CHAPTER 2

PREVIOUS APPROACHES TO TOO-MANY-SOLUTIONS PROBLEMS

2.0 Introduction

In the previous chapter I surveyed the typology of prosody-segmental interactions and argued that the systematic asymmetry in the relationship between prosodic structure and segmental features poses a too-many-solutions problem for OT. This problem has not been systematically addressed in the literature. Apart from an indirect reference to it in de Lacy 2003 (see section 2.2 below), prosody-segmental interactions have escaped the notice of phonological theory.

However, the challenge that the typological data poses for the theory is not entirely new. Several sets of cases where a significant generalization about input-output mappings produces a too-many-solutions problem have been uncovered in recent literature. The problems are, in some ways, analogous to the ones discussed here. In each case, a markedness constraint that states a SURFACE preference is observed to provoke fewer repair strategies than the theory would predict. As Steriade (2001) noted, the problem is akin to a phonological conspiracy in the sense of Kisseberth 1970, although in a new guise. Kisseberth observed that from the point of view of derivational theories, phonological rules conspire to create less marked output structures. Conversely, from the point of view of OT, constraints conspire to avoid certain dispreferred input-output mappings.[5] This way of understanding too-many-solutions problems diagnoses their true origin: OT's radical claim that ALL significant phonological generalizations are located in output structures is too strong. At least some generalizations lie in input-output mappings, and phonological theory must be equipped to handle them.

[5] I will come back to a more detailed discussion of such 'conspiracies' in Chs. 3 and 4.

In this chapter I will discuss three important attempts to handle too-many-solutions problems in OT: Wilson's Targeted Constraints theory (Section 2.1), de Lacy's fixed ranking proposal (Section 2.2), Steriade's P-map hypothesis (Section 2.3), and. In each case, my focus will be on attempting to apply the proposal to prosody-segmental interactions. I will show that the proposed approaches cannot deal with the data in a general way: Wilson's theory fails empirically; Steriade's theory cannot handle covert prosodic structure because of its overly phonetic grounding; and de Lacy's proposal, empirically the best of the three, runs into problems with fixed rankings that are in some cases contradictory and in other cases too parochial.

2.1 Targeted constraints

Wilson (2001) addresses the too-many-solutions problem found in consonant cluster simplification. Apart from the effects of sonority and of morphological structure, Wilson claims that there are no other factors that influence which of two consonants in a cluster will be deleted under simplification. The only influencing factor is a consonant's position in the cluster, and in the configuration VC_1C_2V, it is always the first consonant C_1 that is deleted.

(1) FIRST CONSONANT DELETION (Wilson 2001: 148)
Across languages, deletion processes that apply to intervocalic biconsonantal clusters consistently delete the FIRST consonant (schematically, $VC_1C_2V \rightarrow VC_2V$).

In OT, consonant deletion in clusters is driven by a markedness constraint that militates against sequences of consonants, call it CLUSTERCOND. This constraint is antagonistic to the faithfulness constraint MAX-C that penalizes any deletion of an input consonant. Given the input /VC_1C_2V/, these two constraints cannot in principle choose between the desired output [VC_2V] and the non-occurring output [VC_1V]. Other constraints must make that decision. Generalization (1) cannot be accounted for so long as there is at least one constraint that favors C_1 over C_2, because then the factorial typology will contain the pattern of deletion where C_1 rather than C_2 surfaces. Clearly, such a constraint exists: it could be any one of many markedness constraints that favor C_1 over C_2. The tableau illustrating the unwanted deletion is shown in (2) below.

(2) **M** incorrectly causes deletion of C_2 (Wilson 2001: 148)

	VC_1C_2V	CLUSTERCOND	MAX-C	**M**
	VC_1C_2V	*!		
☞	VC_2V		*	*!
☜	VC_1V		*	

Wilson's point is that markedness IN PRINCIPLE must be made incapable of making the decision about which consonant is deleted and which is preserved in cluster simplification.

> "According to generalization (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 the grounds of MARKEDNESS". (Wilson 2001: 148).

This is a too-many-solutions problem: for a given marked configuration, $/VC_1C_2V/$, OT predicts more repairs than are attested in the typology. The task of Wilson's Targeted Constraints is to render markedness constraints incapable of affecting the cluster simplification pattern.

Wilson's solutions, couched in Steriade's Licensing by Cue theory, is to target constraints to a specific repair. In the configuration VC_1C_2V, the first consonant is deficient as compared to the second in terms of perceptual cues. In general, prevocalic consonants have better cues due to the release burst and formant transitions. Wilson's targeted markedness constraint specifically calls for the deletion of the perceptually weaker non-prevocalic consonant C_1, while deleting the stronger consonant C_2 does not count as a repair of the targeted constraint.

The standard markedness constraint CLUSTERCOND penalizes any candidate that possesses an offending cluster, and it is satisfied by any candidate that does not, no matter what the repair for the cluster is, be it deletion of either consonant, deletion of the whole cluster, or epenthesis. In the harmonic ordering that the constraint imposes, any candidate that contains no cluster is more harmonic than any candidate with a cluster.

(3) a. CLUSTERCOND (Standard)
For any two candidate *x* and *y*, *x* is more harmonic than *y* if *x* contains no consonant cluster and *y* contains a consonant cluster.

b.

VC_1V	$\succ$	VC_1C_2V
VC_2V	$\succ$	VC_1C_2V
VV	$\succ$	VC_1C_2V
VC_1VC_2V	$\succ$	VC_1C_2V

The harmonic ordering asserted by the targeted version of the constraint is weaker: Wilson's NOWEAKCONSONANT (NOWKC) constraint prefers only one specific repair, *viz.* the deletion of the first (weak) consonant in a cluster, while not imposing any harmonic ordering on candidates with other repairs of C_1C_2.

(4) a. NOWKC (Targeted; Wilson 2001: 160)
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. $y \succ x$).

b. $VC_2V \succ VC_1C_2V$

The constraint does not impose any ordering between VC_1C_2V and the other candidates.

The harmonic orderings asserted by targeted constraints may not be expressed with standard violation marks, as Wilson points out in the footnote on pp. 162-163.[6] Wilson's notation therefore departs from the traditional way of using asterisks as violation marks. Instead, in each cell of the tableau, starting from the cells for the highest-ranked constraint, for each candidate the harmonic orderings imposed by the constraint against that candidate are recorded, and cumulative orderings are tallied at the bottom of the tableau. Lower-ranked constraints impose their own orderings, and only those orderings count that are compatible with the cumulative ordering passed down from the higher-ranked constraints. Incompatible orderings are parenthesized.

(5)

	$/VC_1C_2V/$	NoWkC	Max-C	**M**
	VC_1C_2V	$VC_2V \succ VC_1C_2V$!		$(VC_1V \succ VC_1C_2V)$
☞	VC_2V		$(VC_1C_2V \succ VC_2V)$	$(VC_1V \succ VC_2V)$
	VC_1V		$VC_1C_2V \succ VC_1V$!	
	cumul	$VC_2V \succ VC_1C_2V$	$VC_2V \succ VC_1C_2V \succ VC_1V$	

The right-ranked constraint NoWkC only asserts the ordering $VC_2V \succ VC_1C_2V$; it is silent on all other candidates. The faithfulness constraint Max orders VC_1C_2V above both of the deletion candidates, but one of these orderings, $VC_1C_2V \succ VC_2V$, contradicts the ordering passed down by the higher-ranked constraints. Thus the candidates VC_1V incurs a fatal violation. The markedness constraint's contribution is irrelevant.

Wilson illustrates this with a specific example from Diola Fogny, which I repeat here (Wilson 2001: 165). In Diola, complex clusters are simplified by deleting the first member, e.g. /let+ku+jaw/ → *lekujaw* 'they won't go'. The relevant markedness constraints are the place constraints, *Dor ⟩⟩ *Cor. In Standard OT, he higher-ranked *Dor is capable of forcing the deletion of the more marked second member of the cluster /tk/. However, this does not take place if the targeted version of the cluster constraint NoWkC is used.

[6] The reason is that harmonic orderings expressible in terms of violation marks are stratified orders, *i.e.* $a \succ b$ and $a \cong c$ imply $c \succ b$: non-comparable members share order relations. This does not hold for

(6)

	/let+ku+jaw/	NoWkC	Max	*Dor	*Cor
	letkujaw	leku ≻ letku!		(letu ≻ letku)	(leku ≻ letku)
☞	lekujaw		(letku ≻ leku)	(letu ≻ leku)	
	letujaw		letku ≻ letu!		leku ≻ letu
	cumul	leku ≻ letku	leku ≻ letku ≻ letu		

Wilson demonstrates that the addition of the targeted version of the cluster constraint accounts for the typological generalization in (1): the factorial typology of NoWkC, Ident, Max, and markedness does not contain any rankings which lead to the first consonant in a cluster being deleted. Furthermore, the typology is not only restrictive, but sufficient, as it allows vowel epenthesis as a repair for consonant clusters.

In the remainder of the section I will attempt to apply Wilson's targeted constraints theory to the problem outlined in the previous chapter. I will show that targeted constraints cannot account for the relevant typological generalizations, and, following McCarthy's (2002) criticism of Wilson's theory, will diagnose the problem with TC.

Consider the example of stress-aspiration interactions. Given a marked configuration where the stress and aspiration are not located on the same syllable, the Standard OT stress-to-aspiration constraint is unable to rule out the repair where the stress shifts to the syllable with aspiration, while typologically, only the reverse repair, aspiration shifting to the stressed syllable, is attested. In the general case, given a marked structure in the input, and two possible repairs, A and X, where A is typologically attested and X is not, the targeted version of markedness constraint that militates against the marked structure of the input needs to assert the ordering of A above the candidate that violates the non-targeted version of same constraint, saying nothing about the other candidates, including X (McCarthy 2002). Thus, in the stress/aspiration case, the TC would assert that any candidate that repairs the violation by aspirating the stressed syllable is more harmonic than the FFC.

---

targeted constraints, *e.g.* $VC_2V \succ VC_1C_2V$ and $VC_2V \cong VC_1V$ but not $VC_1V \succ VC_1C_2V$.

(7) Targeted STRESS-TO-ASPIRATION (T-ASPIRATE)
Let x be any candidate and α a stressed syllable in x. If candidate y is exactly the same as x except that aspiration is located on α's onset, then $y \succ x$.

Assume that the stress placement constraint is the cover constraint PENULT, and the relevant aspiration faithfulness constraint is NOFLOP-*h*.

For the hypothetical input /patʰa/ there are three relevant candidates: the candidate (a) that violates STRESS-TO-ASPIRATION, [pátʰa], with stress in the default position and aspiration realized faithfully, at the cost of misalignment between the two; (b) the candidate [pʰáta], with stress in the default position and aspiration attracted to it at the cost of violating NOFLOP; and (c) [patʰá], with aspiration surfacing faithfully and stress attracted to it at the cost of violating PENULT.

The theory's task is to ensure that (c) [patʰá] is a perpetual loser. Standard OT fails at this task, as there is a ranking ASPIRATE ⟩⟩ NOFLOP ⟩⟩ PENULT where (c) wins. It turns out, however, that replacing ASPIRATE with its targeted counterpart T-ASPIRATE makes no difference: (c) can still win. This is shown in the following tableau.

(8)

	/patʰa/	T-ASPIRATE	NOFLOP	PENULT
a.	pátʰa	b ≻ a!		
b.	pʰáta		(a ≻ b) c ≻ b!	
☠ c	patʰá			(a ≻ c) (b ≻ c)
	cumul	b ≻ a	c ≻ b ≻ a	

T-ASPIRATE fatally penalizes the candidate (a), where aspiration and stress are disassociated. The remaining two candidates represent the prosodic and the segmental repairs (b) and (c). The segmental repair violates the segmental faithfulness constraint NOFLOP, because it involves aspiration being attracted from its input position to the onset of the stressed syllable. This causes the unwanted candidate (c) to emerge as optimal. The TC theory thus fares no better than Standard OT in failing to rule out the pathological interactions.

This problem is a general one in the domain of stress-segmental interactions where the typological generalization is that only stress influences segments, but not vice versa. McCarthy's (2002) criticism of the Targeted Constraint theory allows us to see why. Observing that the addition of some reasonable constraints to Wilson's analyses undermines the typological predictions of the Targeted Constraint theory, McCarthy formulates the general conditions when replacing a Standard OT constraint with a targeted constraint will account for a typological generalization.

Suppose there is a markedness constraint M, an input /I/, and three candidates, C which violates M, and two candidates that are repairs for M, A and X. Let us assume there is a too-many-solutions problem: candidate A but not candidate X is attested typologically as a repair for violations of M, as illustrated schematically below.

(9)

			M
/I/		C	*
	☞	M	
	☠	X	

The strategy of the targeted constraint approach would be to replace M with a targeted version T-M that imposes the ordering $\{A \succ C|X\}$ on the candidate set. McCarthy shows that this strategy can only ensure that X is a perpetual loser if two conditions hold:

(10) a. A and X are equal in faithfulness, and
b. A is equal to X in markedness on constraints other than T-M.
(McCarthy 2002: 287).

If at least of the two conditions above fails to hold, then there are rankings where X is optimal.

In the case of prosody-segmental interactions under discussion, it is condition (10)a that systematically does not obtain. Consider a markedness constraint M that mentions a prosodic category Π and a segmental property Σ. There are three relevant candidates: candidate C which violates M, candidate A which repairs the violation of M by modifying Σ, and candidate X which repairs the violation of M by modifying Π. Candidate A is the typologically attested segmental repair, and candidate X is the unwanted prosodic repair.

Also, assume that modifying Σ, whatever Σ is, violates the segmental faithfulness constraint FAITH (there must be at least one such constraint). The segmental repair candidate A would violate this FAITH constraint. Neither C nor X would incur violations of faithfulness: the former because it is the fully faithful candidate, and the latter because, by assumption, it involves an unfaithful mapping of prosodic and not segmental structure.

There also must be a prosodic constraint PROS (markedness or faithfulness) which X violates but neither A nor C do: this violation results from 'modifying Π'.[7] Once again, C does not violate it because it is the fully faithful candidate or a candidate with default prosody, while A, by assumption, modifies the segmental but not the prosodic structure in order to satisfy the markedness constraint M.

Finally, the targeted markedness constraint prefers A over C and does not impose any ordering on X. The orderings imposed by the constraints are summarized below.

(11) a. T-M: $\{A \succ C | X\}$
b. FAITH: $\{X \succ A | C \succ A\}$
c. PROS: $\{A \succ X | C \succ X\}$

These orderings indicate that McCarthy's necessary conditions for the targeted constraint to rule out the unwanted candidate do not obtain. Candidates A and X are

[7] If 'modifying Π' means a non-default prosodic structure assignment, then PROS is a markedness constraint. If 'modifying Π' means unfaithfully mapping underlying prosodic structure, then PROS is a faithfulness constraint.

not equal in faithfulness: X satisfies FAITH but A violates it. This means that, ranked high enough, the constraint FAITH can trump the effect of the targeted markedness constraint. The tableau for the general case is given below. It shows that if FAITH outranks PROS, then the unwanted prosodic repair candidate X wins.

(12)

	/I/	T-M	FAITH	PROS
	C	A ≻ C!		
	A		(C ≻ A) X ≻ A!	
☠	X			(C ≻ X) (A ≻ X)
	cumul	A ≻ C	X ≻ A ≻ C	

In other words, replacing the relevant markedness constraint with a targeted version has no effect on the typological predictions of the theory, for the reason that McCarthy made clear: there are constraints that favor the unwanted candidate, and these constraints can trump the effect of the targeted constraint. A and X are not equal in faithfulness, failing to satisfy (10)a: A violates FAITH and X does not.

I conclude that Targeted Constraints do not present a viable strategy for dealing with prosody-segmental interactions.

2.2 Fixed ranking

The most comprehensive attack on too-many-solutions problems as a general challenge for OT comes in Paul de Lacy's work (2003). Recognizing the ubiquitous nature of the problem in phonological interactions, de Lacy proposes a general mechanism of handling cases where a given markedness constraint is observed to condition fewer repairs than OT predicts. He identifies several classes of phonological properties that can and cannot affect each other, mentioning, among other problems, the case of prosody-segmental interactions, although it is not the main focus of the work. De Lacy's strategy is to separate OT constraints into classes based on what type of phonological category they refer to, and to impose fixed rankings between those classes of constraints in order to derive asymmetrical interactions between them. Of all the proposals, de Lacy's constitutes the least radical departure from OT in its formal setup.

In this section I will illustrate de Lacy's approach and attempt to apply it to prosody-segmental interactions. I will argue that the proposal, while empirically sound for many cases, suffers from two weaknesses: first, some cases of prosody-segmental interactions appear to require contradictory fixed rankings, and second, it relies on an arbitrary separation of constraints into categories.

The classes of properties ('representational categories') into which de Lacy divides the phonological world are listed below, together with the classes of OT constraints to which the categories correspond (de Lacy 2003: 2).

(13) a. String structure (the number of segments in a candidate and their order);
MAX, DEP, CONTIGUITY, INTEGRITY
b. Sonority
Sonority-sequencing constraints
c. Prosodic structure
ONSET, NOCODA, FTBIN, etc.
d. Tone
OCP
e. Features
IDENT; Featural markedness

De Lacy addresses three central generalizations about which categories can and cannot influence which other categories.

(14) a. Feature conditions cannot affect string structure or prosody;
Prosody-segmental interactions;
Epenthesis/syncope cannot be sensitive to feature cooccurrence conditions;
b. Tone conditions cannot affect string structure (can affect prosody);
Tonal conditions cannot force epenthesis/syncope
c. Sonority and prosodic conditions can affect string structure and prosody.
Epenthesis/syncope can be conditioned by cluster phonotactics.

De Lacy's proposal is to separate constraints into classes that correspond to the representational categories and to impose fixed rankings between members of those classes. This is shown in (15) below (De Lacy 2003: 5).

(15) THE PROSODIC PRIMACY FIXED RANKINGS (PPFR)

$$
\begin{tikzcd}
 & \text{Faith(StringStructure)} \arrow[dl] \arrow[dr] & & \text{Mk(Prosody)} \arrow[dl, dotted] \arrow[dr] & \\
\text{Mk(Tone)} & & \text{Mk(features)} & & \text{Faith(features)}
\end{tikzcd}
$$

One of the fixed rankings in (15) directly relevant to the problems at hand is Mk(Prosody) ⟩⟩ Faith(Features). De Lacy's supporting example for this ranking is analogous to one of the cases discussed in the previous chapter: the interaction of stress placement with vowel reduction (2003: 9). In a language with reduction of /o/ to [ə] and a default trochaic stress system, a high-ranking IDENT[round] could cause a non-default stress assignment just in case it would circumvent vowel reduction. Below I repeat de Lacy's hypothetical example. Three constraints are necessary: TROCHEE, enforcing default foot structure, REDUCE, penalizing full vowels in unstressed syllables, and IDENT[round], a faithfulness constraint that makes it more costly to reduce to schwa a rounded vowel like /o/ than an unrounded vowel like /a/.

(16) a. TROCHEE
b. REDUCE 'no full vowels in unstressed syllables'
c. IDENT[round]

Because IDENT[round] prefers the reduction of unrounded vowels, it can produce a pattern where stress is attracted to a rounded vowel IN ORDER to prevent its reduction. The following tableau illustrates this hypothetical mapping of /pato/ to *pətó*, with non-default stress assignment in violation of TROCHEE that caters to IDENT[round].

(17) de Lacy (2003: 9)

	/pato/	IDENT[round]	REDUCE	TROCHEE
	páto		*!	
	pátə	*!		
☠	pətó			*

Recall the analogous examples from the previous chapter in a language with reduction of mid vowels, e.g. a language with a five-vowel inventory in stressed syllables and a three-vowel inventory in unstressed syllables. High-ranking faithfulness could cause stress to be attracted to mid vowels to allow them to escape reduction. De Lacy argues that this unwanted interaction can be ruled out by imposing a fixed ranking between prosodic and featural constraints. If all prosodic constraints like TROCHEE are required to rank above all featural constraints like IDENT[round], this type of interaction is ruled out, as shown in the tableau below.

(18) de Lacy (2003: 9)

	/pato/	TROCHEE	IDENT[round]	REDUCE
☞	páto			*
☞	pátə		*	
	pətó	*!		

De Lacy's strategy is a general one. The problematic interactions between consonantal features and stress discussed in the previous chapter are ruled out as long as the faithfulness constraints for those features are required to rank below the relevant stress placement constraints. For example, in the case of stress-flapping interaction, the prosodic constraints are require to rank above the faithfulness constraint violated by the /t/ ↔ [ɾ] mapping, i.e. IDENT[son].

Along the same lines, de Lacy proposes to rank the prosodic markedness constraints above featural markedness. His examples include prosody-segmental interactions of the type where a prosodic category provides the domain for some segmental process, such as agreement in nasality within a foot or in voice within a syllable. The fixed ranking predicts, in de Lacy's words, that "feature conditions cannot force a change in prosodic structure" (2003: 10). The supporting examples are not unlike the examples presented in the previous chapter. AGREE[voice]$_\sigma$, the voicing agreement constraint relativized to the syllable, could force a voice-driven difference in syllabification if sufficiently high-ranked. An input like /adra/ would surface with a complex onset [a.dra], while /atra/ would come out as [at.ra] to satisfy the agreement constraint.

(19)

			AGREE	IDENT[voi]	NOCODA	*COMPLEX
/adra/	☞	a.dra				*
		ad.ra			*!	
/atra/		a.dra		*!		*
		a.tra	*!			*
	☞	at.ra			*	

As mentioned in the previous chapter, attested examples of voice-dependent syllabification show the opposite effect: it is the voiced, not the voiceless consonants that prefer to syllabify as codas, under pressure of syllable contact. What is problematic in the grammar in (19) is the mapping /atra/ → [at.ra], where it is the constraint AGREE which forces the segment [t] into the coda of the preceding syllable, in violation of NOCODA.

De Lacy argues that, once the ranking Mk(Prosody) (NOCODA, *COMPLEX) and Mk(Features) (AGREE) is fixed, the unwanted prediction does not arise.

(20)

			NOCODA	*COMPLEX	AGREE	IDENT
/adra/	☞	a.dra		*		
		ad.ra	*!			
/atra/	☞	a.dra		*		*
	☞	a.tra		*	*	
		at.ra	*!			

If, as in (20), the constraint NOCODA ranks above *COMPLEX, the outputs for both /adra/ and /atra/ have a complex onset. If the reverse ranking holds, both inputs produce forms with a coda. In either case, there is no problematic relationship between syllabification and voicing as in (19).

The following diagram summarizes de Lacy's fixed ranking proposal relevant to the discussion here.

(21)

$$\begin{tikzcd} & \text{Mk(Prosody)} \arrow[dl] \arrow[dr] & \\ \text{Mk(Features)} & & \text{Faith(Features)} \end{tikzcd}$$

In the remainder of this section I will discuss the consequences of de Lacy's proposal for stress-segmental interactions. I will conclude that not all unwanted interactions are ruled out once sonority-stress effects are brought into the picture. I will also discuss some technical problems with de Lacy's proposal.

Along with prosodic constraints, de Lacy calls sonority constraints the 'heavyweights': the phonological categories they refer to are able to influence all other phonological categories. A summary of the sonority constraints is given below.

(22) Sonority constraints
   a. MARGIN{sonority level}, NUC{sonority level}
   b. Syllable contact
   c. Sonority sequencing

Sonority can affect string structure and segments, and therefore all sonority constraints can rank above MAX, DEP, and IDENT constraints. Evidence for Sonority ⟩⟩ StringStructure comes from cases like the deletion of high (i.e. less sonorous) vowels in Arabic, showing that *NUC {i,u} ⟩⟩ MAX. Likewise, there are many examples where sonority factors cause feature changes.

Prosodic constraints, which include syllable structure and foot structure, are likewise able to cause deletion/epenthesis (and thus outrank MAX and DEP), and feature change (and thus outrank IDENT). Sonority constraints must also be able to outrank prosodic constraints, because sonority factors can block prosodic processes, and because sonority considerations can trump prosodic markedness. For example, the constraint *NUC{i,u} must outrank the prosodic constraints NOCODA and *COMPLEX in order to allow processes like the Arabic deletions in /fihimna/ → *fhimna*, /kibirat/ → *kibrat*. The reverse ranking is needed for languages where vowel syncope is blocked by prosodic factors. However, if *NUC{i,u} is allowed to outrank stress constraints, then stress may be attracted to high (less sonorous) vowels in order to avoid syncope – an unattested system. Consider the hypothetical input /pilata/ in a language with default penultimate stress and high-ranking *NUC{i,u}/WEAK (essentially the Arabic constraint).

(23)

		/pilata/	*NUC{i,u}	MAX	PENULT
	a.	piláta	*!		
	b.	pláta		*!	
	c.	pílta		*!	
☠	d.	pílata			*

The candidate (a) is killed right away because it has a high vowel in the unstressed syllable. The next two candidates which try two different syncope patterns are duly eliminated by MAX, and what remains is the pathological candidate where stress is attracted to the high vowel, *pílata*. Note that if there were no high vowel in the input, stress would surface in its default penultimate position, since *NUC{i,u} would not be active.

The reader might wonder whether the bad predictions in (23) arise not because of the ranking of the sonority constraint above stress, but because MAX outranks the stress constraint. Indeed, moving MAX down the ranking at first blush appears to eliminate the problem.

(24)

		/pilata/	*NUC{i,u}	PENULT	MAX
	a.	piláta	*!		
☞	b.	pláta			*
☞	c.	pílta			*
	d.	pílata		*!	

Here, PENULT is ranked high enough to eliminate the unwanted candidate at a point when the reasonable candidates (b) and (c) are still in the running. Other constraints would then decide between them.

De Lacy is silent on the relative ranking of StringStructure constraints like MAX and the prosodic constraints. Let us assume for the sake of argument that there is a fixed ranking Mk(Prosody) ⟩⟩ StringStructure.[8] As I will show, the unwanted candidate (23)d can still win even with this assumption, so the bad predictions must arise from the ranking *NUC{i,u} ⟩⟩ PENULT.

In order for (23)d to win given that PENULT ⟩⟩ MAX, some high-ranked constraint(s) must eliminate the candidates (b) and (c), and these constraints must be of a type that is allowed to outrank MkProsody. It is not difficult to find ways in which (b) and (c) are worse than (a) and (d): for one thing, they have more complex syllable structure and thus violate NOCODA and *COMPLEX, two prosodic constraints that can outrank PENULT. This allows the candidate (23)d to emerge as optimal even if MAX is low-ranked, as illustrated below.

(25)

		/pilata/	NOCODA	*COMPLEX	*NUC{i,u}	STRESS	MAX
	a.	piláta			*!		
	b.	pláta		*!			*
	c.	pílta	*!				*
☠	d.	pílata				*	

[8] This assumption is a stretch: MAX must be able to outrank prosodic constraints like NOCODA – otherwise codas would be universally deleted. DEP must be able to outrank stress constraints like FTBIN, otherwise degenerate feet would not exist.

Just as in the case (23), the input /palata/ would surface with default penultimate stress *paláta* because this candidate does not violate *NUC{i,u}, and we have in effect a stress system where less sonorous vowels attract stress.

In sum, the hypothetical example of {i,u} syncope shows that if sonority constraints are allowed to outrank stress constraints, bad predictions are made. At the same time, sonority constraints must be freely ranked with respect to syllable structure constraints, for reasons outlined above. This calls for a more fine-grained classification if de Lacy's approach is to be maintained: prosodic constraints must be divided into stress (foot structure) and syllable structure, and the two sub-classes come with different fixed ranking stipulations: one but not the other is required to outrank sonority constraints.

Let me now turn to the constraints not mentioned in de Lacy 2003, but discussed by him at length elsewhere (de Lacy 2002): the stress-sonority constraints. These are responsible for quality-driven stress systems. De Lacy proposed the following hierarchies:

(26) a. *HEAD/{ə}, *HEAD/{ə,i,u}, etc.
b. *NONHEAD/{a}, *NONHEAD/{a,e,o}, etc.

There is no a priori way to decide whether these stress-sonority constraints belong to the stress class or to the sonority class (I will come back below to the general problem of finding criteria for sorting constraints into classes). Counting these constraints as prosodic would make for a simpler overall system, for, unlike sonority constraints, *HEAD[ə] and its ilk must be freely ranked with respect to foot structure constraints, in order to account for both sonority-driven and non-sonority-driven stress. It is not crucial that the decision on where to place sonority-stress constraints be made now.

To summarize the discussion so far, below in (27) is the diagram of the more fine-grained structure with foot and syllable prosody separated into two categories.

(27)
$$\begin{tikzcd}[column sep=small]
\text{sonority/stress} \arrow[r, leftrightarrow] & \text{stress} \arrow[r, leftrightarrow] \arrow[dr] & \sigma\text{-struc} \arrow[d, leftrightarrow] \\
 & & \text{other sonority}
\end{tikzcd}$$

However, it turns out that even (27) is not enough. The *NUC{i,u} constraint discussed above is really a sonority-stress constraint from de Lacy's hierarchy, namely *NONHEAD{ə,i,u} (assuming there is no *ə* in the system). This constraint must clearly rank freely with respect to stress constraints, and yet it can still drive the unfortunate stress system where stress is attracted to less sonorous vowels.

So far I have used a cover constraint for stress to refer to several constraints at once, e.g. PENULT should be properly decomposed into AL-FT-RIGHT and FTBIN. Could it be the case that once the more primitive components of constraints like PENULT are considered, it would become possible to establish more parochial, but less contradictory fixed rankings? Once again, the answer is no: consider the ranking between *NONHEAD{ə,i,u} and AL-FT-RIGHT. The sonority-stress constraint must be able to outrank AL-FT-RIGHT in order to account for stress systems where sonority trumps alignment; e.g. in Kiriwina, stress shifts to the antepenultimate syllable whenever the penult contains a low-sonority vowel, showing that *NONHEAD{ə,i,u} ⟩⟩ AL-FT-RIGHT (de Lacy 2003). Conversely, the reverse ranking is needed for a host of languages where sonority plays no role in main stress assignment, such as English or Latin. And yet, despite the fact that AL-FT-RIGHT and *NONHEAD{ə,i,u} must be freely ranked with each other, the pathological stress system of (25) is predicted to exist.

This shows that the problem of asymmetrical prosody-segmental interactions cannot be solved by fixed rankings in the general case: de Lacy's proposal accounts for some cases, but leaves out others that are exactly analogous and call for a unified analysis.

Apart from the empirical incompleteness of de Lacy's account, it faces a difficulty that ultimately makes it a useful summary of the data rather than an explanatory account. One has to be able to decide *a priori* which of the several classes of constraints a particular constraint belongs to. This classification is easy enough with REPRESENTATIONS: we can all agree on what is a feature, what is a tone, what is a segment, and what is prosodic structure. Elementary constraints can be classified straightforwardly as well: the IDENT constraints all refer to features, the MAX and DEP constraints, at least in some versions of OT, refer to segments and string structure.

However, in practice, most useful OT constraints straddle representational boundaries. Such as the constraints of interest for prosody segmental interactions: they all mention two elements, one prosodic and one segmental. Such are the *CODA/[+voice], AGREE[nas]$_{\varphi}$, and many others of the same type. Likewise, the sonority-stress constraints mention both a stress category (headedness) and a featural or sonority category (location on the sonority scale). The same goes for tone-stress constraints.

How to decide where these hybrid constraints should belong? De Lacy's way, which he does not justify, is to count as featural (i.e. belonging to MkFeatures) those constraints that mention features, possibly together with some other property. So, *[+voi] and AGREE[nas]$_{\varphi}$ count as MkFeatures, even though the latter also mentions prosody; AL-FT-RIGHT and NOCODA count as prosodic. Also non-featural are the stress-sonority and stress-tone constraints like *NONHEAD{ə,i,u} and *HEAD-L.

So there seems to be a double standard: if you mention features at least once, you are a featural constraint, but you are a prosodic constraint if you mention NOTHING BUT prosody. 'Features' here must be construed narrowly to exclude tone and sonority. This strategy is required for empirical reasons, but its basis is unclear. De Lacy in effect replaced one stipulation ('features cannot influence prosody') with another ('feature-prosody constraints are MkFeatures').

In sum, de Lacy's fixed rankings proposal is a first serious attempt to handle too-many-solutions problems as a general challenge for OT. It helps to draw a more clear picture of what the generalizations are, but is neither empirically complete nor truly explanatory.

Let me now briefly discuss how de Lacy's fixed ranking proposal might handle the subject of Chapter 4 of this dissertation, the typology of vowel epenthesis and vowel syncope. As I will argue in Chapter 4, syncope and epenthesis are restricted in terms of environments in which they apply. There is a significant procedural generalization about these two processes: epenthesis universally serves to resolve consonant cluster phonotactics, while syncope targets weak (unstressed, unparsed, etc.) vowels. Conversely, epenthesis is not used as a response for violations of metrical constraints such as *CLASH, *LAPSE, FTBIN, and NONFIN. The only case where epenthesis is used for something other than relieving a marked consonant cluster is minimality-driven epenthesis: in some languages, subminimal words are supplied with an epenthetic vowel in order to bring them up to the minimal size. On the other hand, metrically driven syncope can never target stressed vowels. In particular, it cannot be used as a repair for violations of SWP. I leave a detailed discussion of the empirical issues for Chapter 4; for now, I assume the generalizations as stated here.

Let us see what kinds of fixed rankings must be imposed in order to capture these generalizations. The fact that epenthesis is used as a repair for cluster phonotactics means that DEP-V is freely ranked with respect to the constraints against marked clusters. These constraints, as we will see in Chapter 4, include syllable structure constraints such as NOCODA and *COMPLEX, and syllable contact constraints. On the other hand, DEP-V must outrank *CLASH, *LAPSE, FTBIN, and NONFIN, in order to prevent epenthesis as a repair for metrical markedness. The fact that epenthesis is a possible repair for minimality suggests that DEP-V is freely ranked with respect to GRW=PRW.

As for MAX-V, the fixed rankings are quite different. On the one hand, MAX-V must universally outrank SWP in order to prevent stressed vowel syncope. On the other hand, this constraint must rank freely with respect to the other metrical constraints, such as PARSE-σ, in order to produce attested cases of metrically driven syncope. This picture is summarized below: the lines indicate fixed rankings, and the absence of a line connecting two classes of constraints means that the two classes must be freely ranked.

(28) DEP-V MAX-V

SYLLSTRUC METRICAL SWP

Although empirically sound, fixed rankings with this amount of detail become difficult to justify. More seriously, the important generalization on the ENVIRONMENTS of syncope and epenthesis that will be the subject of Chapter 4 is lost in a picture like (28).

2.3 The P-map

The two solutions presented in the previous sections, Wilson's Targeted Constraints and de Lacy's system of fixed ranking were formalist proposals insofar as they attempted a modification of either how the constraints are interpreted (Wilson), or the ways constraints are allowed to interact (de Lacy), with a limited or no attempt to ground the proposals in anything outside the theory.

Another general attempt to limit the power of OT to exclude systematic overprediction of repair strategies comes in Steriade's (2001) P-Map theory. Steriade takes a more directly functionalist view: the formal modifications of OT that she proposes rather mechanically follow from perceptual factors. This section will be devoted to the application of Steriade's ideas to the problems at hand, with special attention to the general conditions when a perceptually-based theory like the P-Map can handle too-many-solutions problems.

Since Lombardi (2001[1995]), it has been known that many laryngeal processes like coda neutralization and final devoicing involve a too-many-solutions problem when analyzed in Standard OT. If a language disprefers final voiced obstruent stops, then THE ONLY way that such stops can be eliminated in the output is by voicing neutralization. However, there are many other imaginable repairs that would also remove the violation of the markedness constraint against final voiced stops, *[voi]#: making the stop into a sonorant, epenthesizing a vowel after it, or deleting it altogether would surely do. And yet, none of these is attested. To be sure, final nasalization, epenthesis, and deletion are common enough, but NOT AS A RESPONSE to final voiced obstruents: if a language deletes final voiced stops, it will also delete voiceless stops as well.

And yet, Standard OT with freely ranked MAX, DEP, and IDENT constraints does not account for the privileged status of devoicing relative to other changes. The rankings that give both the attested pattern of final devoicing and the three impossible patterns are given in (29)a and (29)c-b, respectively. The constraints involved are the standard DEP, MAX, and IDENT constraints, in conflict with the markedness constraint *[voi]#.

(29) a. /tab/ → [tap]
DEP-V, MAX-C, IDENT[nas] ⟩⟩ *[voi]# ⟩⟩ IDENT[voi]
b. /tab/ → ☠ [tam]
DEP-V, MAX-C, IDENT[voi] ⟩⟩ *[voi]# ⟩⟩ IDENT[nas]
c. /tab/ → ☠ [tabə]
MAX-C, IDENT[voi], IDENT[nas] ⟩⟩ *[voi]# ⟩⟩ DEP-V
d. /tab/ → ☠ [ta]
DEP-V, IDENT[voi], IDENT[nas] ⟩⟩ *[voi]# ⟩⟩ MAX-C

Steriade's observation is that preventing IDENT[voi] from ranking above any of the other relevant faithfulness constraints rules out the unattested patterns. The research program is then to couch this fixed ranking in external perceptual factors.

The general strategy is to link the ranking of faithfulness constraints in a grammar to speakers' knowledge of perceptual similarity of potential outputs. This knowledge is encoded in the P-Map, which serves as an interface between speakers' phonetic knowledge and phonological grammars. The novel hypothesis of Steriade's is that given a choice of repair strategies for a marked structure, speakers pick those strategies that involve the least perceptible deviation from the input. In other words, the more perceptibly unfaithful a mapping is, the more cost it incurs, and other things being equal, the least costly modification of the input should be chosen. "Other things" here refers to the markedness violations: another way to state Steriade's hypothesis is that if two outputs fare equally on some markedness constraint, the one that involves a less perceptible deviation from the input is more harmonic.

Applying the proposal to laryngeal interactions, Steriade first demonstrates what the perceptual similarity of the relevant forms is – in other words, constructs a P-Map, arguing from confusion matrices and data on imperfect rhyme. The speakers' knowledge of similarity, in turn, comes from "daily experience with confusability", as well as, potentially, other sources. Steriade arrives at a scale of perceived similarity, with voicing distinctions being less perceptible compared to other feature changes, epenthesis, and deletion. Now there is an explanation for why deletion and other processes like those in (29)c-b are not viable repairs for final voicing: speakers do not pick those repairs because there is always a less perceptible alternative available, namely devoicing.

This explanation depends on [tab] being more similar to [tap] than to [tam], [tabə], or [ta]. The phonological grammar reflects the speakers' knowledge that [tab] is more similar to [tap] than to [tam] though fixed ranking of positional faithfulness constraints. Namely, the constraint against final devoicing, IDENT[voi]/__#, is fixed to rank lower than IDENT[nas]/__#, the constraint against changing the nasality of the final segment, and both constraints rank lower than MAX and DEP.

Is the P-Map applicable to segment-prosody interactions? The structure of the problem is similar to the too-many-solutions situation that Lombardi and Steriade faced: there is a markedness constraint which militates against some configuration, and two ways to avoid a violation of that constraint. Of these two repairs, only one is typologically attested.

The markedness constraint is the familiar prosody-segmental constraint, and the two repairs are the prosodic and segmental modification. In order for the P-map to account for such a too many solutions problem, we would need to establish that the unwanted prosodic repair systematically incurs a greater perceptual cost than the segmental repair.

First, however, a technical issue must be addressed. In the exclusively segmental domain of laryngeal interactions, Steriade could talk about the perceptibility of deviations from INPUTS, because the relevant properties are always present in the input.

The fixed ranking derived from the P-map then held between faithfulness constraints and determined which modifications of the input were better (i.e. less perceptually costly) than others.

This setup sometimes, but not always works for prosody-segmental interactions. Consider first the systems where prosodic structure is contrastive: footing or headedness is encoded in the underlying forms. The generalization then is that input prosodic structure is not modified in order to cater to segmental features. This situation is easily expressible in P-map terms. Take the familiar example, the hypothetical input /pát^{h}a/ with lexical initial stress and aspiration on the second syllable. The fully faithful candidate for this input, [pát^{h}a], would violate the markedness constraint that calls for aspiration to coincide with stress; the two repairs are shifting aspiration or stress from their input locations. A P-map solution to why only the former but not the latter repair is attested would be to establish that the form *pát^{h}a* is more perceptually similar to *p^{h}áta* than to *pat^{h}á*. Here the fully faithful candidate *pát^{h}a* serves as a kind of baseline with which comparisons for perceptual similarity are made. If the perceptual distance facts can be established, then the fixed ranking between faithfulness to stress and faithfulness to aspiration would hold, NOFLOPSTRESS ⟩⟩ NOFLOP-*h*, making it universally more costly to shift stress than to shift aspiration, insofar as both achieve the same markedness results.

The same argument will not quite work with a nearly identical case, where stress is not lexical but predictable, and the conflict is between a stress-segmental constraint and other stress markedness constraints. To modify our example slightly, take the input /pat^{h}a/ in a system with default penultimate stress (rather than lexically specified stress). In order for the P-map to work, we must somehow decide which repairs forced by ASPIRATE incur the least perceptual cost. However, unlike in the faithfulness example, is not obvious which candidate is the baseline of comparison for perceptual similarity, because not all relevant properties are present in the input: stress is not supplied from the underlying form but assigned by constraint raking.

For now let me simply make a workable stipulation about what the baseline should be, so that we can move on to the substance of the P-map proposal rather than the technical details.

(30) To determine the best (i.e. least perceptually costly repairs) forced by a markedness constraint M, the baseline of comparison is the most harmonic candidate that violates M.[9] The candidates which satisfy M are then compared with the baseline, and the most perceptually similar one is the best possible repair.

In the example at hand, we are trying to determine the best way to repair violations of ASPIRATE. The most harmonic candidate that violates this constraint is *pát^{h}a*: it is fully faithful to segments and perfect on stress markedness. This candidate is the baseline of comparison. The next step would be to compare candidates that satisfy ASPIRATE (*p^{h}áta* and *pat^{h}á*) and determine which of the two is more perceptually similar to *pát^{h}a*. The answer would then determine the relative fixed ranking of the constraints violated in the two candidates, PENULT and NOFLOP-*h*, making one universally more harmonic than the other.

Let us assume for now that there is a way to technically extend the P-map approach to pure-markedness cases like most stress-segmental interactions. In the remainder of this section I will discuss the substantive side of Steriade's proposal with respect to the issues at hand.

The general ideology of the P-map theory is surface-oriented, in that surface perceptibility of features is the only factor that determines the choice of repair processes to apply. The conflict that speakers aim to resolve is between their own interests, i.e. articulatory pressures expressed in markedness constraints, and the interests of the hearer, i.e. avoidance of perceptible deviations from established lexical norms.

[9] This definition is superficially similar to McCarthy's definition of the Sympathetic candidate.

> "The view presented here is that speakers are actively concerned with avoiding *perceptible* deviations from established lexical norms, but they are otherwise not averse to linguistic innovation, insofar as it remains covert. […] The P-map serves as the instrument differentiating more from less perceptible innovations". (Steriade 2001: 18)

If Steriade's general proposal is to be extended to the interactions between segments and prosody, the task would be to establish that the attested repairs involve less perceptible deviations than the non-attested ones. Conversely, if it can be established that a certain prosodic repair involves a perceptually greater unfaithfulness than a segmental repair, then the P-Map theory predicts that my generalization should not hold: precisely in those cases, the prosodic repair should be chosen.

Let me start with the limiting case. Unlike most segmental processes, any type of prosodic restructuring can be *covert*, i.e. involve no perceptible consequence at all. On Steriade's preview expressed in the quote above, only *perceptible* deviations are disfavored, while speakers "are not averse to linguistic innovation, insofar as it remains covert". The upshot of this claim is that covert prosodic restructuring should always be preferred to overt segmental changes.

To be more specific, consider coda devoicing again. Suppose that a language has a general dispreference of complex margins, so that inputs like /apra/ are mapped to outputs like [ap.ra]. This indicates the ranking *COMPLEX ⟩⟩ NOCODA. Now suppose that there is a constraint against coda voiced stops, *[+voi]/CODA. There are two logically possible ways of resolving inputs like /abra/: either not have a voiced stop or not have a coda, i.e. either to fix segments or to fix prosody. This would give the outputs [ap.ra] and [a.bra], respectively.

/abra/ → (1) [ap.ra] (Fix segments)
/abra/ → (2) [a.bra] (Fix prosody)

The choice between (1) and (2) depends on the relative ranking of IDENT[voi] and *COMPLEX: if the faithfulness constraint is low-ranked, then the output contains a devoiced stop, but if the faithfulness constraint is high ranked, the output contains a complex syllable margin. This ranking will not affect the mapping of /apra/, as illustrated in the following tableau.

(31)

			*[+voi]/CODA	IDENT[voi]	*COMPLEX	NOCODA
/abra/	☞	[ap.ra]		*(!)		*
	☞	[a.bra]			*(!)	
		[ab.ra]	*!			*
		[a.pra]		*(!)	*(!)	
/apra/	☞	[ap.ra]				*
		[a.pra]			*!	

Once again, the relevant property, syllabification, is not present in the input. In order to use to P-map to determine which repair for the constraint *[+voi]/CODA and input /abra/ is preferable, we need a baseline of comparison. In accordance with the stipulation (30) above, the baseline of comparison is the most harmonic candidate that violates the constraint, namely *ab.ra*. The question then becomes, which of the two repairs, *ap.ra* or *a.bra*, is perceptually more similar to *ab.ra*. Crucially, resyllabification does not necessarily involve any perceptible differences: [ab.ra] and [a.bra] can be pronounced identically. Of the two repairs, the prosodic repair *a.bra* involves no perceptible deviation from the baseline, and the theory predicts it should be universally preferred. The P-map theory not only fails to rule out the unwanted output in this example, but actively ENFORCES it.

There is a way out in examples like these for a proponent of surface-oriented theories like the P-map, which lies in denying the existence of syllable-final devoicing and even syllable structure altogether (Steriade 1999). On such views, the devoicing process that maps /abra/ to [apra] involves preconsonantal rather than coda neutralization, so shifting the syllable boundary to put the offending voiced stop into the onset of the following syllable would not constitute a repair of the markedness constraint violation.

In a more radical theory that denies the existence of syllable structure, the very option of prosodic restructuring does not exist.

However, as I will show below, the P-map solution does not work in the general case, even in situations where it cannot be saved by denying the role of syllable structure in laryngeal neutralization. The general prediction of the P-map is that covert prosodic structure, if it exists at all, should NEVER condition segmental unfaithfulness, because modifying covert structure is by definition perceptually cheaper than modifying anything overt. This is the limiting case of a situation where the choice of the prosodic repair involves a demonstrably smaller perceptual deviation from the input than the segmental repair, and thus, on the P-map view, should be preferred to the segmental repair.

There are many cases in languages when the prosodic boundaries of constituents larger than a syllable are inaudible. For example, the right boundaries of trochaic feet, so long as they are not followed immediately by another trochaic foot, cannot be heard. A form with initial stress and no secondary stresses has at least four possibilities of footing, which are identical on the surface but represent different metrical constituent structure (Kenstowicz 1994).

(32)	a.	$(\acute{\sigma}\ \sigma)\ \ \sigma\ \ \ \sigma\ \ \sigma$	Single trochee
	b.	$(\acute{\sigma}\ \sigma)\ (\sigma\ \ \ \sigma)\ \sigma$	Trochees with no secondary stress realized
	c.	$(\acute{\sigma}\ \sigma\ \ \sigma)\ \ \ \sigma\ \ \sigma$	Ternary foot
	d.	$(\acute{\sigma}\ \sigma\ \ \sigma\ \ \ \ \sigma\ \ \sigma)$	Unbounded foot

Although there is no difference in the realization of stress between the four possibilities of (32)a-d, the covert placement of the right foot boundary could have potentially overt consequences for segmental processes sensitive to footing.[10] Such consequences have been documented for a number of Panoan languages by González (2004), where various vowel and consonant alternations are sensitive to covert feet. One such case is the

[10] Further evidence for the reality of covert constituent boundaries comes from metrics: the meter of Plautus makes reference to covert right foot boundaries of trochaic feet (Blumenfeld 2004).

process of glottal stop deletion in Capanahua (Loos 1967, Safir 1979, Loos 1986, González 2002). Primary stress in Capanahua falls on the second syllable if it is heavy, otherwise on the first syllable. The only reported realization of stress is high pitch, which falls on the stressed syllable and spreads rightward up to the penultimate syllable. Thus Capanahua words can take one of the two shapes schematized in (33).

(33) a. Initial stress: (σ̲́ σ́) (σ́ σ́) σ
b. Peninital stress: σ (σ̲́ σ́) (σ́ σ)

There are no H/L pitches alternating in a binary fashion that would overtly signal the presence of secondary stress feet posited in (33), nor is there any other correlate of secondary stress (Loos 1967, González p.c.). However, the covert feet must be present, because they condition the deletion of the glottal stop. As shown in the following examples in (34), glottal stops are deleted foot-finally, so that the morpheme sequence /ɾaʔ-taʔ/ surfaces either as [ɾaʔ-ta] or as [ɾa-taʔ], depending on whether it is preceded by an even or an odd number of syllables.

(34) a. /ʔotʃiti-ɾaʔ-taʔ-ki/ (ʔotʃi)(ti-ɾa-)(taʔ-ki) 'it is probably a dog'
'dog-prob-decl-cert'
b. /ʔotʃiti-ma-ɾaʔ-taʔ-ki/ (ʔotʃi)(ti-ma-)(ɾaʔ-ta-)ki 'it is probably not a dog'
'dog-neg-prob-decl-cert'

c. /ta-mani-ʔt-wi/ (tama)(niʔwi) 'take a step'
/ka-riʔbi-wi/ (kari)(biwi) 'go again'

According to the predictions of the P-map, examples such as the Capanahua alternation should not exist, because Capanahua sacrifices segmental faithfulness to covert prosody. The prosody-segmental interaction in Capanahua is in the same direction as the prosody-segmental interactions in English and other languages where the conditioning prosodic structure is audible and overt: it is the segments that cater to prosody, not the other way around.

One possible answer to this argument against the P-map would appeal to the possibility that covert features structure is, in general, dispreferred. Languages strive for ways of overtly realizing the covert metrical structure, which is normally done through pitch, intensity, and duration. Any such realization of metrical prominence would involve some unfaithfulness, i.e. some perceptual cost. For example, if a language lengthens its stressed vowels, it makes it more difficult to maintain a length contrast in those positions. There is a tradeoff between how clearly the metrical prominence is realized and how much faithfulness cost is involved in its realization. Consonantal processes sensitive to foot structure are then simply another way of signaling the position of the foot boundaries, and Capanahua, in this line of thinking, can be said to employ segmental cues to metrical structure rather than prosodic ones. The glottal stop deletion process, on this view, would not be an unmotivated instance of unfaithfulness, but rather a functionally motivated means to supplying the listener with cues to the covert metrical structure.

However, the Capanahua example illustrates that in the general case this counterargument is invalid. Glottal stops are contrastive in the language, while secondary (as well as primary) stress is not. Thus there should not be any functional motivation to signal non-contrastive boundaries of feet. And yet, the covert and predictable feet are preserved, while the unpredictable glottal stops are deleted, resulting in their neutralization with ∅.

Furthermore, the argument that segmental alternations like the glottal stop deletion in Capanahua have a function in signaling covert metrical structure makes a more general prediction that such alternations should not occur, or at least occur less frequently, in languages with clearly audible dynamic stress. This is, however, not the case; according to the survey in Bybee et al. (1998), the reverse generalization is true: the more perceptible the metrical structure, the more likely the language is to have segmental alternations sensitive to it.[11]

[11] This generalization indicate that at least some of the segmental alternations conditioned by prosody are in fact conditioned not by the abstract constituent structure, but by its overt correlates such as intensity,

The P-map makes an even stronger and more clearly false prediction. Because innovations in covert structure have no perceptual cost, they should be used whenever possible if the covert innovation leads to a faithful mapping of overt segmental information. For example, if a language has a process of vowel reduction in unfooted syllables, then there is always the possibility of constructing a covert secondary stress foot over those reducible syllables in order to save them from reduction. According to the letter of the P-map proposal, this is what ought to happen: covert innovations are free, while the danger of reducing vowels has perceptual consequences. Likewise, if a language (e.g. Mam, England 1983) has a constraint against multiple long vowels in the same word – i.e. a constraint against long vowels in any position except in the head of a stress foot – nothing prevents the language from constructing covert feet over long vowels in other positions of the word in order to allow them to faithfully surface. Again, on the P-map proposal this should be the solution of choice, because there is no cost in the covert innovation, while the benefit of building a secondary stress foot would be the faithful rendition of a vowel that would otherwise shorten.

Looking at less extreme examples where prosody is covert, the exact predictions of the P-map view are difficult to tease out: there is nothing GENERAL about how prosody and segments ought to interact that follows from the setup of the theory. Rather, the predictions are to be evaluated on a case-by-case basis, determining in each type of interaction whether being unfaithful to prosody or unfaithful to segments incurs the least perceptual cost. One general expectation that comes out of the P-map is that the behavior of a prosodic system may depend on how it is phonetically realized: high-intensity dynamic stress may be more salient and thus more resilient to unfaithful mappings than a pitch accent. This prediction is not borne out. In general, the

duration, and high pitch. This is clearly true of consonantal alternations like Verner's Law, where the voicing of obstruents depends on the accentuation of neighboring syllables. Phonetically plausible accounts of Verner's Law appeal to high pitch, rather than purely abstract prominence, as the conditioning factor. Likewise, aerodynamic factors are likely explanations for English-style alternations like aspiration and flapping. At the same time, as argued in detail in González 2003, there are many alternations that are conditioned by foot structure not reducible to phonetic correlates of prominence. Capanahua, with no phonetic correlates at all, is one such example.

directionality of prosody-segments interaction does not depend on whether the system is one of dynamic stress or pitch accent. Pitch accents cannot be sensitive to purely segmental properties in the same way that dynamic stress cannot be.

One ready diagnostic for perceptual similarity appealed to by Steriade is imperfect rhyme. On the assumption that more perceptually similar strings make better rhymes, it is easy to deduce relative similarity of features and segments from their (non-)use in rhyme. This method can also shed some light on the predictions of the P-map for prosody-segmental interactions. In dynamic stress systems imperfect rhyme practically never involves a stress mismatch, even if very liberal with segmental differences between members of a rhyming pair. In both English and Russian, voicing, place of articulation, and vowel quality may all mismatch in imperfect rhyme, but stress almost never can. The fact that prosodic agreement within the rhyme is an absolute requirement suggests that differences in dynamic stress are more perceptible that differences in segments, leading to the P-map style argument that segmental repairs are preferred to prosodic repairs.

While this prediction is correct for dynamic languages, the argument is not so easy to make if we look at the prosodic systems where prominence is realized in some other way. In Serbo-Croatian the metrical heads surface as pitch accents, and, unlike in English and Russian, imperfect rhyme CAN involve accent mismatch (Eekman 1974).

(35) tvôm ~ mȉlom
kȁmen ~ plên
starína ~ sȑbina (Kostić, Eekman 1974)

And yet, the generalizations about the interaction of pitch accent systems with segments is the same as for dynamic stress systems. This shows that while the imperfect rhyme facts may go with the phonetic details of a prosodic system, the role of that system in the grammar and its interface with segmental phonology does not.

The reason why these incorrect predictions of the P-map arise is due to the surface-oriented nature of the theory: the only factor that plays a role in phonological structure, according to the P-map, is the phonetic realization of a feature. However, metrical constituent structure appears to possess a level of abstractness with which the P-map cannot deal.