The Phonology of Perceptibility Effects: the P-map and its consequences for constraint organization

Donca Steriade, *UCLA*

1. Introduction

This study outlines a proposed revision in the structure of Optimality Theoretic (OT) phonologies (Prince and Smolensky 1993). The proposal is to let a distinct grammatical component, which I call the P-map (P for *perceptibility*) project correspondence constraints and determine their ranking. The P-map is a set of statements about relative perceptibility of different contrasts, across the different contexts where they might occur. For instance, the P-map will be the repository of the speaker's knowledge that the [p]-[b] contrast is better perceived before V's (e.g. in [apa] vs. [aba]) than before C's (e.g. in [apta] vs. [abta]). The point of departure here is the theory of correspondence set forth in McCarthy and Prince 1995, with its distinction between MAX/DEP constraints, which identify the elements of two representations that stand in correspondence, and the Ident F constraints, which require a precise featural match between correspondent elements[*].

The general rationale for the P-map proposal is that attested phonological systems display less diversity than predicted by versions of OT in which correspondence and phonotactic constraints interact freely. In particular, the range of pairings between constraint violation and "repair strategy" is more limited than current versions of OT will lead one to expect. An example of this need for a tighter fit between predictions and typology involves the effect that constraints on obstruent voicing have on phonological systems. Consider a common constraint like (1), an underlying string like /tæb/, which violates (1), and the range of possible responses of the grammatical system to this violation, as sketched in (2).

(1) A phonotactic constraint:

*[+VOICE]/_]$_{word}$: voiced obstruents do not occur at the end of the word[1].

(2) Conceivable grammatical responses to the violation of (1) in UR / tæb/[2]:

Change in UR, to satisfy (1)	*Corresponding constraint ranking*
a. Devoicing: /tæb/ -> [tæp]	*[+VOICE]/_] >> Ident [±voice]
b. Nasalization: /tæb/ -> [tæm]	*[+VOICE]/_] >> Ident [±nasal]
c. Lenition to glide: /tæb/ -> [tæw]	*[+VOICE]/_] >> Ident [±consonantal]
d. C-Deletion: /tæb/ -> [tæ]	*[+VOICE]/_] >> MAX C
e. V-Insertion: /tæb/ -> [tæbə]	*[+VOICE]/_] >> DEP V
f. Segment reversal: /tæb/ -> [bæt]	*[+VOICE]/_] >> Linearity (segments)
g. Feature reversal: /tæb/ ->[dæp]	*[+VOICE]/_] >> Linearity (for features)

The table in (2) should be read on the understanding that the correspondence constraint named in a given cell, is the lowest ranked among the correspondence constraints that are in potential conflict with the phonotactic. Of the changes in (2), only the devoicing in (2.a) is actually attested as a reaction to *[+VOICE]/_] violations. This is not surprising if one consults one's linguistic intuition, but it is unexpected in the context of OT, in its present form: if the ranking between the correspondence constraints in (2.b-g) is free, one expects at least the range of fixes shown in (2).

My claim is not that nasalization and C-deletion etc are unattested processes: but that they are unattested as responses to the voicing problem posed by (1). This means that one does not encounter sound systems in which all the final voiced stops, *and only they*, turn to nasals, delete, or trigger epenthesis or metathesis. (3) indicates what systems of alternations would look like if some of these changes did occur.

(3) Unattested systems (lexically related forms linked by arrows)

Nasalization of final voiced obstruents

(a) Morpheme shapes before V:	**tud**-a,	tat-a,	**tib**-a,	top-a,	**tag**-a,	tek-a
	↕	↕	↕	↕	↕	↕
(b) Word final:	**tun**,	tat,	**tim,**	top,	**taŋ,**	tek

Deletion of final voiced obstruents

(a) Before V:	**tud**-a,	tat-a,	**tib**-a,	top-a,	**tug**-a,	tek-a
	↕	↕	↕	↕	↕	↕
(b) Word final:	**tu,**	tat,	**ti,**	top,	**tu**,	tek

Epenthesis after final voiced obstruents

(a) Before V:	**tud**-a,	tat-a,	**tib**-a,	top-a,	**tug**-a,	tek-a
	↕	↕	↕	↕	↕	↕
(b) Word final:	**tudə,**	tat,	**tibə,**	top,	**tugə,**	tek

The diagnosis for the problem encountered - the fact that devoicing is the only available cure to violations of (1) – starts with the observation that of all the input-output pairs displayed in (2), the one judged most similar is the pair [tæb]-[tæp] in (2.a). (Evidence for the relevant hierarchy of similarity is reviewed in section 4.) The aim, in any departure from the UR, is to change it minimally to achieve compliance with the phonotactics. The modifications in (2.b-g) are less minimal, as they result in greater input-output dissimilarity, than devoicing. This is why they are systematically avoided.

If this is the correct diagnosis, then what is needed is a mechanism that relates rankings between correspondence constraints to perceived differences in degree of similarity. I refer to knowledge of relative phonological similarity as the P-map. The function of the P-map we discuss here is that of guiding the speaker in search of the minimal input deformation that solves a phonotactic problem. A set of similarity rankings will also be needed in lexical access: the listener has to pick which lexical entry most closely resembles a frequently ambiguous auditory input[3]. The grammatical reflex of the P-map is the projection and ranking of correspondence constraints. Thus, if the P-map identifies the pair [p]-[b] as more similar in the context V_] than [b]-[m] for the same context, then the P-map's effect on the grammar will be to rank higher the faithfulness condition corresponding to the more distinctive contrast [b]-[m], hence Ident [±nasal]/V_] >> Ident [±voice]/V _]. This idea is outlined in (4) using the same example as illustration:

(4) P-map effects on the ranking of correspondence conditions

P-map comparisons	*More distinctive contrast* (e.g. [b]-[m] in V_]	 vs.	*Less distinctive contrast* [b]-[p] in V_])
Ranking of correspondence constraints	*Higher ranked constraint* (e.g. Ident [±nas]/ V_]	 >>	*Lower ranked constraint* Ident [±voice]/ V_])

The reader will note the parallel with Prince and Smolensky's treatment of phonetic scales (1993): the grammatical reflex of the physical scale is the fixed ranking between constraints referring to points on the scale.

Consider now present OT. The concept of minimal modification embodied in this theory is the candidate that optimizes satisfaction of correspondence constraints, as ranked in a given grammar. No independent principle determines the ranking of potentially conflicting correspondence conditions. This means, in the context of the example in (2), that either [tæp] or [tæm] will count as minimal modifications of the input /tæb/, depending only on the unconstrained ranking between Ident [±nasal] and Ident [±voice]. I assume below that the phonotactic (1) induces *some* input modification: the question is which.

(5) Devoicing as minimal modification

/ tæb/	Ident [±nasal]	Ident [±voice]
☞ tæp		*
tæm	*!	

(6) Nasalization as minimal modification

/ tæb/	Ident [±voice]	Ident [±nasal]
tæp	*!	
☞ tæm		*

The problem with this view is that for at least some phonological properties, and perhaps for all, there exists a cross-linguistically constant notion of minimal modification: that is why a violation of (1) is only resolved by (2.a) and not in other ways. This study is a contribution to our understanding of this notion.

The difficulty outlined in (2) – which I call "Too-Many-Solutions"– arises with particular clarity in Optimality Theory. This is because OT views phonology as a problem solving system: the problem is the conflict between phonotactic constraints and lexical forms that violate them, such as the UR /tæb/ of (2). The Too-Many-Solutions conundrum arises when the system of constraints and rankings predicts too many resolutions of a given phonotactic problem. But the same difficulty comes up in any other approach to phonology in which changes in the underlying form are seen as the sound system's responses to phonotactic violation[4]. Thus Kisseberth's (1970) insight that conspiracies arise when the sound system aims at a specific target structure via multiple means leads to equivalent questions in the context of rule-based phonology: if phonologies aim to eliminate final voiced obstruents, why don't they employ obstruent nasalization, deletion, metathesis?

2. Sources of evidence for the P-Map

The P-map hypothesis is that speakers possess judgments of relative similarity of the form in (7) and use these to determine a partial ranking of correspondence constraints.

(7) The pair of strings x-y is less similar than the pair w-z

(abbreviated as Δ(x-y) > Δ (w-z), where Δ = difference).

If some of these judgments are cross-linguistically invariant, then the ranking of certain correspondence conditions will be invariant too. The invariant correspondence ranking will give rise to typological laws suggested earlier. We observed that if a final voiced obstruent must be avoided, then the repair is to devoice it rather than nasalizing it. The P-map conjecture is that source of this law is a similarity ranking: Δ (oral C-nasal C) >

Δ(voiced C-voiceless C). This similarity ranking induces a correspondence ranking, Ident [±voice] >> Ident [±nasal]. If the similarity ranking is constant, the correspondence ranking is constant too – with qualifications discussed below - and this fact explains the law.

Empirical tests of the P-Map hypothesis can proceed by observing invariant preferences for one repair strategy over conceivable others. Some of these observations are discussed below and in earlier work (Fleischhacker 1999, Jun 1995, Steriade 2001). What must be verified then is that each invariant preference corresponds to a constant judgment of relative similarity. Conversely, one can start from documented judgments of relative similarity and verify the P-map's prediction that these judgments correlate with preferences for certain input modifications as against others. Both strategies are currently hampered by the scarcity of direct evidence on relative similarity. Few studies address this question and their vast majority compare different contrasts in the same position (say *ba-pa* vs. with *ba-ma*[5]) rather than the same contrasts across positions (e.g. *ba* vs. *pa* compared to *ab* vs. *ap*).

However an additional hypothesis about the nature of similarity judgments expands the range of evidence bearing on the P-map. This second hypothesis is that phonological similarity is evaluated on perceptual representations (cf. Flemming 1995 (2001)[6]). This means that the representations inspected in evaluating similarity are those which encode the perceptual correlates of a contrast – the auditory properties that differentiate its terms. These properties vary with context. The voicing contrast, for instance, is conveyed in part by VOT differences but these are unavailable when the C'ss occur in final position. We can reason then that a voicing contrast like *b* vs. *p* will give rise to different similarity judgments depending on context: the pair *aba* vs. *apa* will register as more dissimilar than *ab* vs. *ap* if one of the voicing correlates is missing word finally. More generally, the hypothesis that phonological similarity is perceptual in nature turns evidence about cue distribution into a source of information about similarity.

As a result, confusion rate data (Miller and Nicely 1955, Wang and Bilger 1973) is potentially relevant to the P-map hypothesis: if the pair of auditory stimuli x-y give rise to a higher rate of confusion than the pair z-w that may be because their auditory representations are less well differentiated. This is at least one factor in the *phonological* similarity judgment Δ(x-y) > Δ(z-w). Confusability and similarity are distinct matters: subjects rating similarity of auditory stimuli do not report confusion, just similarity. But if similarity is evaluated on perceptual representations, then confusability and similarity are related, because they share at least one source: some auditory representations are differentiated by fewer or less salient properties than others. This makes them more similar and, in the limit, more confusable.

In substantiating a claim of relative similarity one can rely on speakers' direct judgments of similarity; or refer to confusion studies indicating that one contrast is more perceptually robust than another; or reason from the observation that in the position considered, one contrast misses an acoustic correlate and the other does not. Finally, there is the evidence of linguistic tasks that imply a judgment of identity or near-identity: rhyming and alliteration. On the former see Zwicky 1976, Hanson (this volume), Steriade and Zhang 2001. All forms of evidence are employed below.

Beyond these observations, no model of similarity computation is offered here. (See Frisch, Broe and Pierrehumbert 1997 for a model of *context-free* similarity evaluated on *articulatory* representations.) Rather, the strategy here is to observe that subjects hold certain judgments of relative similarity – whatever their source may be - and that the preference for certain repair strategies correlates with these judgments.

This study focuses on only some of the predictions of the P-map hypothesis, as they affect voicing neutralization, cluster resolution, epenthesis and V deletion. The aim is to show in each case that there exist preferred methods of resolving phonotactic violations; that these preferences are not accounted for by currently available mechanisms; that each preference for a particular solution is explained by the idea that the least perceptually distinctive contrast whose modification removes the violation is always the one sacrificed; and that a solution can be obtained by ranking correspondence constraints via the P-map.

3. The P-map

The P-map is a mental representation of the degree of distinctiveness of contrasts in various positions. It can be viewed as a set of statements about relative similarity between sounds or other phonological properties. The example below involves a hypothetical set of similarity rankings about voicing distinctiveness in different contexts:

(8) Δ (p-b/V_V) > Δ (p-b/_V) > Δ (p-b/V_) > Δ (p-b/C_C)

(8) states that *p* and *b* are less similar inter-vocalically than prevocalically, even more similar postvocalically, and maximally similar between C's. The statement can be generalized to a broader class, such as all obstruent pairs that differ only in voicing, if similarity rankings do not differ across class members. The P-map fragment in (8) embodies the hypothesis that voicing differences are equally perceptible for all pairs of

obstruents but not equally perceptible across positions, the optimal context for this being the intervocalic position (Steriade 1999).

Several properties of P-map statements like (8) are critical here. As already indicated, they must reflect the effect of the syntagmatic context on perceived similarity. Voicing contrasts are not equally well perceived in all positions, and this has fundamental effects on the phonology of voicing. Likewise, the distinctiveness of a contrast between any specific segment class and its absence (here Ø) is also affected by context. For instance, pairs like [fɪts]-[fɪs] – with a t-Ø difference in the V_C context – register as less similar than pairs like [fɪst]-[fɪs] – with the same t-Ø difference in the C_] context[7]. A corresponding P-map statement appears below.

(9) Δ (t-Ø/V_C) > Δ (t-Ø/C_])

Relative distinctiveness judgments can involve suprasegmental or serial position differences. Thus (10) below records a hypothetical similarity ranking: C_1C_2 reversal under adjacency is less distinctive than the reversal of C's separated by a V.

(10) $\Delta (C_1VC_2\text{-}C_2VC_1) > \Delta (C_1C_2\text{-}C_2C_1)$

According to (10), a pair like *apsa-aspa* is more similar than the pair *apas-asap*. The point here is not to examine the truth of or rationale for (10) but to illustrate how the terms of similarity comparisons may involve strings larger than one segment[8].

While relative similarity judgments can be documented among many string pairs, one must also consider that no knowledge of relative similarity is available in other cases. Another hypothetical comparison illustrates this: is the pair *task-tass* more or less similar than *task-task[ə]*? Suppose that the P-map fails to record a difference in this case: then the P-map cannot be the source of any ranking between correspondence constraints prohibiting either type of modification (a variant of MAX C in *task-tass*; a variant of DEP V in *task-task[ə]*). This doesn't mean that the two constraints cannot be ranked: but if cross-linguistically invariant correspondence rankings emerge only from cross-linguistically invariant similarity rankings, then the MAX C vs. DEP V ranking will be free to vary on a language specific basis.

The final point to emphasize about the P-map is that distinctiveness and its opposite, similarity, are properties of contrasts (Flemming 1995, 2001): the statement "a is more perceptible than b" means "a is more reliably distinguished from a reference term x than b is distinguished from x". It is not the sounds or the articulations a and b that are being compared for perceptibility but the contrasts a vs. x and b vs. x. This point is fundamental to the success of the P-map as an analytic tool and follows from the assumption that the P-map is so structured as to permit a definition of the concept of minimal modification. Consider what information is needed to discover the minimal modification that will render a representation like /tab/ compatible with a constraint like *[+VOICE]/_]. Since several modifications of this input achieve compatibility with this constraint the P-map must indicate which one among these represents the minimal modification. For example, we compare [tap] and [tabə] as potential modifications of input /tab]. The necessary comparison involves then the input-output pairs /tab/-[tap] and /tab/-[tabə/. If there is a guide to the minimal modification, this guide must exist in the form of statements about the relative perceptibility of contrasts like these, or their generalized forms. The contrast is that between the unchanged input string and its modified output correspondent, as it occurs in the context of the modification.

There is another sense of contrast and another sense of perceptibility, distinct from the one used here. Suppose that there are invariant properties that underlie sound classes – either invariant acoustic properties or articulatory gestures common to all manifestations of a given class. Then we can talk about the fact that some context allows a better recovery of these invariants. For instance, suppose that the invariant properties of [b] may be better recovered intervocalically than interconsonantally. What that means is that we can better distinguish [b] in V_V from *all other sounds* that could have occurred there. This is the broad sense of contrast. This may be a useful notion but not for the purpose of defining the minimal modification: it doesn't tell us which pair - /tab/-[tap] or /tab/-[tabə] - is the most similar input-output pair.

4. A P-map account of voicing neutralization

4. 1. Differences of relative similarity

As a preliminary to the P-map analysis of devoicing – our answer to one aspect of the Too-Many-Solutions problem – I outline now the evidence for a hierarchy of perceived similarity between the pairs in (11). Each pair corresponds to the contrast between an input string with a voiced obstruent in final position and its modified counterpart.

(11)
a. D vs. T/V_] D = voiced stop, T = voiceless stop
b. D vs. N/V_] N = nasal
c. D vs. G/V_] G = glide or lateral
d. C vs. Ø/V_]

e. Ø vs. V/C_]

f. C_1VD_2] vs. D_2VC_1]

The present task is to show that, among these, the voicing contrast D vs. T/V_] is least distinctive, i.e. that its terms are perceived as more similar than those of the other pairs. In fact, the voicing contrast (11.a) stands out because it is the only one that lacks its primary perceptual correlate: the VOT value. This is one reason to expect the (11.a) pair to be considered most similar. There are however no studies that compare overt similarity judgments for the relevant five pairs in (11) (a-b, a-c, a-d, a-e, a-f). This gap can be filled by combining rhyming studies, studies of foreign accent perception and similarity studies for CV pairs, where C quality is systematically varied. Regarding the latter, if the study of CV sequences shows that voicing pairs (e.g. [ba]-[pa]) are more similar than nasality-based pairs (e.g. [ba]-[ma]) then we can reason that the same result will obtain a fortiori for the VC pairs, since the voicing contrast is, if anything, further attenuated in VC sequences.

4.1.2. Voicing vs. manner

Studies of imperfect rhyming provide a direct comparison between the final voicing contrast D vs. T/V_] and the manner contrasts D vs. N/V_] and D vs. G/V_]. Zwicky 1976 analyzes 236 half-rhymes in which a consonantal mismatch is ignored, out of a corpus of rock-and-roll rhymed texts with a total of 700 half-rhymes. Relevant here is that next to the 18 instances of a voicing mismatch (pairs like *died-light* or *wise-price*) there are only 5 comparable cases where nasality or obstruency differences are discounted (i.e. *mid-sin*). Hanson (1999) studies *slant* rhyme in the poetry of Robert Pinsky: here V's differ freely in rhyming pairs, while final C's stand under a violable requirement of identity. She notes that of the 128 imperfect slant rhymes in her corpus, 96% differ only in voicing (e.g. *woes-loss*). Pinsky is not an isolated case. Hanson discusses similar data in Pope and Yeats: for the latter, out of a total of 66 rhyming pairs containing a difference in the final C, 94% involve a voicing difference. No rhymes are cited where nasality or laterality is ignored.

The rhyming results are supported by the studies of similarity comparing CV sets (Walden and Montgomery 1975) or isolated C sets (van den Broecke 1975). The first of these studies identifies four dimensions of contrast: sibilant vs. non-sibilant, sonorant vs. obstruent, stop vs. non-stop and, to a much lesser extent, the p-t-k place contrast. Voicing was not a global contrast factor and the overall similarity between voicing cognates (e.g. [pa]-[ba]) emerges as much greater than that between oral/nasal or continuant/non-continuant pairs. Van den Broecke's study records Dutch subjects' impressions of similarity between single isolated C's uttered silently: here too differences based on nasality and sonority emerge as dominant. Conversely, pairs judged to possess the highest degree of similarity are pairs of similar sonority, most of them [p]-[b]-type pairs. Greenberg and Jenkins's (1965) report similar results in one of their experiments, where subjects were asked to list associates to nonsense stimuli like [klæb]. For all forms which, like [klæb], could yield a lexical item through a change of the final C's voicing, the most common responses involved voicing changes. Thus for [klæb], the most commonly mentioned forms were [klæp] (23/46 responses) and *hands* (a clear associate of *clap*: 12/46 responses). Significantly, there were other potential associates that also differ by exactly one feature from the stimulus: for [klæb] a minority associate is [klæm] (11/46). The one feature differentiating the stimulus [klæb] from the [klæm] response is nasality: apparently, however, nasality is more significant than a difference in obstruent voicing, since [klæm] was much less frequently elicited than [klæp].

This brief review indicates that voicing is, in any context, perceived as less distinctive than contrasts based on obstruency differences; and moreover that this weak voicing contrast is being suppressed - by final devoicing - in one position where it is least distinctive to begin with. This supports the proposal that devoicing is preferred to nasalization, gliding or lateralization as a means of complying with the voicing constraint (1), because devoicing induces the smallest input-output dissimilarity.

4. 1.3. Voicing vs. the C/Ø contrast

I consider next evidence on the distinctiveness of voicing as compared with the C-Ø contrast. The aim here is to suggest that dropping the C, to avoid violating (1), is a more salient departure from the input than simply devoicing it. To this end, we could note that the C-Ø contrast involves multiple dimensions of difference (as C and Ø differ in voicing, labiality, obstruency) whereas the voicing contrast involves just one of these dimensions. This relates to theories of correspondence in which feature values stand in correspondence and their mismatch is penalized by MAX and DEP constraints (Casali 1996, Lombardi 1998): in these theories, discussed below, the loss of a segment necessarily violates a superset of the constraints violated by simple featural modification.

However, I claim that the perception of similarity does not reduce to counting features. We thus look for independent support for rankings like Δ (b-p) < Δ (b-Ø). Fleischhacker (1999) solicited from English speakers relative similarity judgments between a target word and a modification of it. Some modifications involved changes of final obstruent voicing while others involved C-loss, metathesis or V insertion. The results relevant to us were of the following type:

(12) Voicing vs. C/Ø similarity differences: Fleischhacker 1999

Reference term	More similar to	Than to
print	*prind*	prin, prit

Fleishhacker also tested possible correlations between, on the one hand, greater perceived similarity between target and modified form and, on the other hand, greater preference for one modification as against another. She did this by insuring that, in some of the sets compared, the more similar form was also phonotactically *disfavored*. Thus *prind* is judged as more similar to *print* that *prin* or *prit* but it is phonotactically disfavored relative to these, both because it violates (1) and because it contains a complex coda. Despite the phonotactic improvement, the preference test correlates with the similarity test, suggesting that preference for a given modification is first and foremost a function of its similarity to the source, and only secondarily a matter of phonotactic well-formedness.

4. 1.4. Voicing vs. precedence relative to the V

I consider now the effects of precedence, bearing in mind that the voicing constraint (1) could be satisfied by displacing the C: [tæb] -> [bæt]. Data presented in Fleishhacker (1999) allows us to indirectly compare these effects to those induced by a voicing difference. In the absence of a direct comparison between voicing and serial position contrasts, one can rely on the assumption that similarity is a transitive relation. Thus if C-Ø differences are more distinctive than differences of voicing, and if serial position differences (Δ (C_1VC_2-C_2VC_1)) are more distinctive than C-Ø differences, then by transitivity, serial position is more contrastive than voicing.

Fleischhacker's (1999) study shows that precedence modifications are judged more significant than either coda or onset C deletion.

(13) Fleishhacker (1999) C/Ø contrast vs. precedence relative to V

Reference term	More similar to	Than to
flip	*fip*	*filp*
gulf	*guf*	*gluf*

Since C-Ø contrasts are more distinctive than voicing we infer that contrasts involving position relative to the V are more distinctive than voicing.

.

4.1.5. Feature transfer

Available similarity data does not bear on the possibility of single feature transfer as an alternative to final devoicing: i.e. /tæb/ -> [dæp] as against /tæb/ -> [tæp]. Here however it is safe to reason without data: whatever the dissimilarity degree of /tæb/ vs. [tæp] might be, that of /tæb/ vs. [dæp]will be greater, since two C's modify their voicing value in the case of featural metathesis, as against only one, in the case of devoicing.

This case appears irrelevant to the discussion of standard correspondence theory: single feature movement of the /tæb/-> [dæp] sort will violate twice Ident[±voice], whereas mere devoicing will violate it only once. In this case the correct preference appears to be built into the existing system. However, the variant of correspondence theory that adopts MAX [αF] constraints - instead of or in addition to Ident [±F] constraints (Casali 1996, Lombardi 1998) - will allow the [dæp] candidate to emerge as the minimal modification of the input /tæb/, under rankings like *[+voice]/_], MAX[+voice] >> Linearity [voice].

(14) Feature reversal in a MAX [αF] theory

/tæb/	*[+voice]/_]	MAX [voice]	Linearity [voice]
tæp		*!	*
☞dæp			**

The aim here is not to argue against MAX [αF] constraints, but to point out that without recourse to a theory of perceived similarity their mere existence causes candidates to emerge that need further weeding out. However, since standard correspondence theory accidentally avoids this issue, I consider in what follows only full segment reversal as an option that needs to be explicitly excluded.

4.1.6. V/ Ø vs. voicing

The last case discussed is the possibility of removing violations of (1) through epenthesis. The question involves the relative distinctiveness of the [tæb]-[tæp] contrast as against [tæb]-[tæbə]. I am discussing only one choice of epenthetic V, [ə], as any other V will likely represent an even more salient departure from the original. On this, see below.

We can directly compare for distinctiveness the voicing and Ø-[ə] contrasts on the basis of data reported by Magen (1998), who sought to determine which features of the Spanish accent are most noticeable to English speakers. Among the more common aspects of Spanish-accented English are schwa insertion (as in [əspik] for *speak* and [klozəd] for *closed*), deletion of final sibilants (as in *stand* for *stands*) and the modified realization of the voicing contrast: medial [z] realized as [s] and initial voiceless stops realized without aspiration and perceived as voiced. Magen asked her English subjects to rate for native quality the original, Spanish-accented utterances as well as edited versions of these originals, in which specific manifestations of the Spanish accent had been targeted and changed, so that the utterances would acquire native-like quality in those specific respects. In this way, one can observe how English speakers rated the V-Ø difference between the original and edited version of forms with epenthesis (e.g. Spanish-accented [klozəd] vs. modified [klozd]) and compare this with the rating difference between the original and edited version of forms with voicing changes (e.g. Spanish-accented [risən] *reason* vs. edited [rizən]). The relevant results are that voicing changes did not elicit statistically significant rating differences; in contrast, C deletion and epenthetic schwa significantly altered the ratings and in fact ranked as the most noticeable differences observed.

We reach the same conclusion about the relative salience of voicing vs. V-Ø in a different way, on the basis of Fleischhacker's (1999) study, supplemented with results of earlier work done on English and Swedish by Wingstedt and Schulman (1988). These researchers did not compare directly devoicing and epenthesis, but rather epenthesis and C deletion. Wingstedt and Schulman's subjects rated C deletion outputs as preferable to the epenthesis outputs: the relevant triplets were in this case three modifications of a base form like *conduct*: *conduc* vs. *condu[kət]* vs. *condut*.

(15) Preference judgments in Wingstedt and Schulman (1988): final C/Ø vs. Ø/V

Reference term	Best modification	Worse	Worst
conduct	*conduc*	*condu[kə]t*	condut

Fleischhacker's results allow us to compare a different version of the same question, as she inserted the V after the last C. She also distinguished similarity and preference ratings and was thus able to show that these ratings correlated.

(16) Fleischhacker (1999) similarity judgments: final C/Ø contrast vs. Ø/V

Reference term	More similar to	Than to	Than to
heft	*hef*	*heftə*	het

(17) Fleischhacker (1999) preference judgments: final C/Ø contrast vs. Ø/V

Reference term	Best modification	Worse	Worst
heft	*hef*	*heftə*	het

In this context, Wingstedt and Schulman's preference data becomes relevant to the issue of similarity. Recall that Fleischhacker had compared the effects of devoicing with those of C deletion and had verified that forms related via C deletion (*print-prin*) are perceived as more dissimilar to the base relative to forms related via voicing (*print-prind*). Reasoning again from the assumption of transitivity, it follows that devoicing will be less distinctive a departure from input than V insertion. From this I deduce that devoicing is preferred to epenthesis because it is a less salient modification of the input. As noted above, Magen's study leads to the same conclusion.

This exhausts all the alternatives to devoicing considered in (2).

4.2. The analysis

The discussion of relative similarity has yielded the dissimilarity hierarchy in (18):

(18) A hierarchy of distinctiveness in contrasts

Δ (C_1VC_2 - C_2VC_1), Δ ([ə] - [Ø]) > Δ (C - Ø) > Δ ([±son] >Δ ([±voice]/_])

Relevant to the discussion of final devoicing is only the fact that that the word final voicing contrast emerges as less distinctive than other contrasts considered. Next I show that this fact alone resolves the Too-Many-Solutions problem as applied to final devoicing.

The solution anticipated earlier is that correspondence constraints are ranked as a function of the relative distinctiveness of the contrasts they refer to. Since it is the P-map that contains information on distinctiveness, the analysis must establish a link between correspondence constraints and P-map statements. There are two aspects of this process. First, if the P-map encodes a similarity ranking between two contrasts – Δ(x-y) > Δ(z-w) – each one of the contrasts must map onto a distinct correspondence constraint, Corresp (x-y) and Corresp (z-w): otherwise there will exist relative similarity rankings that fail to be reflected in the structure of the correspondence system. I formulate this requirement below. It amounts to the claim that the dimensions and degrees of similarity differentiated by the system of correspondence are projected from the P-map.

(19) P-map projects correspondence constraints

• Let (Δ(x-y)_K_i stand for the perceptual difference between members of sound classes x and y in context K_i.

• If Δ(x-y)_K_i > Δ(w-z)/_K_j, then there exist distinct sets of correspondence conditions, Corresp. (x-y/_K_i) and Corresp (w-z/_K_j).

Corresp. (x-y/_K_i) is an abbreviated reference to constraints prohibiting the correspondence of a member of class x, in one representation, to a member of class y in a related representation, when both x and y occur in context K.

The basic requirement on the relation between the P-map and the correspondence system is that the more distinctive contrasts are protected by higher ranked correspondence conditions.

(20) Ranking correspondence constraints by relative distinctiveness

If Δ(x-y)/_K_i > Δ (w – z)/ _K_j then any correspondence constraint referring to Δ(x-y)/ _K_i outranks any parallel constraint referring to Δ(w – z)/ _K_j

The term *parallel constraints* refers to constraints that link the same pair of representations: input-to-output, and varieties of output-to-output correspondence (base-to-reduplicant, unaffixed base-to- affixed base, etc). Thus, if Δ (±nasal)/_K > Δ(±voice)/_K then (20) requires that Ident [±nasal]/_K IO >> Ident [±voice]/_K I-O. However Ident [±nasal]/_K IO may or may not outrank Ident [±voice]/_K Base Derivative, as these two constraints do not link the same pair of representations and thus are not parallel constraints.

I clarify now the reference in (20) to "any correspondence constraint referring to Δ(x-y)/_K_i". Ident [±F] constraints refer to P-map differences of the form Δ(x-y)/_K, where x and y are distinct values of the same feature. In contrast, MAX and DEP constraints refer to P-map differences of the form Δ(x-Ø)/_K, where x is a phonological property (e.g. a segment) and K is a context where x occurs in one representation and its absence occurs in a related representation. Thus MAX segment/_K, could be paraphrased as: "There is no Δ(Ø-C) between I and O, such that I contains C in K and O contains Ø in K' and K corresponds to K'."

Regarding Linearity, it is possible to view the difference Δ(AxB – ABx), where A and B are context strings and x is a segment, as the sum of the differences Δ(x-Ø)/A_B and Δ(x-Ø)/AB_. If so, then under the P-map hypothesis, Linearity refers to the same type P-map differences as MAX and DEP and thus Linearity reduces to conjunctions of context sensitive MAX and DEP constraints. This suggestion has great potential benefits as well as problematic aspects. The observation of phonological similarity judgments implicit in rhyming practices suggests that modifications of linear order induce much greater perceived differences between strings than featural mismatches or segment loss. Rhyming pairs involving metathetic mismatches like *task-fax* are virtually absent from several large corpora of half-rhymes, in clear contrast to subsequence half-rhymes like *desk-mess* (Steriade and Zhang 2001). This observation is explained if metathesis induces a superset of the Δ(x-Ø) differences caused by deletion or insertion: cumulative differences between rhyming lines create worse, hence less frequent half rhymes. A related benefit of this proposal is its ability to explain the extreme rarity of metathesis as a solution to phonotactic violation[9]. The possible drawback is that, by the same token, if Linearity reduces to MAX and DEP conjunctions and all MAX and DEP constraints are context sensitive, it becomes difficult to characterize any systems that prefer metathesis over deletion as a solution to phonotactic violation. In what follows I continue to refer to metathetic differences as Δ(AxB-ABx) and to Linearity as an independent class of constraints. The problem noted is left unresolved here.

Returning now to the analysis of final devoicing, we observe that from the principle in (19) and the distinctiveness hierarchy in (18) it follows that each contrast distinguished by (18) gives rise to a distinct set of correspondence conditions. From (20) it follows that correspondence conditions extracted from (18) are ranked by distinctiveness as below:

(21) Ranking of I-O correspondence constraints by the distinctiveness scale (18)

Linearity (C_1VC_2 vs. C_2VC_1), DEP (ə vs. Ø)/C_] >> MAX (C vs. Ø)/V_] >>

Ident [±son]/V_] >> Ident [±voice]/ V_]

Recall now that the correspondence constraints in (21) are the only ones whose violation could in principle satisfy the *[+VOICE]/_] constraint, for inputs that violate it. Our starting point was the observation that each one of the five constraints in (21) can, in the present version of correspondence theory, be ranked lower than the others, thus predicting at least five distinct solutions to violations of (1). The constant ranking in (21) – derived from the constant similarity relations in (18) - eliminates this difficulty. Although, there are in principle six different ways of ranking *[+VOICE]/_] relative to members of the correspondence hierarchy in (21), only two sets will yield distinct effects: one set contains the five rankings in which *[+VOICE]/_] >> Ident [±voice]/ V_], all of which amount to final devoicing; the other set contains the ranking in which Ident [±voice]/ V_] >> *[+VOICE]/_]. Only two distinct outcomes are thus predicted: violate the phonotactic or apply final devoicing. This is the result we were aiming to derive.

4.3. Consider the alternatives

The next step is to consider if alternative views on correspondence and constraint interactions, unaided by the P-map, can achieve the desired result of cutting down appropriately on the number of solutions to phonotactic constraint violation.

Two of the solutions listed in (2), nasalizing the voiced stop or leniting it to an approximant, will have the effect of changing not one feature but possibly two: in both cases an oral stop ([-son, -nasal, -cont]) becomes a sonorant (either [+son, +nas, -cont] or [+son, -nas, +cont]). One might hope to discover a formal solution to the Too-Many-Solutions problem, by noting that a one-feature modification (i.e. violating once an Ident F constraint) is better than a two-feature modification (violating two Ident F constraints). It is not clear how this idea can be implemented, as rankings of the form Ident F >> Ident G, Ident H cannot be ruled out in principle. However there is independent reason to believe that the cause of our problem does not reside in the count of features being modified. This can be shown by observing that in languages like Turkish (Inkelas and Orgun 1995) where stops – not fricatives - are subject to final devoicing, the active constraint must be (22).

(22) The Turkish version of *[+VOICE]/_]: *[+voice,-son, -cont]/_]

This constraint can be, in principle satisfied by turning voiced stops into fricatives to avoid devoicing. But Turkish reacts to violations of (22) exactly as Russian or Dutch react to violations of (1): by final devoicing. Underlying forms like /kitab/ are devoiced ([kitap]), not lenited (*[kitaβ][10] or *[kitav]). The real generalization is that stricture contrasts are not being sacrificed when the phonotactic problem at hand is readily solved by voicing adjustments. As discussed above, changes in stricture induce greater perceived differences than changes of voicing.

The same point arises in connection with place phonotactics. Certain heterorganic obstruent clusters – among them *tp, dp, tk, dk*– are frequently disfavored or impermissible, as in Korean, Ancient Greek or Classical Latin. Consider now ill-formed /dk/ inputs (e.g. Latin *ad-kelera:re*, surface [ak:elera:re]). Such inputs lend themselves to multiple fixes: [ak:elera:re] vs. *[askelera:re], *[alkelerare] etc. One of these solutions is widely attested – gemination or place assimilation – while the others are simply unheard of. The mapping /adkelerare/ -> [ak:elera:re] entails the loss of two features (voicing and coronality): it is unclear what feature counting alternative to the P-map will establish this mapping as preferable to /adkelerare/ -> *[alkelera:re] mapping, in which obstruency is lost and laterality is added. The generalization here is that when the same phonotactic problem can be addressed by adjusting either place or stricture, the solution is to change place.

Consider now the viability of C-deletion as a solution to final voicing violations. Recall that underlying forms like /tæb/ could – but never do – satisfy (1) by dropping the final voiced stop altogether. Here too one might imagine that a different modification of the theory of correspondence, one that substitutes MAX [αF] for Ident F constraints (Casali 1996, Lombardi 1998), will explain the preference for devoicing. Thus the output of devoicing, [tæp], violates only MAX [+voice], while the output of C-deletion, [tæ], violates MAX [+voice], plus MAX [labial], MAX [-cont], MAX [-nasal] etc. On this view, the C deleting candidate loses under any ranking of the MAX constraints. But this cannot be the answer either. Consider the constraint against stop+non-coronal stop sequences (*tp, tk, dp, dk*) active in Ancient Greek. Most such inputs arise at the boundary between the verb root and the perfect ending –*ka* and the constraint is satisfied in this case through t/d deletion: e.g. *ke-komid-ka* -> [kekomika] 'I have eaten'. Thus the Greek solution to the *tp, tk* problem is not to place-assimilate, as in Latin or Korean, but rather to drop the first stop altogether, i.e. to violate MAX [α voice], MAX [Coronal], MAX [-cont] etc. If we look at this problem in terms of the *number* of features being sacrificed from the input, we cannot understand why the [d] of *komid-* had to drop, when it could well have been turned into [l], [r], [s] yielding well-formed *[kekomilka], *[kekomirka], or *[kekomiska]. Each one of these alternatives contains fewer MAX F violations than the solution actually adopted, which was to eliminate the [d] altogether.

(23) Failed attempt at C-deletion in system with MAX F and no Ident F

/ke-komid-ka/	MAX [-cont]	MAX coronal
kekomika	*	*!
☞kekomilka	*	

A theory that relies exclusively on MAX[αF] cannot explain any pattern in which a segment is deleted in contexts where the phonotactic violation can be met by modifying a subset of its features. If we adopt, along with the MAX [αF] constraints, DEP[αF] constraints, then the ranking DEP [+strident], DEP [+nasal], DEP [+continuant] >> MAX [-son], [-cont] can induce t/d deletion. This move however brings back the original problem: in a system where every feature value possesses its DEP constraint, devoicing violates DEP [-voice]. Then what rules out DEP [-voice] >> MAX [-son], [-cont]? What emerges from this discussion is that some hierarchy of features must be assumed in any approach: one must recognize that modifications of voicing, especially final voicing, matter less than modifications of obstruency. This is the first step on the way to the P-map.

For the P-map analysis, the Greek t/d deletion process raises not the formal problem faced by MAX F analyses but an empirical question: is the contrast between unreleased *t/d* and Ø in pre-stop position judged less distinctive than that between *t* and *s* or *t* and *n* or *t* and *l* in the same position? If yes, then the we predict that, in such a context, straight deletion is more likely than fricativization or lenition to a sonorant. We lack similarity data bearing on this: but studies of assimilation and cluster simplification (Wilson 2001, Steriade 1999b) suggest that a perceptibility based solution will be fruitful for this case too.

5.1. Size-of-cluster constraints

Many languages constrain agglomerations of C's, when they exceed some specified size. If the constraint responsible for size-of-cluster phenomena prohibits strings of the form C_iC_j/_K, where K specifies a context, segmental or prosodic, then a representation violating it can achieve compliance in at least three ways: by deleting C_i, by deleting C_j, by modifying either of them or by adding a V, the insertion of which will yield further choices regarding site and V quality.

In this section I briefly suggest that this wealth of apparent choices in dealing with size-of-cluster constraints fails to reflect phonological reality: the actual solution comes much closer to being pre-determined by the composition of the string containing the violation. While the choice between V insertion and C deletion might remain free in resolving a size-of-cluster violation, other decisions (which C to delete; which C to modify, and how; where to insert a V and which V to insert) are partly or fully predictable. They are predictable largely in terms of the relative similarity between the input and the modified output: it is the most similar input-output pair that is predominantly selected.

The issue of predictability in intervocalic CC cluster simplification has been independently identified by Wilson (2001), whose formal proposal differs from ours but whose discussion raises points related to those made here. The partial predictability of epenthesis site in initial clusters is analyzed in a framework akin to the P-map by Fleishhacker (1999). Here I extend Wilson's observations by considering briefly the choice of C's to delete in sequences more complex than VCCV.

5.2 Similarity with Ø in cluster simplification

Wilson (2001) notes that when C-deletion targets a sequence VC_iC_jV, the lesser perceptibility of C_i leads to its loss. I consider here the choice between deletion of C_i and C_j in $VC_iC_jC_kV$. To simplify matters, I assume, with Wilson, that the prevocalic C_k is undeletable here. My basic empirical point is that the target of deletion is predictable from considerations of confusability, not from its prosodic position or its adjacency to the V.

The P-map's predictions for cluster simplification in $VC_iC_jC_kV$ are derived from the degree of distinctiveness of two contrasts: C_i vs. Ø/V_C, and C_j vs. Ø/C_C. To be judged sufficiently distinct from Ø in a given context, the sound must in fact be sufficiently distinct from *both* of the elements adjacent to it. To see this suppose that C_j in $VC_iC_jC_kV$ is confusable with C_i: the percept resulting from this confusion is $VC_iC_iC_kV$ or $VC_i{:}C_kV$. The effect of shortening-in-clusters (Haggard 1973, Klatt 1973) renders $VC_i{:}C_kV$ confusable, in turn, with the simplified VC_iC_kV. Thus C_j's similarity to C_i leads, under C_i-C_j adjacency, to C_j's similarity to Ø. The same holds if C_j is similar to C_k. Likewise, postvocalic C_i in $VC_iC_jC_kV$ is similar to Ø, if it is too similar with either the preceding V (a confusion leading to the $V{:}C_iC_kV$ percept) or to the following C_j. Finally consider a sequence C_jC_kV

in utterance initial position: the initial C_j is confusable with Ø if it is confusable with either the absence of sound that precedes it or the C_k that follows. Mutatis mutandis, the same holds for confusion with Ø of C_k in an utterance final VC_jC_k. Similarity with Ø means then similarity with either one of the adjacent elements, whether silence or sounds.

From this we predict that position relative to the syllable boundary or a V will not guarantee that C_i is less confusable with Ø than C_j. The simplest example illustrating this is the loss of postvocalic liquids in systems where other postvocalic C's are preserved. In such cases what identifies the element deleted is not proximity to the V or syllable position but *similarity* to the neighboring V. Cho (2000) presents an example of this type in his analysis of postvocalic l deletion in Korean VlCCV.

A more complex illustration of the same point is the difference between the confusability with Ø of interconsonantal stridents and stops. Consider first the case in which the sequence $VC_iC_jC_kV$ contains three stops. As a stop, C_i is sufficiently distinguishable from the immediately preceding V. Moreover, since this V carries C_i 's transitional cues to place and voicing, it provides information distinguishing C_i from other C, *including* C_j, which might have occurred in the same position. Therefore C_i is not confusable with either the V or the following C_j, hence it is not confusable with Ø. The medial C_j, on the other hand, is confusable with Ø: as no V is adjacent to C_j, the string $VC_iC_jC_kV$ contains less information allowing the listener to differentiate C_j from any other stop that might have occurred in the $VC_i\_C_kV$ position, *including from* C_i *or* C_k. If C_j is confusable with either one of the adjacent stops, then the string $VC_iC_jC_kV$ is confusable with $VC_i{:}C_kV$ or $VC_iC_k{:}V$ and hence with the shortened variant VC_iC_kV. And therefore C_j is more confusable with Ø, than C_i. We reach in this way the conclusion that, if cluster simplification targets the C that is most confusable with Ø, then it will operate in this case at the expense of the medial C_j. However, this holds only for cases in which all but position relative to the V is equal between the three C's. Suppose that C_j in our $VC_iC_jC_kV$ string is a strident: then the distinctiveness difference between Δ(C_i vs. Ø)/V_C, and Δ(C_j vs. Ø)/C_C might be obliterated or reversed, as the inherent noisiness of the strident C_j identifies it as distinct from any adjacent non-stridents, even in the absence of vocalic transitions. If so, the relevant correspondence constraints (MAX C/V_C and MAX strident/C_C) either remain unranked or MAX strident/C_C might in fact rank higher.

5.3. Test cases

Several predictions follow from this: first, a language may delete interconsonantal stops but not interconsonantal stridents. This pattern occurs in Dihovo Macedonian (Groen 1977). To analyze it we need a size-of-cluster constraint, interpreted as the requirement that each C be adjacent to a V:

(24) C//V: Every C is adjacent to an V.

The Dihovo pattern of cluster simplification corresponds to (25):

(25) Stops, not stridents, are deleted between stops in VCCCV

MAX [-cont] /V_C, MAX strident/C_C >> C//V >> MAX[-cont] /C_C

The clear effect of the P-map in this case is the ranking MAX stop /V_C, MAX strident/C_C >> MAX stop/C_C. The position of the phonotactic C//V relative to the correspondence constraints is left undetermined by the P-map and this allows us to predict variation in the patterns of cluster simplification. Thus the modified hierarchy in (26), where C//V has climbed higher, requires that *some* cluster simplification take place even in V-stop-strident-stop-V clusters.

(26) All VCCCV clusters reduced to VCCV.

C//V >> MAX [-cont] /V_C, MAX strident /C_C >> MAX[-cont] /C_C

Further elaboration of (26) yields two types of simplification for $VC_iC_jC_kV$ sequences where C_i is a stop and C_j a sibilant: either $VC_iC_jC_kV$ -> VC_iC_kV or $VC_iC_jC_kV$ -> VC_jC_kV. What is however invariant is that, if the middle C_j is a stop surrounded by obstruents, it will always be deleted. Colloquial Latin illustrates the more revealing pattern: inter-obstruent stops are lost, whereas inter-obstruent [s] is preserved at the expense of the stop preceding it.

(27) Two types of cluster simplification in Latin (Niedermann 1952)

VC_i stop C_k (V) -> V C_i C_k(V)	VC_i s C_k (V) -> V s C_k(V)
pa:sktus -> pa:stus	sekstus -> sestus
nokts -> noks	obstendo -> ostendo
temptare -> tentare	apsporto -> asporto
lampterna -> lanterna	sekskenti: -> seskenti
k^winktus -> k^wintus	pinstus -> pi:stus

This cluster reduction pattern suggests that the strident-Ø contrast is more distinctive, even in the absence of contextual cues, than the postvocalic stop-Ø contrast.

(28) Simplified cluster reduction hierarchy for Latin[11]

C//V >> MAX strident/C_C >> MAX [-cont] /V_C >> MAX[-cont] /C_C

(29) Cluster reduction in *obstendo*

/obstendo/	MAX strident/C_C	MAX[-cont]/V_C
optendo	*!	
☞ostendo		*

(30) Cluster reduction in *kwinktus*

/k^winktus/	C//V	MAX[-cont]/V_C	MAX[-cont]/C_C
k^wiktus		*!	
☞k^wintus			*

The Latin asymmetry between postvocalic stops and sibilants as targets of cluster simplification is encountered in several languages: among them Finnish, Catalan (Wheeler 1979) and colloquial Polish (Madejowa 1992). The alternative pattern of deletion, where every interconsonantal obstruent deletes, whether it is a stop or a sibilant, is perhaps also attested, in Greek, Sanskrit (Steriade 1982) and Korean (Kim-Renaud 1974) but alternative interpretations are available for these cases. In particular the analysis of cases like Korean *kaps-to* 'price-and' -> [kapto] must take into account the fact that all fricatives are prevocalic in Korean: if [p] had deleted, the actual outcome would have to be *[katto], not *[kasto]. The tableau below indicates that the Latin ranking of correspondence constraints need not be changed to derive this case.

(31) Cluster reduction in *kapsto*

/kapsto/	*s/_C	MAXsibilant/C_C	MAXstop/V_C
kasto	*!		*
katto		*	*!

☞kapto		*	

The Greek and Sanskrit instances of deleted inter-obstruent [s] are sparsely attested and involve exclusively suffixal [s]. It is possible then that the Latin reduction pattern represents the general case.

My general claim however is more modest. A P-map account, as sketched here, predicts only this: insofar as a C-cluster contains one and only one C whose similarity with Ø is greater than that of the other cluster members, cluster reduction will target this one C. Similarity with Ø means similarity with an adjacent element. We have seen that an inter-obstruent stop – or a stop flanked by a nasal and an obstruent – can be identified as the most confusable with Ø among all components of its C-cluster. This corresponds to the observation that stops in such contexts are the systematic, invariant targets in cluster simplification. It may also turn out that the inherent salience of stridents renders the Δ(strident-Ø)/C_C contrast more distinctive than Δ (stop-Ø)/V_C. If so, a stronger prediction is made: the postvocalic stop will always be deleted, unless morphological factors intervene, in V-stop-s-C sequences.

6. Insertion and the ranking of DEP constraints

The P-map account of the choice of epenthetic segments derives from the hypothesis of a context-dependent hierarchy of similarity between individual segments and Ø. If a phonotactic constraint requires insertion of a segment in some context K, then the segment most confusable with Ø in K is predicted to be the choice of insertion. I outline now how this prediction follows from proposals made thus far. The class of correspondence constraints violated by insertion take the form in (32):

(32) DEP (I-O) schema

There is no Δ(Ø-x) between I and O, such that I contains Ø in K and O contains x in K' and K corresponds to K'.

Like all correspondence constraints, the DEP constraints are projected from the P-map: this means that if $\Delta(\text{Ø-x})/\_K_i > \Delta(\text{Ø-y})/\_\_K_j$, then corresponding to this similarity ranking there exists a ranking of correspondence constraints so that the more confusable contrast with Ø corresponds to the lower ranked DEP constraint (cf. (20)). It follows that the outcome of phonotactically motivated insertion is to a large extent predetermined. This prediction is mitigated by the effect of conflicting phonotactics and thus the effect of markedness on epenthesis has to be considered as well. I outline next only one aspect of the evidence bearing on this point: the selection of epenthetic segment *quality*.

6.1. Epenthetic glottals

The typology of epenthetic C's has been usefully outlined by Lombardi (1997), who identifies a general pattern, insertion of [ʔ], and minor deviations from it, due either to structure preservation (in the form of constraints forbidding [ʔ]), morphological constraints, or the dispreference for rhyme [ʔ]. The central question is what accounts for the preference for inserted [ʔ]. Lombardi assumes that the relevant factor is markedness: [ʔ] is the least marked among C's. But what fact other than its propensity to get inserted reflects [ʔ]'s extreme unmarkedness? This is a harder question: the standard evidence for markedness, the implicational universals, suggest otherwise: [ʔ]'s presence in an inventory is not asymmetrically implied by the presence of all other C's, or indeed by the presence of all other members of its stricture class. Under the P-map hypothesis, [ʔ] has, with [h], a uniquely favorable property for an *epenthetic* C: it does not possess an oral constriction and thus it will fail to induce coarticulatory changes on neighboring segments, unlike the orally articulated C's. If we compare input-output pairs of the form V(input)-CV(output) the most similar ones will be V-ʔV or V-hV or V-GV, where G is homorganic to V. Both epenthesis of [h] and epenthesis of homorganic glides represent in fact the only widely attested epenthesis patterns, along with the more common case of [ʔ] insertion. Thus, if the lack of coarticulatory V modification translates into the similarity hierarchy in (33), then the P-map hypothesis predicts the preference for [ʔ] as epenthetic segment, regardless of how it rates in markedness.

(33) Δ(Ø-t)//V; Δ(Ø-k)//V; Δ(Ø-p)//V > Δ(Ø-ʔ)//V

The view that markedness determines the choice of epenthetic segments runs into independent difficulties. Several of the languages Lombardi cites, where [ʔ] occurs as an exclusively epenthetic C, must in fact be assumed to rate the markedness of [ʔ] as higher than that of all their other C's. Thus to explain the fact that German allows only epenthetic [ʔ] we have two options. We can assume that a constraint *[V >>*[ʔ] >> MAX C while at the same time assuming MAX C >> *[p], *[k] and *[t]. But this contradicts the universal markedness ranking Lombardi assumes: *[p], *[k] >> *[t] >> *[ʔ]. The alternative is to rely on the P-map-based ranking: MAX/DEP [p], [t], [k] /_V >> MAX/DEP [ʔ]/_V. The German choice of epenthetic [ʔ] follows when this correspondence ranking is embedded in the complete analysis: MAX [V, MAX/DEP [p], [t], [k] /_V >> *C >> *[V >> MAX/DEP

[ʔ]/_V. On this view, it is correspondence, not markedness, that differentiates C types. The same conclusion follows from Lombardi's analysis of Asheninca, one of the rare languages where something other than a laryngeal or a glide is inserted in hiatus. Lombardi argues that [t] is inserted in Asheninca because *[ʔ] is undominated. If we grant this, it follows that *[ʔ]>>*[t], since *[t], but not *[ʔ], is outranked by MAX C. This too is incompatible with the claim of unmarked status for [ʔ]. I conclude that there is either no constant context-free, all-purpose preference for glottal as against other stops, or, if there is a preference, it is the opposite from the one needed to predict the proper choice of epenthetic C. The choice is predicted by the P-map.

6.2. Epenthetic schwa

I document next the status of [ə] as the inserted V of choice. This too follows from a hierarchy of similarity with Ø. What defines for our purposes the class of schwa-vowels is not their mid central quality (Romanian and English [ʌ], for instance, are not epenthetic, while Romanian [ɨ] is) but rather the fact that the schwa-like V is significantly shorter and more variable in quality than all other V's in an inventory. This characterization allows for some diversity in the actual quality of a language's neutral V, while permitting us to make specific predictions about what differentiates it from other V's. A systematic difference of duration between schwa and other V's of Dutch is documented in Koopmans-van Beinum 1994; and known informally to obtain for English and French schwa. Further, Dutch schwa is also more variable in its F2 values than other Dutch V's (Koopmans-van Beinum 1994; van Bergem 1995).

Assuming then that the defining properties of schwa are shortness and variability, the preference for schwa as an epenthetic element follows from the fact that it is, in both duration and relative absence of invariant articulatory properties, the closest thing in a V system to no segment at all, i.e. to Ø. Note that this is not the same as saying that schwa has no properties. First, it *is* a vowel. When it does occur, speakers count an extra syllable. This is invariant. Further, schwa in Dutch is less durationally variable than other V's (Koopmans-van Beinum 1994): it is least subject to contextual or context-free lengthening. In many languages where schwa is unstressable, as in Dutch, Indonesian, and French, this can be attributed to the fact that schwa cannot be lengthened. In that respect then it *does* have a second invariant property: short duration. Thus we cannot explain why schwa is preferentially inserted by assuming that it possesses no properties, or a subset of the invariant properties of other vowels. Schwa epenthesis is preferred because the P-map identifies it as the closest thing to no epenthesis at all.

The preference for schwa insertion may manifest itself in a language independently of the composition of the V inventory. However, when structure preservation does not constrain its occurrence, i.e. in languages possessing contrastive or non-alternating schwas, this preference for inserting schwa is absolute: the statement in (34) holds of all relevant cases I have encountered, some of which are listed below.

(34) If a language contrasts schwa and zero in some context, or if it contains non-alternating forms with schwa, and if it resolves clusters through epenthesis, then the choice of productive epenthetic vowel is limited to schwa.

[Indonesian (Adisamito 1993), Romanian (Avram 1990 and below), German (Giegerich 1987), Damascene Arabic (Bohas 1986), French (Dell 1978 and below), Meitei (Chelliah

1997), Miya (Schuh 1996), Welsh, English, Dutch (Booij 1995; Kuijpers, Donselaar, Cutler 1997), Berber (Kossman 1995, MacBride 1999)]

The comments about the markedness of [ʔ] apply here too: schwa insertion does not stem from a context-free preference for this segment. It most clearly does not in languages where schwa is permissible only where a V is otherwise needed. In Berber, for instance, schwa – but no other V – must be prevented from occurring in open syllables (MacBride 1999); in Miya, it cannot occur after a sonorant (Schuh 1996). Phenomena of this sort require *[ə] constraints, whose high ranking is not paralleled by other *V conditions: this precludes a claim of unmarked status for schwa. What explains the V-epenthesis generalization is the existence of a hierarchy of DEP(V) constraints containing, at its bottom, DEP[ə]. Here too I speculate that the source of this DEP(V) hierarchy are speakers' judgments of relative similarity between individual V's and Ø.

7. All-purpose segments

I turn next on a different respect in which the P-map proposal tightens the theory of correspondence. The suggestion here will be that the segments most likely to be inserted are also most likely to be deleted. The behavior of [ə] will illustrate this, but reports about specific segments being both preferentially inserted and deleted go beyond this case[12]. This phenomenon is not predicted by the classic theory of correspondence but it follows from the same principle (20) that solved the Too-Many-Solutions puzzle involving final devoicing. (20) predicts that if Δ(x-Ø)/_K > Δ(y-Ø)/_K then not only will DEP(x)/_K >> DEP(y)/_K but also MAX(x)/_K >> MAX(y)/_K. Then y has priority for both insertion and deletion over x. More concretely, this means that if some y-Ø contrast is identified as more confusable than other segment-Ø contrasts, the insertion and deletion of y will be the preferred response to all phonotactics for which insertion or deletion of segments from y's class represent a potential solution.

The general class of situations described is this: a language avoids hiatus, hence it must delete a V when adjacent to others or insert a C between them. The same language avoids CCC clusters, hence it must insert some V in such clusters, or else delete a C. As it appears impossible to predict the preference between C insertion and V deletion, or between C deletion or V insertion, these choices are settled on a language-specific basis. The language we are interested in eliminates clusters by V insertion and hiatus by V deletion. Under these conditions, the P-map hypothesis predicts that the V deleted in hiatus is the same as the V inserted as a cluster resolution strategy. That is because the criterion that selects a V for deletion is the same as the one selecting it for insertion: this criterion is the greater similarity of the contrast between that V and Ø relative to all other V's. I summarize this below, using [ə] as the V judged to be most confusable with Ø.

(35) If, for any choice of V≠ə, Δ(V-Ø) > Δ(ə-Ø)

Then MAX (V), DEP (V) >> MAX(ə), DEP (ə)

Note that, aside from the P-map, nothing guarantees that the MAX and DEP constraints corresponding to different V's will be ranked as pairs. Nothing excludes rankings like (36), which predict that schwa is deleted but that [a] is inserted.

(36) Schwa deleted and [a] inserted in system lacking P-map, where MAX/DEP constraints exist for individual segments:

MAX (a), DEP (ə) >> Phono-constraints >> MAX (ə), DEP (a)

There exist systems in which hiatus is resolved at the expense of certain V's only (Pulleyblank 1988, Casali 1966; below). As Casali argues, these systems reflect a hierarchy of distinct MAX constraints, independent of the markedness scale: deletion targets the V associated with the lowest ranked MAX V. This dispenses with alternatives to the P-map analysis which assume that targets of deletion/insertion are determined by markedness conditions alone, interacting with a monolithic MAX V, DEP V.

One can document systems where both V insertion takes place and specific V's are deleted, either to avoid hiatus or to shorten the word: in all such cases, it is the prediction of the P-map analysis, (35), that is upheld. The pattern emerges more clearly if we restrict our attention to productive, lexically unrestricted insertion and deletion. Thus in French, it is schwa that deletes in hiatus, regardless of its location relative to the other V. Schwa is also the vowel inserted, optionally, in C_CC clusters:

(37) French schwa deletion

(a) Optionally deleted in VC_CV contexts: **no other V deletes**.

la pelouse [lapluz] 'the lawn' cf. phrase initial *pelouse* [pəluz]

pas de role [padʁol] *'no role'* cf. phrase initial *de role* [dəʁol]

(b) Obligatorily deleted in hiatus; **no other V deletes**.

t'entendre 'to hear you' cf. *te remercier* [təʁmɛʁsje] 'to thank you'

vivre ailleurs [vivʁ ajœʁ] 'live elsewhere', cf. *vivre là* [vivʁə la] 'live there'

Compare *ni entendre* [ni ɑ̃tɑ̃dʁə]'neither to hear'; *vivra ailleurs* [vivʁa ajœʁ] 'will live elsewhere', all of which surface with hiatus, in the absence of a deletable vowel. Schwa is also inserted, optionally, to avoid clusters of obstruents.

Romanian [ɨ] gives rise to identical patterns:

(38) Romanian schwa:

(a) [ə] optionally deleted next to a V: **no other V deletes**

vine 'ndatʌ 'comes immediately' cf. *ɨndatʌ* 'immediately'

vine 'nainte 'comes before' cf. *ɨnainte* 'before'

(b) Non-deleting V's:

vine odatʌ 'comes once'; *vine ʌla 'comes that one-masc.'*

vine aja 'comes that one-fem.'

(c) [ə] optionally inserted in obstruent clusters CCC(C): **no other V is inserted**

opt-spre-zece [optsprezetʃe] ~ [optɨsprezetʃe] '18' ('eight-to-ten')

Comparable patterns appear in Dutch(Booij 1995; Kuijpers, Donselaar, Cutler 1997), Meithei (Chelliah 1997) and English (Bybee 1978).

8. Conclusion

I note in closing potential points of dispute relating to the P-map.

At the most basic level one can dispute the premise this account shares with most modern phonology, namely that phonology is a problem-solving system, or – as Goldsmith (1993) puts it – "an intelligent system". If the phonotactic in (1) is not viewed as a problem to be solved, or as a standard of well-formedness independent of the lexicon's contents, but rather as a generalization over the words that happen to be attested in one's language, then no Too-Many-Solutions problem arises: learners, on this view, do not seek to find a solution to (1) but to learn whatever patterns happen to be instantiated by their lexicon.

Similarly, one may question whether the Too-Many-Solutions problem arises in the initiation of sound change. The view presented here is that innovators may aim to improve a sound system and that they do so in the safe regions of confusability identified by the P-map: we assume, for instance, that speakers who initiate final devoicing have a choice of methods to satisfy (1) – or a choice of spontaneously occurring speech variants to promote (Lindblom et al. 1995)- and choose final devoicing because it is the least departure from the established speech norms. But it may be possible to look at the initiation of sound change in different terms if most naturally occurring variants to an established lexical form represent its common misperceptions. In that case, innovators have the more passive role of simply favoring the *more commonly noted* deviations from the norm, without reflecting on their phonotactic virtues or on their similarity to canonical forms. This possibility has been discounted here on the strength of evidence that speakers know not only what are the more common deviations from the norm but also which deviations are more similar to the norm. We have seen that knowledge of similarity is displayed in rhyming practices and experiments seeking overt similarity judgments. It remains to extablish however that the available knowledge of similarity is exploited by speakers in constructing grammars and in favoring one change over others. (See also Steriade 2001.)

A different class of possible objections to the P-map involves the fact that there is, at least at first sight, a considerably greater variety of alternations than a theory of perceived sound similarity may predict. If the case for predictability of C-deletion or V-insertion appears overstated, a way of testing the P-map proposals is to focus on fully productive, not yet lexically entrenched processes. For the moment, it seems necessary only to acknowledge the existence of parochial constraints governing alternations, in addition to phonotactics and P-map generated correspondence constraints.

Finally, I have focussed here on aspects of perceived similarity that correspond to broad cross-linguistic generalizations: and for this reason it may appear that a claim of universality is made regarding the contents of the P-map. This is not the intention. If the perception of similarity is governed, in part, by "the contents of the universe of discourse" (Tversky, cited in Frisch, Broe, Pierrehumbert 1997), then the same pairs of sounds will rate differently for similarity, when embedded in different systems. The existence of such effects is not denied: the development of a first-approximation version of the P-map will hopefully allow one to identify them.

References

Adisamito, Niken 1993 "Syllable structure and the nature of schwa in Indonesian," Studies in Linguistic Sciences, 23, 1-19, University of Illinois, Urbana.

Avram, Andrei 1990 "Sur les voyelles neutres, en roumain, en albanais et dans les langues romanes occidentales," Revue Roumaine de Linguistique, 35, 19-27

Blevins, Juliette and Andrew Garrett 2002 "The Evolution of Metathesis," to appear in Bruce Hayes, Robert Kirchner and Donca Steriade (eds.) *The Phonetic Bases of Phonological Markedness*, Cambridge University Press

Bohas, G. 1986 "Accentuation et effacement de e dans le parler du Damas," Revue Québecoise de Linguistique, 16, 39-60

Booij, G. 1995 *The Phonology of Dutch*, Oxford University Press

Calabrese, A. 1995 "A Constraint-Based Theory of Phonological Markedness and Simplification Procedures," *Linguistic Inquiry*, 26, 373-463

Casali, R. 1997 "Vowel elision in hiatus: which vowel goes?" *Language* 73, 493-533.

Chelliah, S. 1997 *A Gramar of Meithei*, Mouton de Gruyter

Cho, T. 2000 "Intra-dialectal variation in Korean consonant cluster simplification: A stochastic approach. *Chicago Linguistics Society 35* (CLS 35).

Connine, C. D. Titone, T. Deelman, Dawn Blasko "Similarity Mapping in Spoken Word Recognition,"*Journal of Memory and Language*, Vol. 37, No. 4, Nov 1997, pp. 463-480

Dell, F. 1978 "Epenthese et effacement de schwa dans les syllabes contingues en Français," in B. de Cornulier & F.Dell (eds.) *Etudes de Phonologie Française*, 75-81

Fleischhacker, Heidi 1999 "The Location of Epenthetic Vowels with Respect to Consonant Clusters: An Auditory Similarity Account" UCLA MA thesis

Flemming, E. 1995 *Perceptual Features in Phonology*, UCLA Ph.D. dissertation

Frisch S., M. Broe and J.Pierrehumbert 1997 “Similarity and Phonotactics in Arabic” *Rutgers Optimality Archive 223*

Giegerich, H. 1987 “Zur schwa-Epenthese im Standard Deutchen,” *Linguistische Berichte* 112, 449-469

Goldsmith, J. 1993 *The Last Phonological Rule*, University of Chicago Press

Greenberg J. and J.J Jenkins 1964 “Studies in the Psychological Correlates of the Sound System of American English,” *Word* 20, 157-177

Groen, B. M. 1977 *A structural description of the Macedonian dialect of Dihovo : phonology, morphology, texts, lexicon*, Lisse: The Peter de Ridder Press

Haggard, M. 1973 “Abbreviation of consonants in English pre- and post-vocalic clusters” *Journal of Phonetics*, 1, 9-4.

Hansen, K. (this volume) “Formal Variation in the Rhymes of Robert Pinsky’s *The Inferno of Dante*,” Inkelas, S. and C.O. Orgun 1995 "Level ordering and economy in the lexical phonology of Turkish" *Language* 71:4.

Ito, J. 1986 *Syllable Theory in Prosodic Phonology*, University of Massachusetts, Amherst Ph.D. Dissertation. New York, Garland Press 1988

Jun. J. 1995 *A Constraint Based Analysis of Place Assimilation Typology*, UCLA Ph.D.Dissertation.

Kisseberth, C. 1970 "On the functional unity of phonological rules," *Linguistic Inquiry* 1, 291-306

Klatt, D. 1973 "Durational characteristics of pre-stressed word initial C clusters in English," MIT RLE Quarterly Progress Report, 108, 253-260

Koopmans-van Beinum, Florien 1994 "What's in a Schwa?" *Phonetica* 51:68-79

Kossman, Marten Gossling 1995 "Schwa en berbere," Journal of African Languages and Linguistics, 16, 71-72

Kuijpers, Cecile, Wilma van Donselaar and Anne Cutler 1997 "Facilitatory effects of vowel epenthesis on word procesing in Dutch," *Jounal of Memory and Language*

Lindblom, Bjorn, Susan Guion, Susan Hura, Seung-Jae Moon and Raquel Willerman 1995 “Is sound change adaptive?” *Rivista di Linguistica* 7, 5-37.

Lombardi, L. 1997 “Coronal epenthesis,” *Rutgers Optimality Archive* 247

Lombardi, L. 1998 “Evidence for MAX feature constraints from Japanese*,” Rutgers Optimality Archive* 247

Lombardi, L. 1999 “Restrictions on the direction of voicing assimilation, direction in assimilation*” Rutgers Optimality Archive* 246

Madejowa, M. 1992 “Zasady wspolczesnej wymowy polskiej oraz ich przydatnosc w praktyce szkolnej”, *Jezyk Polski*, 72, 187-98.

MacBride, A. 1999 “Phonological Issues in the Distribution of Berber [ə],” UCLA MA.

Magen, H. 1998 “The perception of foreign-accented speech” *Journal of Phonetics*, 26, 381-400

McCarthy, J 1993 "A Case of Surface Constraint Violation," in Carole Paradis and Darlene LaCharité, eds., *Constraint-Based Theories in Multilinear Phonology*, special issue of *Canadian Journal of Linguistics* 38, 169–195.

McCarthy, J. and A. Prince 1995 "Faithfulness and Reduplicative Identity," *University of Massachusetts Occasional Papers in Linguistics* 18: Papers in Optimality Theory. Ed. by J. Beckman, S. Urbanczyk and L. Walsh Dickey. Pp. 249–384.

Miller, G.A. and P.E. Nicely 1955 "An analysis of perceptual confusions among some English consonants". *Journal of the Acoustical Society of America*, 27, 338-352

Niedermann, Max 1953 *Précis de phonétique historique du Latin*. Paris, C. Klincksieck.

Paradis, Carole 1988 "On constraints and repair strategies" *Linguistic Review* 6,71-97.

Prince, Alan and Paul Smolensky 1993 *Optimality Theory*, RuCCS TR-2, Rutgers University.

Schuh, R. 1996 *A Miya Grammar*, University of California Publications in Linguistics, Berkeley.

Steriade, D. 1999 "Alternatives to the syllabic analysis of consonantal phonotactics," in Osamu Fujimura and Bohumil Palek (eds.) *Proceedings of the 2nd Linguistics and Phonetics Conference*, Charles University Press, Prague.

Steriade, D. 2001 “Directional asymmetries in place assimilation” in Elizabeth Hume and Keith Johnson (eds.) *The role of speech perception phenomena in phonology*, Academic Press.

Steriade, D. and J.Zhang 2001 “Context-Dependent Similarity: Phonology & Phonetics of Romanian Semi Rimes” paper presented at the Chicago Linguistics Society.

van Bergem, Dick R. 1995 *Acoustic and Lexical Vowel Reduction*, Institute for Functional Research into Language and Language Use, Amsterdam

van den Broecke, M.P.R. 1976 *Hierarchies and Rank Orders in Distinctive Features*. Van Gorcum, Assen/Amsterdam

Vitz, P. and B.S. Winkler 1972 “Predicting the judged “Similarity of Sound” of English words”, *Journal of Verbal Learning and Verbal Behavior* 12, 373-388.

Walden, B. and A. Montgomery 1975 “Dimensions of consonant perception” *Journal of Speech and Hearing Research*, 18, 445-455

Wang, M. D and Bilger, R. C 1973 “Consonant Confusions in Noise” *Journal of the Acoustical Society of America* Nov. Vol. 54 (5): p. 1248-1266

Wilson, C. 2001 "Consonant Cluster Neutralization and Targeted Constraints," Phonology.

Wingstedt, M. and R. Schulman 1988 "Listener's judgments of simplifications of consonant clusters," *Linguistics* 26, 105-123

Zwicky, Arnold 1976 "Well, this rock and roll has got to stop. Junior's head is hard as a rock," in S.Mufwene, C.Walker and S.Steever (eds.*) Papers from the 12th Meeting of the Chicago Linguistic Society*, Chicago Linguistic Society, University of Chicago.

[1*] Acknowledgments: I am grateful to editors, Kristin Hanson and Sharon Inkelas and to Bert Vaux for detailed and helpful comments on a first draft. Questions from audiences at Berkeley, Stanford, UCLA and the 1st North American Phonology Conferences at Concordia University, Montréal have helped clarify some of the issues discussed here. Most of the writing was done in fall 1999, while I was in residence at the Miller Institute for Basic Science, at UC Berkeley. Many thanks to Andrew Garrett, Larry Hyman and Sharon Inkelas for engineering my stay at the Miller.

This study is dedicated to Paul Kiparsky, whose writings converted me in the late 70's from Indo-European studies to phonology, and whose teaching and example since then have inspired my work.

[1] What exactly causes final obstruents to devoice is a matter of debate: see Lombardi 1999 and Steriade 1999 for two recent proposals, among many. The precise nature of the phonotactic does not affect the problem outlined in this study.

[2] The correspondence constraints cited are those proposed by McCarthy and Prince1995. The argument carries over to other views on correspondence. See section 5.3 below.

[3] See Connine, Titone, Deelman and Blasko 1997.

[4] See Ito 1986, Paradis 1988, Goldsmith 1992, Calabrese 1995.

[5] Mohr and Wang 1968, Walden and Montgomery 1974, van den Broecke 1976.

[6] Flemming's point is that it is faithfulness that's evaluated on perceptual representations: "The output must *sound* like the input". The same holds for any similarity relation.

[7] Cf. Wingstedt and Schulman 1988 and Fleischhacker 1999 for relevant experimental evidence. A further indication of the same point is that partial subsequence rhymes of the form VC_1C_2-$VC_1Ø$ (e.g. *best-less*)are frequent in English (Zwicky 1976) and elsewhere (Steriade and Zhang 2001) while subsequence rhymes of the form VC_1C_2-$VØC_2$ (e.g. *bets-less*)are extremely rare.

[8] See however section 4.2.

[9] See Blevins and Garrett (2002) for relevant discussion and references.

[10] Readers who object that [kitaβ] is a non-structure preserving change – as Turkish lacks bilabial fricatives – will recall that under the ranking Ident [±voice] >>*β, Ident [±cont] structure preservation should be irrelevant. Indeed voicing adjustments can be non-structure preserving in languages like German and Catalan, where final devoicing is incomplete and does not obliterate the contrast. Thus a formal ranking generating [kitaβ] exists under the standard account. The question is why it is unattested.

[11] I are not claiming that C//V is undominated in Latin: there exist clusters like *mbr*, *ltr, str* etc. But the focus here is on the fate of medial C-obstruent-obstruent sequences in which cluster simplification did occur regularly in the spoken language. To obtain the more accurate account, one must assume MAX constraints that outrank C//V. For these MAX constraints, the P-map's claim is that they involve clusters whose individual members are better distinguishable from Ø than the C's that do in fact delete.

[12] See Pulleyblank 1988 for the observation that the same vowel may be both the prevalent target of deletion and the preferred inserted element in selected languages. For a particularly interesting case of epenthetic/deletable C see McCarthy 1993, who discusses post-vocalic r-insertion and deletion in New England varieties of English. Not surprisingly, postvocalic [r] in most varieties of American English is an approximant hardly distinguishable from the end of a preceding low back vowel: it may thus be the closest thing to Ø in that context.