mathrm{V}\alpha_1\beta_2\mathrm{V} \arrow[l, "① \textsc{NoWeakCons}"']
\end{tikzcd}$$

Circular harmonic orderings are illegitimate, I assume, for the simple reason that they do not contain any optimal candidate. (Recall that a candidate *x* is optimal according to the order-based definition of optimality in (18) iff there is no candidate *y* that is more harmonic than *x*. Note also that the term 'circular ordering' is a convenient oxymoron; 'circular relation', or better yet 'non-asymmetric relation', would be more precise.) Thus, diagram (33) tells us that *either* [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] *or* [V$\beta_2\alpha_1$V] $\succ$ [Vβ_2V], but not *both*, could take part in determining the cumulative harmonic ordering. The remaining question is why it is actually [V$\alpha_1\beta_2$V] $\succ$ [Vα_1V] that does so, as shown in (32).

I propose that there is a general principle that governs cases of potential circularity such as this one. Suppose that a given constraint asserts two pairwise harmonic orderings that are *individually* consistent with the harmonic orderings that have been established by higher-ranked constraints. Suppose further that, as in this case, adding *both* of these pairwise orderings to the cumulative harmonic ordering leads to circularity. Then the grammar gives priority to the pairwise ordering that favours the *more harmonic* candidate. More precisely, priority is determined as follows:

(34) *Priority of the more harmonic*

Let **C** be any constraint, and let $f \succ x$ and $g \succ y$ be any two harmonic orderings asserted by **C**. If f is more harmonic than g according to the hierarchy (i.e. if the highest-ranked constraint that prefers one of the candidates over the other favours f), then $f \succ x$ takes priority over $g \succ y$ when adding both of them to the cumulative harmonic ordering would lead to circularity.[17]

In the case under consideration, the priority of the more harmonic principle assigns priority to $[V\alpha_1\beta_2V] \succ [V\alpha_1V]$ over $[V\beta_2\alpha_1V] \succ [V\beta_2V]$, because the highest-ranked constraint that prefers one of the two candidates $[V\alpha_1\beta_2V]$ and $[V\beta_2\alpha_1V]$ over the other is NoMetathesis (the two candidates tie on all the other constraints), and NoMetathesis prefers $[V\alpha_1\beta_2V]$ over $[V\beta_2\alpha_1V]$. Consequently, only $[V\alpha_1\beta_2V] \succ [V\alpha_1V]$ – the pairwise ordering with higher priority – is added to the cumulative harmonic ordering. It follows that the optimal candidate is in fact the correct one ($[V\beta_2V]$), as shown in diagram (32). More generally, the main result of the basic factorial typology is maintained even now that metathesis candidates are included in the candidate set. Non-contextual markedness constraints such as $*\beta$ cannot force deletion of the second member of an illegal intervocalic cluster (inventory restrictions aside); deletion consistently eliminates the first member of such a cluster, in accordance with the generalisation in (1).[18]

The proposed analysis of consonant deletion has now been fully presented. The central claim of the analysis is that the contextual markedness constraints responsible for deletion must be *targeted*

[17] The priority of the more harmonic principle is clearly related to, and was no doubt partially inspired by, Sympathy Theory (McCarthy 1999, to appear). See Wilson (2000) for a formal discussion of the connection between the two proposals.

[18] In addition to all of the candidates considered so far, candidate sets are standardly assumed to include representations that are derived from the input by fusion (e.g. $/VC_1C_2V/ \rightarrow [VC_{1,2}V]$). I tentatively suggest that such candidates, which involve many-to-one correspondence relations, should in fact not be produced by Gen (except perhaps in the case of fusion of tautomorphemic identical segments, the province of the classical OCP (McCarthy 1986)). The purported empirical motivation for fusion comes from cases in which, descriptively speaking, one segment assimilates in some feature(s) to another segment that is nevertheless not present in the output. However, unless all assimilation should be analysed with fusion, as seems highly unlikely, such cases appear to call instead for a theory of phonological opacity. (The connection between fusion and opacity is mentioned in McCarthy 1998: 45.) Approaches to opacity in OT are currently the subject of intense research; indeed, the targeted-constraint approach developed here also offers an account of at least certain opacity effects, as I demonstrate in §5.2 below.

constraints. Specifically, these constraints must prefer only deletion of the first (i.e. poorly cued) member of an intervocalic biconsonantal cluster. On the substantive side, this preference is grounded in the weak element principle (12), which states that the less complex of two auditorily/perceptually similar representations is less marked. On the formal side, the introduction of targeted constraints into the OT framework is supported by an order-based definition of optimisation ((18); see also the Appendix).

The targeted-constraint analysis captures the generalisation about consonant deletion in (1). Most importantly, it rules out deletion processes that violate the generalisation by consistently removing the less marked consonant *regardless* of its position in the cluster. This result follows immediately when the candidate set is restricted to forms derived from the input by deletion and feature change. And the result holds when this restriction is dropped, given one additional principle: priority of the more harmonic (34).

4 Discussion of alternative analyses

In this section, I consider conceivable alternatives to the targeted-constraint approach presented above. The alternatives come in two varieties, reflecting the major dichotomy between markedness and faithfulness: alternatives based on contextual markedness constraints are considered in §4.1; alternatives based on contextual faithfulness constraints are considered in §4.2.

4.1 Untargeted contextual markedness

Because the targeted constraints proposed here (e.g. NoWeakCons) are contextual markedness constraints of a certain kind, the first type of alternative approach is limited to *untargeted* contextual markedness constraints. This type of alternative was already shown not to be viable in §1 above. The key point to recall is that any untargeted contextual markedness constraint **C** – whether formulated in terms of perceptual cues or prosodic positions – can be satisfied by deletion of *either* member of an intervocalic biconsonantal cluster. Whichever consonant survives in the output can be 'strong' and/or 'licensed' (i.e. by being released into a vowel and/or being in the onset of a syllable) and therefore does not force a violation of **C**. Consequently, **C** always lets other constraints (even very low-ranked ones) decide which consonant is deleted. The decision can even be made by some non-contextual markedness constraint that is ranked below both Max and Ident, and which therefore does not restrict the inventory of the language in question. Non-contextual markedness constraints by their very nature do not consistently favour the second member of a cluster over the first. Therefore, alternatives based on untargeted contextual markedness constraints cannot capture the typological generalisation in (1) (or the empirically equivalent implication in (20).

4.2 Contextual faithfulness

The other type of potential alternative approach makes use of contextual faithfulness constraints (also known as 'positional' faithfulness constraints). Many researchers have argued that the theory of contextual neutralisation must include such constraints, either in addition to or instead of contextual markedness constraints (Alderete 1999, Beckman 1998, Jun 1995, Kirchner 1998, Lombardi 1997, 1999, Padgett 1995; compare Zoll 1998, which presents arguments for the necessity of contextual markedness).

The general idea behind contextual faithfulness constraints is that 'licensing' contexts impose more stringent faithfulness requirements. For example, the pattern of voice neutralisation in Lithuanian (see Steriade 1997), which was (partially) analysed with the contextual markedness constraint *[α voice]/V __ # in §2.1 above, can be alternatively analysed with a contextual faithfulness constraint such as IDENT(voice)/__ [+son] ('an output segment that appears before a sonorant must be faithful to its input correspondent's voice specification'). The ranking IDENT(voice) /__ [+son] ≫ *[α voice] preserves contrastive voice on presonorant obstruents (see the following tableau), while the ranking *[α voice] ≫ IDENT(voice) ensures that the obstruent-voice contrast is neutralised elsewhere, as should be clear even without a tableau. (See Beckman 1998 for general discussion of hierarchies like this one, in which a contextual faithfulness constraint dominates a markedness constraint that in turn dominates a non-contextual faithfulness constraint.)

(35) *Contextual faithfulness preserves voice before a sonorant in Lithuanian*

auglingas	ID[vce]/ __[+son]	*[α vce]	ID[vce]
☞ a. auglingas		*	
b. auKlingas	*!		*
c. auklingas	*!	*	*

It is easy to see that contextual faithfulness constraints suffer from exactly the same inadequacy as untargeted contextual markedness constraints with respect to the generalisation in (1). Let IDENT(strong) be any IDENT constraint that applies to output consonants that are 'strong' (i.e. in a context that proves them with robust perceptual cues) and/or 'licensed' (say, in the onset of a syllable). As shown in the following tableau, which again uses the Diola input /let+ku+jaw/ for purposes of illustration, IDENT(strong) can be satisfied *regardless* of which of the relevant consonants is deleted. (Indeed, IDENT(strong) can be satisfied even if *neither* consonant is deleted, as long as the strong/licensed consonant in the output is featurally faithful to its input correspondent. Deletion must be forced by some other constraint, which for convenience I have taken to be *[−son]/ __ C.)

(36) *Incorrect deletion of non-coronal* C *in Diola*

let+ku+jaw	Id(strong)	*[−son]/__C	Max	*Pl(lab,dor)
a. letkujaw		*!		*
☞ b. lekujaw			*	*!
☹ c. letujaw			*	

The strong/licensed obstruent in both [lekujaw] and [letujaw] is featurally faithful to its input correspondent, therefore Ident(strong) does not distinguish between the two candidates. The untargeted contextual markedness constraint *[−son]/ __ C and Max do not distinguish between the two candidates either. Therefore, the decision about which consonant to delete is made – incorrectly in this case and in general – by a lower-ranked non-contextual markedness constraint (here, *Pl(lab,dor)).

Most of the contextual faithfulness constraints proposed in the literature cited above are versions of Ident(strong), and are thus inadequate given the preceding discussion. But there are two other logically possible types of contextual faithfulness constraint to consider. One, which I will refer to as Max(strong), would require every strong and/or licensed input segment to have an output correspondent (see §4.2.1). The second, which I will refer to as Faith(CV), would require faithfulness to CV sequences (see §4.2.2).

4.2.1 *Faithfulness to strong input consonants*. Upon initial inspection, Max(strong) may appear to be sufficient to account for consonant deletion processes such as the one in Diola. Consider that, in the input /let +ku+jaw/, only one of the relevant consonants (/k/) is prevocalic. If Max(strong) states that every prevocalic consonant in the input must have an output correspondent, then ranking it anywhere above *Pl(lab,dor) will guarantee that the correct output, [lekujaw], is more harmonic than the incorrect output *[letujaw].

This approach, although apparently straightforward, both (i) requires additional assumptions in order to account for simple cases like deletion in Diola and moreover (ii) fails to account for slightly more complex cases. Notice first that the approach cannot account for the generalisation about consonant deletion in (1), even if only simple cases like Diola are considered, without a universally fixed ranking. In particular, Max(strong) must be assumed to universally outrank Max in order for implication (20) to hold in the predicted typology (i.e. in order for it to be the case that, absent any inventory restrictions, only the first member of an intervocalic cluster is susceptible to deletion).

Even granted this universal ranking, the approach either foregoes the insights of the licensing-by-cue theory of contextual neutralisation (see §2.1) or requires an additional assumption about the set of universal inputs. Recall from the discussion of the licensing-by-cue theory that a prevocalic consonant is perceptually 'strong' only by virtue of being released into the following vowel. The provisional definitions of MAX(strong) that was assumed above did not incorporate the requirement that a consonant must be released (in addition to being prevocalic) in order to qualify as 'strong'. Without the release requirement, the constraint relies on an unmotivated notion of 'strength'. But if the constraint is redefined so that it applies only to input consonants that are prevocalic *and* released, then the generalisation in (1) cannot be captured without the additional stipulation that every prevocalic input consonant is released. The following tableau illustrates the problem that would exist in the absence of this stipulation about inputs. (In this tableau, the diacritic '˺' marks lack of release.)

(37) *Incorrect deletion of non-coronal* C *in Diola* (*hypothetical input*)

let˺+k˺u+jaw	MAX(strong)	*[−son]/__C	MAX	*PL(lab,dor)
a. let˺kujaw		*!		*
☞ b. lekujaw			*	*!
☹ c. letujaw			*	

The two obstruents in the input for this tableau are both unreleased, therefore deletion of either one vacuously satisfies the redefined version of MAX(strong). Consequently, the decision about which consonant is deleted incorrectly passes to *PL(lab,dor). As already mentioned, this problem could in principle be solved by stipulating that every prevocalic input consonant is released (thus ruling out inputs like the one above). However, this universal restriction on *input* consonants would be redundant with (i.e. would duplicate) whatever set of assumptions ensures that *output* consonants are systematically released in prevocalic position. Note that the input restriction is not sufficient by itself, since whether or not an output consonant is released cannot in general be predicted from any property of the input. (Note also that redefining 'strength' in terms of the onset position would not circumvent the duplication problem: a stipulation to the effect that every prevocalic input consonant is in the onset would, of course, also be redundant with independently necessary output constraints.)

In contrast to the MAX(strong) approach, the targeted-constraint approach does not depend on any new universal rankings, as demonstrated in §3. And it incorporates the insights of the licensing-by-cue framework without placing new restrictions on inputs, for the simple reason that it does not use an input property of a given consonant to predict whether or not that consonant is deleted to avoid an output cluster. In addition to these theoretical advantages, the targeted-constraint approach also has an empirical advantage with respect to cases in which (descriptively speaking) vowel deletion creates an intervocalic cluster that is then resolved by consonant deletion. Two cases of this sort were cited in (5) of §1, which is repeated below.[19]

(38) a. *Carib*

/s+enaa**pɨ**+**s**a/	senaa**s**a	'I eat it'
/s+enee**pɨ**+**s**a/	senee**s**a	'I bring it'

b. *Tunica*

/tiˈtih**ki**+**t**ʔɛ/	tiˈtih**t**ʔɛ	'a river'
/tiˈtih**ki** **p**iˈrʔutakʔahča/	tiˈtih**p**iˈrʔutakʔahča	'it will turn into a bayou'

The Carib examples in (38a) conform exactly to the generalisation about consonant deletion in (1): the intervocalic cluster that would be created by vowel deletion (e.g. /s+enaapɨ+sa/ → *[senaapsa]) is resolved by deletion of the *first* consonant (*[senaapsa] → [senaasa]). The Tunica examples in (38b) conform to a slightly extended version of generalisation (1), and are equally relevant to the present discussion. Vowel deletion as in /tiˈtihki piˈrʔutakʔahča/ → *[tiˈtihkpiˈrʔutakʔahča] would create a *triconsonantal* cluster [hkp], which is avoided by deletion of the *medial* consonant (*[tiˈtihkpiˈrʔutakʔahča] → [tiˈtihpiˈrʔutakʔahča]). The extension of generalisation (1) that applies to triconsonantal clusters as well as biconsonantal clusters is this: deletion processes that affect intervocalic clusters do not delete the *final* (i.e. prevocalic) member of the cluster. As far as I am aware, the extended generalisation about consonant deletion holds cross-linguistically (with the same sonority-based and root-dominance qualifications discussed in §1 and §3.3.1, respectively).[20] Crucially for the present discussion, the extended generalisation holds even when the cluster is created by vowel deletion, as in (38) above and similar cases found in the Billiri dialect of Tangale (Kidda 1995, Charette 1990: 247) and in Erromangan (also known as Sye (Southern Vanuatu); Crowley

[19] I do not take the motivation for vowel deletion to be directly relevant for the present discussion, which is concerned only with the outcome of consonant deletion. The following descriptions will, however, be relevant for a more complete analysis of these two cases. In Carib, vowel deletion is found with 'the vast majority of verbs ending in *pï*, *tï*, *kï*, *rï*, *mï*, *ku*, or *ru*, or with the verb *wï:to* 'to go'' (Hoff 1968: 59, also cited in Gildea 1995: 67). In Tunica, 'stems ending in hki, ši, ni, li, or ri may syncopate the i (except when their penult is stressed) when they come to stand before a grammatical element beginning in a consonant' (Haas 1946: 343) and 'words ending in hki, hku, ši, ni, li, or ri (unless they have a stressed penult) usually lose the i or u when followed by another word in the same phrase' (Haas 1946: 345).

[20] See Steriade (2000) for a review of cases of deletion in triconsonantal clusters that supports the generalisation in the text. Although I lack space in the present paper to address this issue in detail, the targeted constraint approach to deletion in intervocalic biconsonantal clusters extends straightforwardly to the examples of deletion in triconsonantal clusters cited by Steriade (see Wilson 2000), and possibly also to deletion in word-initial and word-final clusters as well.

1998: 28). The MAX(strong) approach fails in such cases, as is illustrated in the following tableau.

(39) *Incorrect optimality of final consonant deletion in Tunica*

tiˈtihki piˈrʔutakʔahča	*[−son]/__C	MAX(strong)	MAX	*PL(lab,dor)
a. tiˈtihkpiˈrʔutakʔahča	*!			**
b. tiˈtikpiˈrʔutakʔahča	*!		*	**
☞ c. tiˈtihpiˈrʔutakʔahča		*	*	*
👎 d. tiˈtihkiˈrʔutakʔahča		*	*	*

For expository purposes, I have suppressed the constraint interaction that is responsible for vowel deletion (see note 19). As shown, MAX(strong) is violated *both* by deletion of the final member of the derived cluster ([p]) *and* by deletion of the second member ([k]); this is true under both of the definitions of input 'strength' that were considered above. Therefore, in this type of case the contextual faithfulness constraint MAX(strong), just like the general faithfulness constraint MAX, incorrectly allows lower-ranked constraints to determine which consonant actually deletes. In this particular case, lower-ranked *PL(lab,dor) is satisfied equally by deletion of either consonant, leading to the incorrect prediction that deletion of the final consonant of the cluster ([p]) is a grammatical option in Tunica.

Unlike the approach based on MAX(strong), the targeted-constraint approach carries over straightforwardly to cases like those in (38). The relevant optimisation is very similar to previous ones, and is shown in (40) and (41), using a hypothetical input (for reasons of space only). Note that the constraint 'Syncope' stands in for whatever actual constraint drives vowel deletion in the relevant interconsonantal context (again, see note 19).

(40) *Vowel deletion feeds first consonant deletion* (*hypothetical example*)

leti+ku	'Syncope'	⇒NOWEAKCONS	MAX	*PL(lab,dor)
a. letiku	letku, leku, letu ≻ letiku !			letu ≻ letiku
b. letku		leku ≻ letku !	(letiku ≻ letku)	(letu ≻ letku)
☞ c. leku			(letiku, letku ≻ leku)	(letu ≻ leku)
d. letu			letku ≻ letu ! (letiku ≻ letu)	
cumulative ordering	letku, leku, letu ≻ letiku	leku ≻ letku ≻ letiku letu ≻ letiku	leku ≻ letku ≻ letu ≻ letiku	

(41) letku ← ① 'Syncope' letiku

② NOWEAKCONS ↙ ↖ ③ MAX

☞ leku (④ *PL(lab,dor) →) letu

In this order-based tableau, highest-ranked 'Syncope' asserts that the candidate containing the medial vowel ([letiku]) is less harmonic than all of the other candidates. The next constraint in the hierarchy, NoWeakCons, asserts only that the candidate containing the consonant cluster ([letku]) is less harmonic than the candidate that is exactly like it except that the weak consonant [t] has been removed (i.e. [leku]). The cumulative harmonic ordering is thus minimally refined by placing [letku] below [leku]. All of the harmonic orderings asserted by Max are contradicted by the previous cumulative harmonic ordering, except one: [letku] $\succ$ [letu]. Once this ordering has been incorporated into the cumulative ordering, [leku] is the only candidate that remains unbeaten by any other candidate. Therefore, [leku] is optimal. The harmonic orderings asserted by *Pl(lab,dor) are irrelevant, as desired. Note that this optimisation proceeds exactly like the optimisation for /let + ku + jaw/ $\rightarrow$ [lekujaw] in Diola (see (19) in §3.1)), once 'Syncope' has placed [letiku] below all of the other candidates in the cumulative harmonic ordering. Thus the targeted-constraint approach provides a unified analysis of clusters that are created by vowel deletion and clusters that are not.[21]

4.2.2 *Faithfulness to consonant–vowel sequences.* The final contextual faithfulness approach that I consider in this paper is based on the constraint Faith(CV), which is given a correspondence-theoretic definition below.

(42) Faith(CV)

Every CV sequence in the input must correspond to a CV sequence in the output, and *vice versa* (i.e. $/\ldots C_i{}^\frown V_j\ldots/ \Rightarrow [C_i{}^\frown V_j\ldots]$ and $[\ldots C_i{}^\frown V_j\ldots] \Rightarrow /\ldots C_i{}^\frown V_j\ldots/$, where '$\Rightarrow$' stands for implies, $X{}^\frown Y$ indicates adjacency of X and Y, and order is relevant).

Given an input of the form $/V_1C_2C_3V_4/$ or $/V_1C_2V_0C_3V_4/$, Faith(CV) prefers candidate $[V_1C_3V_4]$, which lacks an output correspondent for the first member of the potential $[C_2C_3]$ cluster, over candidate $[V_1C_2V_4]$, which lacks an output correspondent for the second member. Candidate $[V_1C_3V_4]$ preserves the $[C_3V_4]$ sequence of the input. In contrast, $[V_1C_2V_4]$ both destroys the $[C_3V_4]$ sequence of the input and creates a new CV sequence (namely, $[C_2V_4]$) – a double violation of Faith(CV). This contextual faithfulness constraint would therefore seem to capture the generalisation that the final member of an intervocalic consonant cluster

[21] The theory developed in this paper does not determine exactly which member of an intervocalic triconsonantal cluster deletes. But, unlike the Max(strong) approach, the present approach does capture the fact that the *final* consonant cannot be deleted. Note also that certain cases of deletion in triconsonantal clusters also rule out an alternative approach based on special faithfulness to morpheme-initial segments. The relevant type of case is illustrated by the following example from Inuit (Bobaljik 1996: n. 11 and references cited there): /tupiq + pni/ $\rightarrow$ [tupiqni] 'in your house'. In this example, the suffix-initial [p] deletes, exactly the opposite of what morpheme-initial faithfulness would predict. In contrast, the general approach developed here relates deletion of [p] to the fact that it would be poorly cued (because it would not release into a vowel) in the more faithful output *[tupiqpni].

cannot be deleted, even for cases in which the cluster is created by vowel deletion (40).

However, the critique of FAITH(CV) is very similar to that of MAX(strong). First, note that simply positing the constraint is not sufficient to account for the typological generalisation about consonant deletion: if FAITH(CV) is allowed to permute freely with the other constraints in the typology of possible grammars, then violations of the generalisation are inevitable. Therefore, the FAITH(CV) approach must be augmented with some universal ranking. Second, consider that, within the licensing-by-cue framework, the only substantive motivation for according special faithfulness status to CV sequences is the fact that vowels provide strong perceptual cues for consonants that are released into them. If the FAITH(CV) approach does not stipulate that every prevocalic input consonant is released, then CV sequences in the input will not reliably have the cue structure that justifies greater faithfulness to them (as opposed to input VC or CC sequences). But if the FAITH(CV) approach does make this stipulation, then it encounters the duplication problem discussed in the previous subsection (§4.2.1). (As before, the same problem holds if FAITH(CV) is assumed to be motivated by a special licensing property of syllable onsets.)

In contrast, the targeted-constraint approach does not depend on any new universal rankings (see §3). And it maintains the insights of the licensing-by-cue framework without a new restriction on the set of universal inputs (because it does not refer to the cues that a consonant receives in the input, only to those that it receives – or fails to receive – in various candidate outputs).

In addition, empirical considerations appear to favour the targeted-constraint approach over the approach based on FAITH(CV), although space considerations limit me to the following brief remarks. The most straightforward universal ranking that would allow the FAITH(CV) approach to capture the implication in (20), which relates deletion processes to segmental inventories, is FAITH(CV) $\gg$ MAX. I now discuss two undesirable empirical consequences of this universal ranking. (The same points also apply if FAITH(CV) is assumed to be universally undominated.)

(i) If FAITH(CV) universally dominates MAX, then vowel epenthesis as in the Ponapean example /ak + p^{w}uŋ/ → [akup^{w}uŋ] (27) is predicted to be cross-linguistically impossible. This type of epenthesis creates a CV sequence, [ku], that does not exist in the input, and consequently violates FAITH(CV). The alternative mapping in which the first consonant of the potential cluster is deleted, /ak + p^{w}uŋ/ → *[ap^{w}uŋ], violates only lower-ranked MAX, and therefore harmonically bounds the vowel-epenthesis candidate. (See Prince & Smolensky 1993: §9.1.1 for general discussion of harmonic bounding. Recall that the targeted-constraint approach identifies deletion and epenthesis as alternative ways of satisfying a single type of contextual markedness constraint; see §3.3.)

One could in principle solve this first problem by stipulating that FAITH(CV) only applies in the input-to-output direction (i.e. / … $C_i \frown V_j$ … /

$\Rightarrow [\ldots C_i\frown V_j \ldots]$). However, there is a second empirical problem that even the reformulated constraint cannot avoid.

(ii) The well-studied cases of VC infixation in Sundanese (Anderson 1972), Tagalog (Prince & Smolensky 1993) and other languages both destroy CV sequences that exist in the input and create new CV sequences in the output. Schematically, the relevant input–output mappings have the following form: $/V_aC_b + C_1V_2X/ \rightarrow [C_1V_aC_bV_2X]$. A specific example from Tagalog is /**um** + tawag/ → [t**um**awag] 'call (pf., actor trigger)' (Prince & Smolensky 1993: 34). If FAITH(CV) universally dominates MAX, then such infixation should be cross-linguistically impossible. An alternative candidate that trades the FAITH(CV) violation for a MAX violation – thus harmonically bounding $[C_1V_aC_bV_2X]$ – is $[V_aC_1V_2X]$ (i.e. deletion of the affixal consonant). Depending on other constraint rankings, $[V_aC_bC_1V_2X]$ (i.e. simple prefixation) could also win the competition. But, contrary to fact, FAITH(CV)-violating infixation should never be observed.

In contrast, the targeted-constraint approach accounts for infixation of this type in exactly the same way that it accounts for the other non-deletion 'repairs' discussed in §3.3. (I omit the proof of this; one relevant hierarchy would be NOWEAKCONS ≫ MAX, DEP ≫ ALIGN(*vc*-L, PrWd), CONTIGUITY, where *vc* stands for the infix in question.)

In summary, and to conclude this section, the targeted-constraint approach proves superior to all of the alternative approaches – those based on (untargeted) contextual markedness as well as those based on contextual faithfulness – that have been proposed in the literature or that can be straightforwardly extrapolated from existing proposals.

5 Two extensions: feature neutralisation and phonological opacity

Before concluding the paper, I will briefly touch on two further and potentially far-reaching consequences of the targeted-constraint approach (see Wilson 2000 for extensive discussion). The first concerns patterns of contextual feature neutralisation and their relation to the patterns of contextually determined consonant deletion analysed above (see §5.1). The second consequence is that the targeted-constraint approach captures certain cases of phonological opacity in a way that (i) does not require multiple optimisations (cf. the Sympathy Theory of McCarthy 1999, to appear) and (ii) relates the opacity to independent patterns of contextual neutralisation (see §5.2).

5.1 Contextual voice neutralisation

Recall the facts about Lithuanian voice neutralisation from §2.1 above: voiced and voiceless obstruents contrast immediately before a sonorant, but the contrast is neutralised word-finally (Steriade 1997). Some relevant examples were cited in (9), which is repeated below.

(43) *Contextual voice neutralisation in Lithuanian*

a. *Voice contrast maintained before sonorants*

/au**k**le/	au**k**le	'governess'
/au**g**lingas/	au**g**lingas	'fruitful'

b. *Voice contrast neutralised word-finally*

/dau**g**/	dau**k**	'much'
/ka**d**/	ka**t**	'that'

Steriade's (1997) analysis of this pattern makes use of *[α voice]/V __ #, a high-ranking untargeted contextual markedness constraint that is grounded in the weakness of the perceptual cues for voice in word-final position (as opposed to presonorant position); see tableau (10). This kind of analysis can be directly translated into the targeted-constraint approach, as shown in (44).

(44) *Word-final voice neutralisation in Lithuanian*

daug	⇒NoWeakVce	Id[vce]
a. daug	dauK ≻ daug !	
☞ b. dauK		(daug ≻ dauK)
c. dauk	dauK ≻ dauk !	daug ≻ dauk
cumulative ordering	dauK ≻ daug, dauk	dauK ≻ daug ≻ dauk

NoWeakVoice

Let *x* be any candidate and *α* be the [voice] feature of a word-final obstruent (if any) in *x*. If candidate *y* is exactly like *x* except that *α* has been removed, then *y* is more harmonic than *x*.

Following as before Steriade's (1997) description of the Lithuanian data, the outcome of word-final voice neutralisation is taken here to be a segment that is neither [+voice] nor [−voice] (i.e. a segment that has *no* specification for voice). Given a candidate that contains a [voice]-specified obstruent in word-final position (here, [daug] and [dauk]), the targeted constraint NoWeakVoice prefers the candidate that is exactly the same except that the [voice] feature of the word-final obstruent has been removed (here, [dauK]). Ranking this constraint above Ident(voice), as in the tableau, guarantees that deleting the [voice] specification of an obstruent is more harmonic than preserving it in word-final position.

Like NoWeakCons, the targeted constraint NoWeakVoice should be thought of as a formal mechanism for turning certain markedness relations that are established by the weak element principle (12) into harmonic orderings that can affect optimisations. As always, the weak element principle states that, given two representations that differ only with respect to a poorly cued element – and that are therefore difficult to distinguish from an auditory/perceptual perspective – the representation with less structure (here, one fewer [voice] feature) is less marked.

Given the analysis in (44), it is apparent that the targeted-constraint approach formally unifies at least certain patterns of contextual feature neutralisation with the consonant deletion patterns discussed earlier. Future research will test the claim that the new approach is sufficient to account for the typological generalisations (e.g. about direction of assimilation) that pertain only to the contextual neutralisation of features.

5.2 Phonological opacity

In this subsection, I present a targeted-constraint analysis of a counterbleeding pattern observed in Nancowry (Austroasiatic) reduplication. Previous research on reduplication in this language includes Steriade (1988) and Alderete *et al*. (1999); my discussion relies heavily on the latter paper, which presents an OT analysis. Some relevant forms are given below.

(45) *Opacity and transparency in Nancowry reduplication* (Alderete *et al*. 1999)

a. *Transparent interaction: consonant-to-vowel place assimilation in reduplicant*

cɯt	**ʔit**-cɯt	'to go, to come'
kəp	**ʔup**-kəp	'to bite, sting'

b. *Opaque interaction: continuant consonant deletion counterbleeds place assimilation*

tus	**ʔi**-tus	'to pluck out'
hɔw	**ʔu**-hɔw-a	'cave'

The reduplicants in these examples are placed in bold. Each of them begins with a glottal stop – an 'emergence of the unmarked' effect, according to Alderete *et al*. (1999). What interests us, however, is another 'emergence of the unmarked' effect that involves the quality of the *vowel* in the reduplicant. Suppose that coronal consonants and front vowels share the place feature Coronal, and that labial consonants and back rounded vowels share the place feature Labial (Clements & Hume 1995; note that back vowels also have a Dorsal place feature, a detail that is not relevant for the present discussion). Then the quality of the reduplicant vowel in examples such as [ʔit-cɯt] and [ʔup-kəp] (45a) can be determined by local place assimilation: the place feature of the reduplicant-final consonant, which is a copy of the root-final consonant, spreads to the reduplicant vowel. Thus the reduplicant vowel in [ʔit-cɯt] is specified Coronal in agreement with the reduplicant-final [t], just as the reduplicant vowel in [ʔup-kəp] receives its Labial specification from the reduplicant-final [p].

To give a consonant-to-vowel assimilation account of the quality of the reduplicant vowel in examples such as [ʔi-tus] or [ʔu-hɔw-a] (45b), however, would seem to require an intermediate representation. It is a fact about Nancowry that root-final continuant consonants, unlike root-final stops, are not copied by consonants in the surface form of the reduplicant. It is also a fact that the reduplicant vowel in examples like [ʔi-tus] and [ʔu-hɔw-a] has the quality that would be given by place assimilation if the root-final continuant *were* copied in the reduplicant. A serial derivation could make sense of these two facts by initially copying all root-final consonants, then performing consonant-to-vowel place assimilation, and finally deleting reduplicant-final continuants.

(46) *Derivational account of reduplicant vowel quality*

input	a. RED + cɯt	b. RED + tus
Reduplication	ʔut-cɯt	ʔus-tus
Place Assimilation	ʔit-cɯt	ʔis-tus
Continuant Deletion	—	ʔi-tus

The lines of this derivation basically follow the 'full-copy' analysis outlined in Steriade (1988). Most important for the present discussion is the counterbleeding interaction between Place Assimilation and Continuant Deletion, as seen in the derivation of [ʔi-tus] (46b). Place Assimilation changes the place of the reduplicant vowel to match that of the reduplicant-final consonant ([ʔus-tus] → [ʔis-tus]). And then Continuant Deletion removes the reduplicant-final consonant, rendering assimilation to it opaque ([ʔis-tus] → [ʔi-tus]). The counterbleeding rule is *Continuant* Deletion because it spares reduplicant-final stops; thus the system also produces transparent instances of assimilation such as in the derivation of [ʔit-cɯt] (46a).

In their OT analysis of Nancowry reduplication, Alderete *et al*. (1999) propose that there is not, in fact, a counterbleeding interaction between assimilation and deletion in examples such as [ʔi-tus]. To account for the quality of the reduplicant vowel, they propose that the vowel itself stands in base–reduplicant (BR) correspondence with the final consonant of the root (i.e. [ʔi_3-$t_1u_2s_3$]). The vowel is therefore compelled to have the same place as the consonant by the faithfulness constraint IDENT(place)$_{BR}$. (Another IDENT$_{BR}$ constraint prevents a vowel in the reduplicant from corresponding with a root-final stop; see the work cited for details.)

I will now provide an alternative analysis that takes the counterbleeding interaction in (46) to be real, though it does not capture it with a serial derivation. This analysis, which illustrates the general ability of the targeted-constraint approach to account for opaque interactions, has the advantage of relating the quality of the reduplicant vowel in examples such as (45a) to the general assimilation process that is transparent on the surface in examples such as (45b). Comparison with the analysis of Alderete *et al*. (1999), and with other OT approaches to opacity (Benua 1997, Burzio 1996, 1998, 2000, McCarthy 1998, 1999, Goldrick & Smolensky 1999), is beyond the scope of this paper (see Wilson 2000).

In order to focus on the constraint interaction that gives rise to counterbleeding in Nancowry reduplication, I assume the following provisional markedness constraints. AGREE(place), which is adopted from Alderete *et al*. (1999), is violated when the nucleus and the coda of a given syllable do not share any place feature. And the constraint **M**(i) is violated when the reduplicant contains a high front vowel (i.e. [i]). This second constraint represents a basically arbitrary decision to treat the high back vowel [u] as the 'default' reduplicant vowel; changing the default would not affect the aspects of the analysis that are important here. AGREE(place) must outrank **M**(i) in order for place assimilation as in [ʔit-cɯt] (*vs*. *[ʔut-cɯt]) to be optimal.

The remaining two constraints at work in the analysis are not provisional. The first is the targeted contextual markedness constraint NOWEAKCONTINUANT (NOWEAKCONT), a special version of NO WEAKCONS that applies only to continuant consonants. Given a candidate that contains a continuant that is not released into a vowel, NOWEAKCONT prefers the alternative candidate that is exactly the same except that the continuant has been removed (e.g. [ʔi-tus] ≻ [ʔis-tus]). The second and final constraint is R-ANCHOR$_{BR}$, the faithfulness constraint that requires the rightmost segment in the base (which in Nancowry is the root) and the rightmost segment of the reduplicant to stand in BR-correspondence. This constraint is called upon here, as in the Alderete *et al*. (1999) analysis, to force copying of root-final stops (e.g. [ʔit-cɯt]). NOWEAKCONT must outrank R-ANCHOR$_{BR}$ in order to account for the generalisation that root-final continuants remain uncopied on the surface (a generalisation that I continue to refer to as 'continuant deletion'). R-ANCHOR$_{BR}$ must in turn outrank **M**(i) in order for copying of a root-final stop, which leads to place assimilation (e.g. [ʔit-cɯt]), to be more harmonic than failure to copy the stop, which would allow the 'default' reduplicant vowel to emerge (e.g. *[ʔu-cɯt]).

(47) *Rankings for place assimilation and continuant deletion*
  a. Place Assimilation: AGREE(place), R-ANCHOR$_{BR}$ ≫ **M**(i)
  b. Continuant Deletion: NOWEAKCONT ≫ R-ANCHOR$_{BR}$

The rankings in (47) account for the individual place assimilation and continuant deletion processes in standard OT fashion. I now show that further rankings of the *same* constraints derive the counterbleeding interaction between the two processes – a result that would not be possible in standard OT (i.e. if NOWEAKCONT were untargeted). The hierarchy for counterbleeding is given below, followed immediately by the tableau for the opaque mapping /RED+tus/ → [ʔi-tus].

(48) *Hierarchy for counterbleeding interaction*
Agree(place), NoWeakCont $\gg$ R-Anchor$_{BR}$ $\gg$ **M**(i)

(49) *Optimality of the opaque candidate*

red+tus	Agree(pl)	⇒NoWeak Cont	R-Anchor$_{BR}$	**M**(i)
a. ʔus$_3$-t$_1$u$_2$s$_3$	ʔis, ʔu, ʔi ≻ ʔus !	ʔu ≻ ʔus		
b. ʔis$_3$-t$_1$u$_2$s$_3$		ʔi ≻ ʔis !		
c. ʔu-t$_1$u$_2$s$_3$			(ʔus ≻ ʔu) ʔis ≻ ʔu !	
☞ d. ʔi-t$_1$u$_2$s$_3$			(ʔus ≻ ʔi) (ʔis ≻ ʔi)	(ʔus, ʔu ≻ ʔi)
cumulative ordering	ʔis, ʔu, ʔi ≻ ʔus	ʔi ≻ ʔis ≻ ʔus ʔu ≻ ʔus	ʔi ≻ ʔis ≻ ʔu ≻ ʔus	

All of the candidates in this tableau have the same base of reduplication (namely, *tus*), which is therefore omitted from the harmonic orderings.

This optimisation pivots in an interesting way on the relative harmony of the two candidates that satisfy R-Anchor$_{BR}$ (namely [ʔus$_3$-t$_1$u$_2$s$_3$] and [ʔis$_3$-tu$_2$s$_3$]). Neither of these candidates is optimal; but the hierarchy harmonically orders them with respect to each other and the other candidates, and these orderings have consequences for the selection of the optimal candidate. The highest-ranked constraint Agree(place) puts [ʔus$_3$-t$_1$u$_2$s$_3$] below all of the other candidates in the cumulative harmonic ordering. Note that the ordering [ʔis$_3$-t$_1$u$_2$s$_3$] ≻ [ʔus$_3$-t$_1$u$_2$s$_3$] in particular conforms to the general character of Nancowry reduplication: if the reduplicant ends in a consonant, then the reduplicative vowel must assimilate to that consonant's place.

The targeted constraint NoWeakCont, which shares the highest-ranked position with Agree(place), establishes that [ʔi-t$_1$u$_2$s$_3$] is more harmonic than [ʔis$_3$-t$_1$u$_2$s$_3$], and that [ʔu-t$_1$u$_2$s$_3$] is more harmonic than [ʔus$_3$-t$_1$u$_2$s$_3$]. Combining these two orderings with those established by Agree(place) gives the following result. Candidate [ʔi-t$_1$u$_2$s$_3$] is more harmonic than *both* of the candidates that satisfy R-Anchor$_{BR}$. Consequently, R-Anchor$_{BR}$ is unable to place [ʔi-t$_1$u$_2$s$_3$] below any other candidate in the cumulative harmonic ordering. In contrast, candidate [ʔu-t$_1$u$_2$s$_3$] is only more harmonic than the less harmonic candidate that satisfies R-Anchor$_{BR}$ (namely [ʔus$_3$-t$_1$u$_2$s$_3$]). Therefore, R-Anchor$_{BR}$ is able to place [ʔu-t$_1$u$_2$s$_3$] below [ʔis$_3$-t$_1$u$_2$s$_3$]. By transitivity of harmonic ordering, this order also places [ʔu-t$_1$u$_2$s$_3$] below [ʔi-t$_1$u$_2$s$_3$]. Indeed, *every* other candidate lies below [ʔi-t$_1$u$_2$s$_3$] in the cumulative harmonic ordering that is established by the three highest-ranked constraints (Agree(place), NoWeakCont and R-Anchor$_{BR}$). Therefore, [ʔi-t$_1$u$_2$s$_3$] is the only optimal candidate. It wins the competition despite the fact that it incurs a violation of **M**(i) that is not motivated on the surface by place assimilation with a reduplicant final consonant.

The following summary diagram recapitulates the harmonic orderings that are crucial for the optimality of the opaque candidate [ʔi-$t_1u_2s_3$].

(50) *Optimisation summary for counterbleeding interaction*

ʔis_3-$t_1u_2s_3$ ← ① AGREE(pl) ʔus_3-$t_1u_2s_3$

② NOWEAKCONT ↙ ↖ ③ R-ANCHOR$_{BR}$ ↙ ② NOWEAKCONT

☞ ʔi-$t_1u_2s_3$ (④ **M**(i) →) ʔu-$t_1u_2s_3$

(*opaque*) (*transparent*)

One critical path of increasing harmony in this diagram is [ʔus_3-$t_1u_2s_3$] (AGREE(place)) → [ʔis_3-$t_1u_2s_3$] (NOWEAKCONT) → [ʔi-$t_1u_2s_3$], which completely defeats R-ANCHOR$_{BR}$ as far as [ʔi-$t_1u_2s_3$] is concerned. This particular path contains exactly the same representations as the serial derivation in (46). But, because this is an optimisation rather than a derivation, there are many other paths that lead to the same output. Particularly relevant is [ʔu-$t_1u_2s_3$] (R-ANCHOR$_{BR}$) → [ʔis_3-$t_1u_2s_3$] (NOWEAKCONT) → [ʔi-$t_1u_2s_3$], which establishes the non-optimality of the transparent candidate [ʔu-$t_1u_2s_3$].

A somewhat more intuitive description of this optimisation is as follows. The 'general' grammar of Nancowry reduplication, corresponding to the ranking AGREE(place), R-ANCHOR$_{BR}$ $\gg$ **M**(i), prefers reduplication of a root-final consonant and place assimilation (as in [ʔis_3-$t_1u_2s_3$]) over non-reduplication of the root-final consonant (as in [ʔu-$t_1u_2s_3$]) or non-assimilation (as in [ʔus_3-$t_1u_2s_3$]). In addition, a 'special' grammar, corresponding to the ranking NOWEAKCONT $\gg$ R-ANCHOR$_{BR}$, prefers deletion (formally, non-copying) of a continuant consonant (as in [ʔi-$t_1u_2s_3$] or [ʔu-$t_1u_2s_3$]) over more complete reduplication (as in [ʔis_3-$t_1u_2s_3$] or [ʔus_3-$t_1u_2s_3$]). The optimisation finds the correct output ([ʔi-$t_1u_2s_3$]) by interpolating between the 'general' grammar and the 'special' one: it takes the place-assimilated reduplicant vowel from the former and the failure to copy the root-final continuant from the latter. The ranking that is necessary to achieve this result, AGREE(place) $\gg$ R-ANCHOR$_{BR}$, corresponds roughly to an ordering of the 'general' grammar before the 'special' one.

It would be easy to show that the hierarchy in (47) also accounts for transparent mappings such as /RED + cɯt/ → [ʔit-cɯt] (45a). It would also be easy to show that demoting AGREE(place) below R-ANCHOR$_{BR}$ in the hierarchy would make *all* the relevant mappings transparent (e.g. the new output for the input shown in tableau (49) would be [ʔ-$t_1u_2s_3$]). Rather than working through these details, I conclude this subsection by mentioning two strong theoretical advantages of the targeted-constraint analysis of Nancowry reduplication. First, it accounts for the counterbleeding interaction between place assimilation and continuant non-copying with a single optimisation, thus maintaining the most restrictive version of parallelism. Second, it formally connects this case of phonological opacity with independent patterns of contextual neutralisation. Suppression of continuants in positions where they are not cued by a following vowel, as seen in Nancowry (45b), is a fairly common phenomenon. For example, as McCarthy (1998: 44) notes, languages such as Korean and Kiowa ban 'coda fricatives'. This type of contextual restriction can be accounted for in the present framework with the same targeted constraint that was operative in the analysis above (namely, NoWeakCont).[22]

6 Summary and conclusions

A general problem faced by work in Optimality Theory, as in any other theoretical framework, is that of restricting the typology of predicted languages. In this paper, I have addressed a particular typological problem for previous OT approaches to contextually determined consonant deletion. Previous approaches do not capture the generalisation that deletion processes that apply to intervocalic biconsonantal clusters canonically remove the first consonant (1). Instead, they predict unattested processes that consistently remove the consonant that is more marked on some non-contextual dimension (e.g. they predict a process that deletes the consonant with the more marked place specification). I have traced this incorrect prediction back to a deep flaw in previous formulations of the contextual markedness (and faithfulness) constraints. Although previous approaches correctly identify the first consonant in an intervocalic cluster as the 'weak' or 'unlicensed' one, the contextual constraints that they posit do not adequately prefer deletion of the weak/unlicensed consonant over deletion of the following strong/licensed consonant.

The solution that I have proposed employs targeted constraints that prefer *only* deletion of the weak/unlicensed member of a cluster. More generally, I have proposed that the constraints responsible for contextual neutralisation only prefer 'repairs' that affect a weak/unlicensed element itself, not 'repairs' that affect the surrounding context. The substantive basis of this family of targeted constraints is the weak element principle (12), which establishes markedness relations among representations that have nearly identical auditory/perceptual components (and that are therefore easily confused by the hearer). The weak element principle is in turn based on the main insight of the licensing-by-cue framework: namely, that the pressure for contextual neutralisation derives from the absence of perceptual cues. The formal basis of the targeted-constraint

[22] In Wilson (2000), I present a pair of ranking schemas (one for counterbleeding and one for counterfeeding) that are in principle capable of accounting for all cases of phonological opacity of which I am aware. Each of the ranking schemas requires at least one targeted markedness constraint, therefore the generality of the present approach to opacity hinges on the issue of exactly which constraints are targeted. I have not encountered any difficulties in positing the targeted constraints that are required for specific opaque interactions, but clearly further research in this area is called for.

approach is a new definition of OT optimisation that refers directly to the harmonic orderings that individual constraints assert, rather than to the violations (or 'marks') that are standardly used to represent those harmonic orderings ((18), Appendix). The substantive and formal aspects of the targeted-constraint approach have been combined into an explicit account of: the typological generalisation about consonant deletion with which the paper began (see §§2 and 3); other attested processes that 'repair' consonant clusters without deletion (see §3.3); certain aspects of the pattern of contextual voice neutralisation observed in Lithuanian (see §5.1); and a case of phonological opacity (specifically, counterbleeding) in Nancowry (see §5.2).

There are two general conclusions to be drawn from this research. First, TARGETED CONSTRAINTS have been established as an important means for restricting predicted typologies within OT. Second, the family of targeted constraints developed here implies a new role for phonetics in phonology. In addition to being necessary for identifying MARKED STRUCTURES, phonetic factors (here, auditory/perceptual factors) also play a crucial role in directing phonological optimisations toward particular PREFERRED ALTERNATIVES.

Appendix: Order-based optimisation by a constraint hierarchy (formal version)

Let $[\mathbf{C}_1 \gg \mathbf{C}_2 \gg \ldots \gg \mathbf{C}_n]$ be a constraint hierarchy and K be a set of candidate outputs. Each constraint in the hierarchy asserts a (possibly empty) set of pairwise harmonic orderings of the form '$x \succ y$', where x and y are both members of K. For each constraint $\mathbf{C}_i$, let that set of harmonic orderings be designated by $\mathbf{C}_i(\mathrm{K})$.

For each constraint $\mathbf{C}_i$ in the hierarchy, there is a corresponding cumulative harmonic ordering, which is designated by O_i. The cumulative ordering that corresponds to the highest-ranked constraint ($\mathbf{C}_1$) is the transitive closure of the set $\mathbf{C}_1(\mathrm{K})$.[23] Every other cumulative ordering is defined in terms of the previous cumulative ordering, as follows:

(51) Remove every member of $\mathbf{C}_i(\mathrm{K})$ that contradicts the previous cumulative ordering, O_{i-1}. That is, if $x \succ y$ is a member of $\mathbf{C}_i(\mathrm{K})$ and $y \succ x$ is a member of O_{i-1}, then remove $x \succ y$ from $\mathbf{C}_i(\mathrm{K})$. Then add any remaining members of $\mathbf{C}_i(\mathrm{K})$ to O_{i-1} and take the transitive closure. The result is the new cumulative ordering, O_i.

The harmonic ordering established by the hierarchy is defined as the cumulative ordering that corresponds to the lowest-ranked constraint, O_n. A candidate x is OPTIMAL according to the hierarchy iff there is no other candidate y such that $y \succ x$ is a member of O_n.

[23] Definition: The TRANSITIVE CLOSURE of a harmonic ordering P is defined as the ordering P' such that:
(i) For any two candidates x and y, if $x \succ y$ is in P, then $x \succ y$ is in P'.
(ii) For any three candidates x, y and z, if $x \succ y$ and $y \succ z$ are both in P', then $x \succ z$ is also in P'.
(iii) Nothing else is in P'.

REFERENCES

Akinlabi, A. & E. Urea (1993). Foot structure in the Ibibio verb. Ms, Rutgers University & University of Uyo, Nigeria.

Alderete, J. (1999). Head dependence in stress–epenthesis interaction. In B. Hermans & M. van Oostendorp (eds.) *The derivational residue in phonological Optimality Theory*. Amsterdam: Benjamins. 29–50.

Alderete, J., J. Beckman, L. Benua, A. Gnanadesikan, J. McCarthy & S. Urbanczyk (1999). Reduplication with fixed segmentism. *LI* **30**. 327–364.

Anderson, S. R. (1972). On nasalization in Sundanese. *LI* **3**. 253–268.

Anderson, S. R. (1974). *The organization of phonology*. New York: Academic Press.

Beckman, J. (1998). *Positional faithfulness*. PhD dissertation, University of Massachusetts, Amherst. Published 1999, New York: Garland.

Benua, L. (1997). *Transderivational identity : phonological relations between words*. PhD dissertation, University of Massachusetts, Amherst. Published 2000, New York: Garland.

Bobaljik, J. (1996). Assimilation in the Inuit languages and the place of the uvular nasal. *IJAL* **62**. 323–350.

Burzio, L. (1996). Surface constraints versus underlying representations. In J. Durand & B. Laks (eds.) *Current trends in phonology : models and methods*. Salford: ESRI. 123–141.

Burzio, L. (1998). Multiple correspondence. *Lingua* **104**. 79–109.

Burzio, L. (2000). Cycles, non-derived-environment blocking, and correspondence. In J. Dekkers, F. van der Leeuw & J. van de Weijer (eds.) *Optimality Theory : phonology, syntax, and acquisition*. Oxford: Oxford University Press. 47–87.

Casali, R. (1997). Vowel elision in hiatus contexts: which vowel goes? *Lg* **73**. 493–533.

Charette, M. (1990). Licence to govern. *Phonology* **7**. 233–253.

Clements, G. N. (1990). The role of the sonority cycle in core syllabification. In Kingston & Beckman (1990). 283–333.

Clements, G. N. & E. Hume (1995). The internal organization of speech sounds. In Goldsmith (1995). 245–306.

Crowley, T. (1998). *An Erromangan (Sye) grammar*. Honolulu: University of Hawai'i Press.

Fleischhacker, H. (2000). The location of epenthetic vowels with respect to initial consonant clusters: an auditory similarity account. MA thesis, UCLA.

Flemming, E. (1995). *Auditory representations in phonology*. PhD dissertation, UCLA.

Fortescue, M. (1980). Affix ordering in West Greenlandic derivational processes. *IJAL* **46**. 259–278.

Geytenbeek, B. & H. Geytenbeek (1971). *Gidabal grammar and dictionary*. Canberra: Australian Institute of Aboriginal Studies.

Gildea, S. (1995). A comparative description of syllable reduction in the Cariban language family. *IJAL* **61**. 62–102.

Goldrick, M. & P. Smolensky (1999). Opacity, turbid representations, and output-based explanation. Paper presented at the Workshop on the Lexicon in Phonetics in Phonology, University of Alberta, Edmonton.

Goldsmith, J. A. (ed.) (1995). *The handbook of phonological theory*. Cambridge, Mass. & Oxford: Blackwell.

Haas, M. (1941). Tunica. In F. Boas (ed). *Handbook of American Indian languages* 4. New York: J.J. Augustin. 1–143.

Haas, M. (1946). A grammatical sketch of Tunica. In C. Osgood (ed.) *Linguistic structures of Native America*. New York: Viking Fund. 337–366.

Hankamer, J. & J. Aissen (1974). The sonority hierarchy. In A. Bruck, R. A. Fox & M. W. La Galy (eds.) *Papers from the parasession on natural phonology*. Chicago: Chicago Linguistic Society. 131–145.

Hoff, B. (1968). *The Carib language*. The Hague: Martinus Nijhoff.
Hualde, J. (1987). On Basque affricates. *WCCFL* **6**. 77–89.
Itô, J. (1986). *Syllable theory in prosodic phonology*. PhD dissertation, University of Massachusetts, Amherst. Published 1988, New York: Garland.
Itô, J. & A. Mester (1993). Licensed segments and safe paths. *Canadian Journal of Linguistics* **38**. 197–213.
Jenkins, C. (1984). Some aspects of work-formation in a polysynthetic language. *BLS* **10**. 104–115.
Jun, J. (1995). Place assimilation as the result of conflicting perceptual and articulatory constraints. *WCCFL* **14**. 221–237.
Keating, P. (1984). Phonetic and phonological representation of stop consonant voicing. *Lg* **60**. 286–319.
Kidda, M. (1995). *Tangale phonology: a descriptive analysis*. PhD dissertation, University of Illinois, Urbana.
Kingston, J. (1990). Articulatory binding. In Kingston & Beckman (1990). 406–434.
Kingston, J. (1994). Change and stability in the contrasts conveyed by consonant releases. In P. A. Keating (ed.) *Phonological structure and phonetic form: papers in laboratory phonology III*. Cambridge: Cambridge University Press. 354–361.
Kingston, J. & M. E. Beckman (eds.) (1990). *Papers in laboratory phonology I: between the grammar and physics of speech*. Cambridge: Cambridge University Press.
Kiparsky, P. (1973). 'Elsewhere' in phonology. In S. R. Anderson & P. Kiparsky (eds.) *A Festschrift for Morris Halle*. New York: Holt, Rinehart & Winston. 93–106.
Kirchner, R. (1998). *An effort-based approach to consonant lenition*. PhD dissertation, UCLA.
Lombardi, L. (1997). Why Place and Voice are different: constraint-specific repairs in Optimality Theory. Ms, University of Maryland, College Park.
Lombardi, L. (1999). Positional faithfulness and voicing assimilation in Optimality Theory. *NLLT* **17**. 267–302.
McCarthy, J. (1986). OCP effects: gemination and antigemination. *LI* **17**. 207–263.
McCarthy, J. (1998). Sympathy and phonological opacity. Ms, University of Massachusetts, Amherst. (Preliminary version of McCarthy 1999.)
McCarthy, J. (1999). Sympathy and phonological opacity. *Phonology* **16**. 331–399.
McCarthy, J. (to appear). Sympathy, cumulativity, and the Duke-of-York gambit. In C. Féry & R. van de Vijver (eds.) *The syllable in Optimality Theory*. Cambridge: Cambridge University Press.
McCarthy, J. & A. Prince (1993). The emergence of the unmarked: optimality in prosodic morphology. Ms, University of Massachusetts, Amherst & Rutgers University.
McCarthy, J. & A. Prince (1995). Faithfulness and reduplicative identity. In J. Beckman, L. Walsh Dickey & S. Urbanczyk (eds.) *Papers in Optimality Theory*. Amherst: GLSA. 249–384.
Ohala, J. J. (1990). The phonetics and phonology of aspects of assimilation. In Kingston & Beckman (1990). 258–275.
Padgett, J. (1995). Partial class behavior and nasal place assimilation. In K. Suzuki & D. Elzinga (eds.) *Proceedings of the 5th Arizona Phonology Conference*. 145–183.
Paradis, C. & J.-F. Prunet (eds.) (1991). *The special status of coronals: internal and external evidence*. San Diego: Academic Press.
Prince, A. (2000). Comparative tableaux. Ms, Rutgers University. Available as ROA-376 from the Rutgers Optimality Archive.
Prince, A. & P. Smolensky (1993). *Optimality Theory: constraint interaction in generative grammar*. Ms, University of Colorado, Boulder & Rutgers University. To appear, Cambridge, Mass.: MIT Press.
Rehg, K. & D. Sohl (1981). *Ponapean reference grammar*. Honolulu: The University Press of Hawai'i.

Rischel, J. (1974). *Topics in West Greenlandic phonology: regularities underlying the phonetic appearance of wordforms in a polysynthetic language*. Copenhagen: Akademisk Forlag.

Sapir, J. D. (1965). *A grammar of Diola-Fogny*. London: Cambridge University Press.

Samek-Lodovici, V. & A. Prince (1999). Optima. Ms, University of London & Rutgers University, New Brunswick. Available as ROA-363 from the Rutgers Optimality Archive.

Silverman, D.(1995). *Phasing and recoverability*. PhD dissertation, UCLA. Published 1997, New York: Garland.

Steriade, D. (1982). *Greek prosodies and the nature of syllabification*. PhD dissertation, MIT.

Steriade, D. (1988). Reduplication and syllable transfer in Sanskrit and elsewhere. *Phonology* **5**. 73–155.

Steriade, D. (1993). Closure, release, and nasal contours. In M. Huffman & R. Krakow (eds.) *Nasals, nasalization, and the velum*. San Diego: Academic Press. 401–470.

Steriade, D. (1995). Underspecification and markedness. In Goldsmith (1995). 114–174.

Steriade, D. (1997). Phonetics in phonology: the case of laryngeal neutralization. Ms, UCLA.

Steriade, D. (2000). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. Ms, UCLA.

Wilson, C. (2000). *Targeted constraints: an approach to contextual neutralization in Optimality Theory*. PhD dissertation, Johns Hopkins University.

Zec, D. (1994). Coda constraints and conditions on syllable weight. Ms, Cornell University.

Zec, D. (1995). The role of moraic structure in the distribution of segments within syllables. In J. Durand & F. Katamba (eds.) *Frontiers of phonology: atoms, structures, derivations*. London & New York: Longman. 149–179.

Zoll, C. (1998). Positional asymmetries and licensing. Ms, MIT.