de Lacy, Paul (2007). Freedom, Interpretability, and the Loop. In Sylvia Blaho, Patrik Bye, and Martin Krämer (eds.). Freedom of Analysis? Studies in Generative Grammar (SSG) 95. Mouton de Gruyter, pp. 86-118.

Chapter 8
Freedom, Interpretability, and the Loop

Paul de Lacy

1. Introduction

The aim of this chapter is to propose the 'Interpretive Loop'. The Loop is a limited feedback mechanism: if the winning phonological candidate is phonetically uninterpretable, the Loop eliminates it from the candidate set and EVAL finds a new winner; this process of elimination and re-EVALuation continues until an interpretable form wins. The Loop is needed to maintain the broadest conception of Freedom of Analysis. The alternative is to prevent GEN from creating uninterpretable forms; I will argue that such prevention is neither empirically nor conceptually desirable.

In its broadest conception, Freedom of Analysis allows GEN to create all possible candidates, limited only by restrictions on the formal properties of the available objects and relations (Prince & Smolensky 2004, McCarthy & Prince 1993). Therefore, GEN can create phonetically uninterpretable phonological candidates as well as interpretable ones. An interpretable candidate is one that can be converted by the post-phonological interpretive components (i.e. the phonetic module(s)) into motor commands; an uninterpretable candidate cannot be so converted.

Current theories of phonological representation allow a myriad of uninterpretable forms to be created. For example, a segment with the features [+high, +low] is phonologically well formed. However, it cannot be interpreted (i.e. converted into motor commands) as it requires the tongue to be in two different positions at once (from an articulatory phonetics point of view – Chomsky & Halle 1968:305; the same sort of contradiction arises in an auditorist phonetic implementation). Such cases of phonologically well-defined but phonetically contradictory forms have been and continue to be present in all generative theories (see e.g. Hale et al. 1977). Freedom of Analysis also allows GEN to create interpretively incomplete forms: candidates that lack crucial specifications such as features, prosodic structure, precedence relations, and so on.

The standard view is that an uninterpretable output dooms the derivation (i.e. causes 'crash'). For example, if a winning candidate has a [+high, +low] vowel it cannot be converted into motor commands by the phonetic component. So, there will be no speech output for the particular input involved (often called 'absolute ungrammaticality' or 'ineffability').

Absolute illformedness through uninterpretability is far from innocuous because some uninterpretable candidates are harmonic bounds for many interpretable attested ones: in some (perhaps all) situations uninterpretable candidates always win. Section 2 discusses the pervasiveness of the problem, showing that phenomena such as types of harmony, assimilation, and metathesis always favour uninterpretable candidates, and so should not occur under standard conceptions of grammar. In short, uninterpretable candidates must be eliminated otherwise the grammar will not function.

Uninterpretable candidates can be eliminated through the 'Interpretive Loop'. The Loop prevents the derivation from failing if the phonological output is uninterpretable. If the winner is uninterpretable, it is deleted from the candidate set. Eval is then applied again and a new winner is picked. The Loop is described fully in section 3.

(1) *The Interpretive Loop*

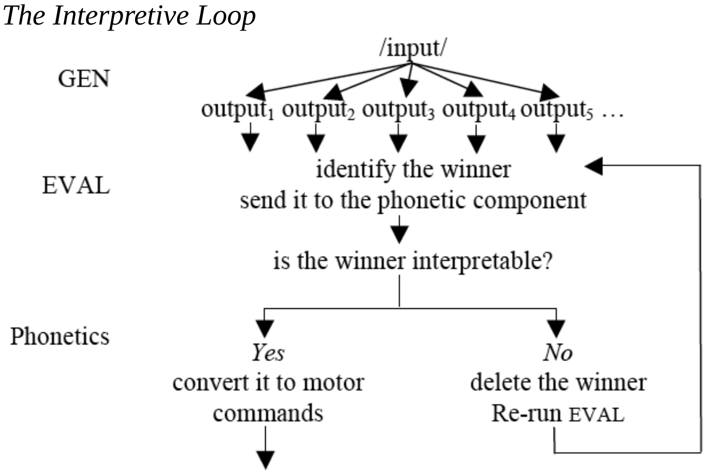

Section 4 discusses alternatives. One is to restrict Freedom of Analysis and ban GEN from producing uninterpretable candidates. Another is to allow the phonetic component to repair uninterpretable forms. I will argue that Freedom of Analysis should not be restricted and the phonetic component cannot repair uninterpretable forms. The Loop's implications are discussed in section 5: it allows many putative phonological constraints to be eliminated. Section 6 presents conclusions.

To avoid unnecessary entanglement later on, it's necessary to mention a few assumptions up front. The Loop is proposed for a unidirectional two-level OT theory of phonology based on innatist (i.e. not functionalist) principles. For comments on bidirectionalism see footnote 7 and on the Loop's compatibility with functionalism and number of levels see section 4.1. I also assume full featural specification, as argued in Steriade (1995), de Lacy (2006:sec.8.4) and references cited therein.

2. The Uninterpretability Problem

This section identifies the problem with allowing GEN to create uninterpretable forms. The main point is that uninterpretable forms are harmonic bounds for attested interpretable ones. They should therefore block those interpretable forms from ever appearing, contrary to observation. Section 2.1 provides some background to interpretability. Section 2.2 identifies ways in which uninterpretable forms affect the derivation.

2.1. What is interpretability?

Interpretability is an issue in all models of phonology in which phonological representation does not directly specify motor commands (i.e. every generative theory of phonology). In such theories, the phonological component generates abstract representations. These representations are converted into motor commands by the 'phonetic component' (Keating 1988, 1990);[1] the conversion process is 'phonetic interpretation'. A phonological form is uninterpretable when the phonetic component cannot convert it into motor commands. Every generative theory proposed so far has the capability of producing uninterpretable forms.

It will be useful to distinguish two types of uninterpretability. One is 'contradiction': where a phonological structure is phonetically interpreted as requiring opposing motor commands. For example, a [+high, +low] vowel requires the tongue to be in two different positions (or in an acoustic approach for different levels of F_1 to be produced simultaneously). Therefore [+high,

+low] cannot be converted into coherent motor commands.

Structures can be contradictory too. The representation in (2) poses an unresolvable problem for interpretation (Sagey 1988, Hammond 1988, Bird & Klein 1990, cf Boersma 2003). ‖H‖ must temporally precede ‖L‖ (‖H‖ refers to the interpretation of the phonological symbol 'H' – it is read as "the phonetic realization of H"), and ‖a‖ must temporally precede ‖o‖. However, the association lines require ‖H‖ and ‖o‖ to occur at the same time, and the same for ‖L‖ and ‖a‖. There is therefore a contradiction: for ‖H‖ and ‖o‖ to be simultaneous and for ‖a‖ to precede ‖o‖, ‖H‖ must follow ‖L‖.

(2) *Structural contradiction: crossed association lines*

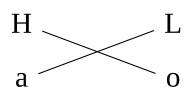

There are many interpretively contradictory structures. Many segmental feature combinations require incompatible realizations (e.g. [−sonorant, +approximant], [−voice, implosive], [strident, glottal]). All of these combinations are phonologically well-formed: they are simply combinations of representational primitives, and so are not formally distinct from interpretable combinations like [+voice,−continuant], [+ATR,+high] and so on.

The other major type of uninterpretability relates to having too little structure. For example, a phonological form that lacks prosodic structure is uninterpretable because it does not provide enough information to create a complete set of motor commands relating to pitch, loudness, and duration. Such forms are 'interpretively incomplete'. (Recall that full specification is assumed here – see section 5 for further discussion).

There may seem to be an obvious and easy way to avoid uninterpretable phonological outputs: simply prevent GEN from creating them. The problems with this approach are discussed in section 4.1.

2.2. Uninterpretability in Optimality Theory

In all models of phonology (and syntax) I am aware of, if an uninterpretable form wins, the derivation ends: if an uninterpretable form is sent to the phonetic component it cannot be converted into motor commands, so no speech will result. If GEN is allowed to create uninterpretable forms, significant problems arise because in many cases – perhaps all – uninterpretable forms will win. In many cases, uninterpretable forms are harmonic bounds for interpretable ones. Examples involving interpretive contradictions are easy to identify: schematically they involve cases where an underlying segment is [αF,αG], a phonological process forces the segment to become [βG], and [αF, βG] is interpretively contradictory. For example, [high] harmony should produce uninterpretable [+high, +low] vowels given the right conditions. Pasiego Montañes Spanish provides a relevant example (McCarthy 2002). McCarthy states that "… all nonlow vowels in a word must agree with the stressed [nonlow] vowel in the value of the feature [high]." The examples in (3) provide alternations. (There is an independent requirement that all vowels agree in tenseness, and a vowel reduction process affects word-final vowels (or perhaps more accurately vowels in the weak position of a foot), requiring them to be one of [a e u ʊ] (e.g. [sentému̲s], *[sentémo̲s])). The underlying value of the root vowel can be seen when the stressed vowel is [a].

(3) Pasiego Montañes [high] harmony

a. harmony with [+high]

/beb/ 'drink'	
[bib-íːs] 2. PL. PR. IND.	[beb-ér] INF.
[bib-íu] PAST PPL. MASC. SG. COUNT.	[beb-ámus]1. PL. PR. SUB.
/kox/ 'take'	
[kux-í] 1SG. PERF.	[kox-ér] INF.
[kux-irían] 3. PL. COND.	[kox-áis] 2. PL. PR. SUB.

b. harmony with [−high]

/sint/ 'feel'	
[sent-émus]1. PL. PR. IND.	[sint-ír] INF.
	[sint-áis] 2. PL. PR. SUB.

The restriction in Pasiego relates to [high] and not to another vowel feature. It is clearly not [ATR] harmony as there is a separate requirement that all vowels in a word agree in tenseness: e.g. [bindiθír] cf. [pɪθígʊ]; [abidúl] 'birch tree' cf. [trænkílʊ] 'quiet (count)'. The requirement that all vowels agree in [high] applies both to tense mid vowels and lax ones: the lax mid vowel [ɔ] is permitted, but it never appears in the same word with a [+high] vowel (e.g. [ækɔlɔdræʊ] 'long, thin').

Tableau (4) presents a straightforward analysis of [high] harmony: a constraint AGREE[high] outranks IDENT[high]. The winner is [bibís]; it contains no interpretive contradictions (or incompleteness), so can be successfully converted to motor commands. σ́-IDENT[high] is responsible for the stressed syllable retaining its underlying [high] specification (i.e. for eliminating *[bebéːs]); its ranking cannot be determined in this competition.

(4) *Pasiego Montañes [high] harmony ranking*

/beb-íːs/	AGREE[high]	IDENT[high]
☞ a. bibíːs		*
b. bebíːs	*!	

The low vowel [a] now presents a problem. Low vowels do not undergo height harmony: e.g. /sal/ 'leave', [sal-ír] {INF.}, [sal-íːs] {2. PL. PR. IND..}, [sal-émus] {1pl. PR. IND.}. However, the ranking requires that with input /sal-iːs/, the output should be the uninterpretable [s❧líːs], where [❧] is used here to symbolize a [+high, +low] vowel. Tableau (5) illustrates; the uninterpretable winner is marked with ☠.

The winner (a) is the least marked and most faithful. It does not violate AGREE[high] because all its vowels are [+high] – i.e. [i] and [❧]. The attested form (b) [salíːs] fatally violates AGREE[high] by having a [+high] [i] and a [−high] [a]. Candidate (c) is an important competitor: it is the most faithful interpretable candidate that has [high] harmony by changing the /a/s [+low] feature to [−low], producing [e]. However, candidate (a) is a harmonic bound for (c): candidate (c) violates IDENT[high] and IDENT[low], while (a) violates only IDENT[high]. In other words, the uninterpretable [❧] is the most faithful way of satisfying the markedness constraint AGREE[high].

(5)

/sal-iːs/	IDENT[low]	AGREE[high]	IDENT[high]
☠ a. s❧líːs			*
b. salíːs		*!	
c. selíːs *!	*		*

So, if there is [high] harmony and an input has a low vowel, the most faithful and least marked form will always be one with a [+high, +low] vowel [❧]. If an uninterpretable winner means the end of the derivation, the practical effect is that if a language has [high] harmony, all inputs with a [+low] vowel and a [+high] vowel should be unpronounceable. Forms like [bibía] show this prediction to be false.

There may seem to be a couple of easy solutions to this problem. For example, why not invoke a constraint *{+high,+low} that outranks IDENT

[low]? Unhappily, there are two problems with this approach. One is that it would require CON to contain constraints that ban uninterpretable structures; this again raises the issue of how CON came to contain such constraints (see section 4.1). It also makes the wrong typological predictions. If there is a constraint *{+high,+low}, it is possible that in some language with [high] harmony it *will* be ranked below IDENT[low], so that every input with a low and high vowel will produce a winner with a [❦]; therefore, for every input with a low and high vowel, no speech should result. No such language has been reported: in every case of harmony, potential uninterpretability is always avoided through failure to undergo harmony or alteration to an interpretable vowel. Having a constraint *{+high,+low} is therefore an inadequate solution.

Does this whole line of argumentation rely on the existence of the features [high] and [low]? Not at all. [high] and [low] are merely used here to illustrate a general property of all extant feature systems: the possibility for interpretive contradiction. It's possible to appeal to 'local' solutions for this particular case, like adopting features like [1height], [2height], [3height]. However, weeding out interpretive contradiction from the entire feature system is much harder, probably impossible (Hale et al. 1977); the Loop makes doing so unnecessary.

In short, under current conceptions of OT and phonology-phonetics interaction, /sal-iːs/ should have the output [s❦́líːs] and be phonetically uninterpretable; therefore, /sal-iːs/ should have no pronouncable output. The general problem this case illustrates – where an uninterpretable candidate is the most faithful least marked candidate – cannot be avoided by introducing constraints to CON or by redefining features so that uninterpretable combinations are impossible.

The Pasiego Montañes Spanish situation arises in all cases where an underlying segment S is specified as [αF, αG], S must become [βG] through an assimilation/harmony/neutralization process, and [αF, βG] segments are interpretively contradictory.

Lokaa ([lòkə́ə́]) provides an analogous example, but instead with autosegmental relations. Akinlabi (to appear) describes a process of tonal metathesis whereby /H+LH/ → [HHL] (for arguments against alternative accounts, see Akinlabi's work). Examples are given in (6).

(6) *Lokaa tone metathesis*

a. /H+LH/ → [HHL]
[lèdʒí] ‘palmseed’ cf. [ɛ́plá jə́ lédʒȉ] ‘palmseed market’
[lètú] ‘head’ cf. [úkpò wə́ léːtù] ‘head towel’

b. Other sequences remain faithful
[kétə̀m] ‘lizard’ cf. [ɛ́plá jə́ kéːtə̀m] ‘lizard’s market’
[éfém] ‘crocodile’ cf. [úkwá wə̀ éːfém] ‘crocodile’s canoe’

The attested response to input /HLH/ is to metathesize to [HHL]. Metathesis involves two changes: (a) reversing the order of tones, and (b) reassociating the tones to new segmental sponsors. However, such a candidate is harmonically bounded by an uninterpretable candidate with crossed association lines, as illustrated in (7). Candidate (a) satisfies the constraint that motivates metathesis (called *HLH here – see Akinlabi to appear for a complete analysis) by being minimally unfaithful – it only violates LINEARITY-tone (the constraint that preserves underlying precedence relations between tones – after McCarthy & Prince 1995). In contrast, the attested interpretable candidate (c) satisfies *HLH, but in doing so violates both LINEARITY-tone and MAX-Association (Myers 1987). In short, candidate (a) is a harmonic bound for candidate (b).

(7) *Tone metathesis should doom the derivation*

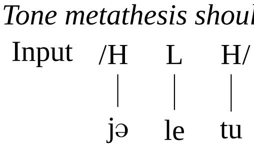

a. Most faithful candidate that avoids [HLH]

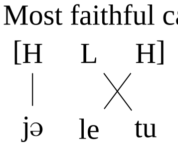

– avoids HLH, preserves input associations

b. Best interpretable candidate that avoids [HLH]

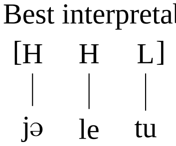

– avoids HLH, *unfaithful* to input associations

As with Pasiego Montañes Spanish, the most faithful way of satisfying a constraint results in an uninterpretable winner: one that places contradictory ordering demands on the tonal and segmental tiers.

One way to avoid this result would be to claim that there is a ban on crossed association lines, either in GEN or as a constraint in CON. The problem with appealing to a restriction in GEN is that it is then necessary to explain how GEN came to have a restriction whose sole purpose is to avoid uninterpretability (more on this in section 4.1). An additional problem with proposing that there is a constraint against crossed lines is that it predicts that there could be some language in which an input failed to have any pronouncable output every time it presented the environment for metathesis. No such language has been reported, to my knowledge.[2]

To generalize, in all cases of metathesis the most faithful response is to reverse the order of the offending tones/features only. However, such a local reverse results in an uninterpretable form. The consequence is that metathesis should not occur in natural language as every time metathesis is motivated the winning candidate is uninterpretable.

The same issue arises with interpretively incomplete forms, but in a more striking way. The constraint ONSET is violated when the leftmost node in a syllable is part of the nucleus (or moraic, depending on the theory). The most minimal way to satisfy ONSET through epenthesis is to insert a root node with no features: i.e. /i/ → [•i], where • is a featureless root node. [•i] is a harmonic bound for all other [Ci] candidates (where C is a more fully specified segment) as it does not violate any featural markedness constraints. For example, a prime competitor is [ti], but it violates constraints against coronals and AGREE[Place] (which favours [tʃi] over [ti]).[3] Therefore, the most harmonic candidate with epenthesis will always contain [•]. However, [•] is interpretively incomplete (cf. proposals in underspecification theories – see de Lacy 2006§8.4 for discussion and references). Therefore, epenthesis should not occur: every input which is forced to undergo epenthesis will have an uninterpretable winner, which of course cannot be pronounced.

The same point can be made for many – perhaps all – phonological processes. For deletion, the most faithful candidate would delete just the offending feature, not the entire segment (deleting the segment results in violations of MAX and CONTIGUITY, deleting the feature alone violates nothing in faithfulness theories with IDENT[F] and not MAX[F]).

An even more surprising point is that uninterpretable candidates should always win regardless of the input. Prosodic structure is non-contrastive (apart from length and main stress), so either prosodic structure is not present in input forms or there are no faithfulness constraints that preserve prosody. Either way, the implication is the same: the most harmonic form will have no prosodic structure. For an input like /tata/, the interpretable output [{(.tá.ta.)}] – i.e. syllabified, footed, and enclosed in a PrWd – inevitably violates several markedness constraints, including constraints on the relation between sonority and syllable constituents, syllable structure, foot form, and PrWd shape. In contrast, the output [tata], with no prosodic structure at all, is equally as faithful as [{(.tá.ta.)}], but violates a proper subset of markedness constraints.[4] Therefore, the most harmonic candidate for any input is one without any prosodic structure; such a candidate is interpretively incomplete, so all words and sentences should be unpronounceable, given current assumptions about OT and the fate of uninterpretable forms.

3. The Interpretive Loop

The preceding section has argued that uninterpretable forms cause a serious problem in OT. At worst, they win in every derivation, meaning that nothing should be pronouncable. At best, they should always win in processes such as epenthesis, metathesis, harmony/assimilation, deletion, and so on; such processes should therefore not be visible. In short, uninterpretable forms are a problem.

There are several potential places to eliminate uninterpretable candidates: GEN, CON, and in a post-phonological component. They could be eliminated in GEN by restricting Freedom of Analysis or in CON by devising constraints against uninterpretable feature combinations and structures. Section 4 will examine the GEN and CON approaches in detail, and will reject them. Instead, a post-phonological mechanism will be proposed here called the ‘Interpretive Loop’. Section 3.1 outlines the form of the Loop and shows how it deals with uninterpretable forms. Section 3.2 discusses the implications of an infinite candidate set for the Loop.

3.1. The Interpretive Loop

The source of the uninterpretability problem is the assumption that uninterpretable winners spell doom for the derivation. I propose that when a phonological winner is uninterpretable, the derivation does not end. Instead, the winner and its associates are deleted from the candidate set, and EVAL runs again with the new candidate set.[5] This process of deletion and candidate set re-evaluation continues until an interpretable winner is found.

The Loop can be illustrated with Pasiego Montañes Spanish [high] harmony. The problem identified in section 2 was that the form [s❦líːs] (with an uninterpretable [+high, +low] vowel [❦]) will always beat every other interpretable candidate that has undergone [high] harmony, such as [selíːs]. It will also beat the candidate in which [high] harmony has been blocked: i.e. [salíːs]. However, the Loop allows [s❦líːs] to be eliminated, summarized in (8).

(8) *The Loop in Pasiego Montañes Spanish [high] harmony*

a. GEN(/sal-iːs/) → candidates {. . . , [salíːs], [s❦líːs], [selíːs], . . . }
   GEN produces a candidate set.
b. EVAL({. . . , [salíːs], [s❦líːs], [selíːs], . . . }) → [s❦líːs]
   EVAL identifies the winner.
c. INTERP([s❦líːs]) → failure
   The interpretive component cannot produce an output.
d. WINNOW → delete [s❦líːs] from the candidate set.
   The uninterpretable winner is deleted from the candidate set.
e. EVAL({. . . , [salíːs], [selíːs], . . . }) → [salíːs]
   EVAL is run again without [s❦líːs]. [salíːs] wins.
f. INTERP([salíːs]) → success
   The interpretive component produces an output.
   The derivation ends.

GEN creates a candidate set related to the input /sal-iːs/ (8a). The candidate set contains [salíːs] (the attested form in which [high] harmony has been blocked), [selíːs] (an interpretable form in which [high] harmony has taken place), and the uninterpretable [s❦líːs], among many other candidates (e.g. [slíːs], [saléːs], etc.). EVAL identifies the winning candidate (8b). As shown in tableau (5), the ranking AGREE[high] ≫ IDENT[high] means that [s❦líːs] will win as all its vowels are [+high] and they are minimally unfaithful. EVAL does not rank-order the entire candidate set (Prince & Smolensky 2004 cf Coetzee 2004). Therefore, at this point [salíːs] and [selíːs] are both losers and are not ordered with respect to each other. The winner [s❦líːs] is sent to the interpretive component (8c). As it cannot be interpreted, a function winnows down the candidate set (i.e. WINNOW) (8d). The winnowing involves deleting the uninterpretable winner [s❦lí:s]. After the candidate set has been cut down, EVAL runs again (8e). In this pass, the primary competition is between [salí:s] and [selí:s]. Both are interpretable. It just so happens in Pasiego Montañes Spanish that blocking of [high] harmony is preferred to altering a low vowel by means of the ranking IDENT[low] ≫ AGREE[high], so [salí:s] beats [selí:s].

(9) EVAL, *pass II*

	/sal-i:s/	IDENT[low]	AGREE[high]	IDENT[high]
☞	a. salí:s		*	
	b. selí:s	*!		*

As tableau (9) shows, [salí:s] wins. As it is also interpretable (8f), it can be successfully converted into motor commands and the derivation ends.

In summary, post-phonology candidate deletion allows interpretable losers to ultimately win. Of course, the result generalizes to all situations where the 'first' winner is uninterpretable. For Lokaa, the candidate with crossed tone association lines wins, but is eliminated by the Loop, and eventually the interpretable metathesized candidate wins.

3.2. The size of the candidate set

The Loop continues to apply as long as EVAL picks uninterpretable winners. Alan Prince (p.c.) has raised a significant issue: what guarantee do we have that, for some input, an interpretable winner will ever be found? Could there be a situation where all interpretable forms are harmonically bounded by an *infinite* number of uninterpretable forms? In such a case, the derivation would loop back forever. This will be called the 'Infinity Problem' here.

The Infinity Problem is different from the oft-discussed issue of whether there are an infinite number of candidates. It is not the number of candidates that is at issue here, it is whether it is ever possible to find an interpretable winner under certain rankings.

The Infinity Problem is easy to construct. For Pasiego Montañes Spanish, it is evident in the candidate [ʔ̲i̲s❦lí:s], which has an initial epenthetic [ʔi]. If DEP ranks below AGREE[high] and IDENT[low], [ʔ̲i̲s❦lí:s] will be the next most harmonic form after [s❦lí:s]. The next most harmonic form after [ʔis❧líːs] will be [ʔiʔis❧líːs], and so on with an infinite number of uninterpretable forms with epenthesis [ʔi…ʔis❧líːs].

In short, the Infinity Problem means that in some cases no interpretable form will ever win; the derivation will continue looping indefinitely.

For Pasiego Montañes, the easy solution is to have DEP outrank AGREE [high], then [salíːs] will beat all candidates with epenthesis […ʔis❧líːs]. However, in some languages DEP is necessarily ranked below featural faithfulness constraints to generate other phenomena. In these languages, the Infinity Problem is unavoidable.

The Infinity problem arises in segmental epenthesis, splitting, and perhaps in prosodification. It does not affect segmental deletion or featural change. The problem does not arise with deletion or feature change because for any length string the number of candidates without epenthesis is finite (these are candidates in which every segment corresponds to a segment in the input form). The finitude follows from the fact that (a) deletion has an upper bound of the length of the input string and a lower bound of 0 and (b) there are a finite number of features.

Consequently, if DEP outranks MAX in a language an interpretable candidate will ultimately be reached. There is always one deletion candidate that can win – one without any phonological content at all (‘∅’). The ∅ candidate is interpretable as ‘no motor commands’, and exists in every candidate set.

The Infinity Problem may apply to prosodification. If there are a finite number of prosodic nodes and prosodic recursion is banned in GEN, then the number of different ways a form can be prosodified is also finite. However, if GEN allows recursion of prosodic nodes, then such structures also pose an Infinity Problem. For example, if the string [ta] can have a candidate with one syllable node, a candidate with two syllable nodes (one on top of the other $[[\text{ta}]_\sigma]_\sigma$), one with three ($[[[\text{ta}]_\sigma]_\sigma]_\sigma$), and so on, there is the potential for an infinite number of uninterpretable forms to beat all interpretable forms.

The simplest way to avoid the Infinity Problem is to place an arbitrary upper bound on segmental epenthesis and splitting, and restrict prosodification. For example, if GEN contains a finite number of prosodic nodes and prosodic recursion is banned (or limited to some arbitrary high finite number), for a string of any length there will be a finite number of possible prosodifications. The idea that there are a finite number of prosodic categories is implicitly accepted in most work. Recursion has been extremely limited as well – it has fallen out of favour for prosodic categories such as the σ and Ft nodes (e.g. Hayes 1981 cf. Hayes 1995), but is still accepted in a limited way for PrWds

(Selkirk 1995, Peperkamp 1997) and perhaps higher units.

Segmental epenthesis can be similarly restricted. Łubowicz (2003:§1.2.1.2) proposes that "there can be only as many segments added to the underlying form as there are segments in it plus 1." So, for an input of length *n* candidates can contain *2n+1* segments or fewer.

I suggest that the restriction be raised somewhat, to at least *4n*. The size of the restriction does not matter, as long as it is finite. The motivation for raising the *2n+1* restriction is that Łubowicz bases it on the idea that "there is only one spot adjacent to each segment in a string of segments available for an epenthetic filler." However, adjacent epenthetic elements do occur: e.g. Slave /t͡saɣ/ 'he was crying' is augmented to [h̲ɛ̲t͡saɣ] for minimal word reasons (Rice 1989:133), and the same for Axininca Campa /na/ → [nat̲a̲] 'carry' (Payne 1981). While the Slave and Axininca Campa outputs do end up with fewer epenthetic segments than the number of input segments and so – strictly speaking – fit within Łubowicz's proposal, the restriction predicts that an input like /n/ in Axininca would not be able to augment to [na̲t̲a̲] as this contains three epenthetic segments while the input has only one. One can imagine a case where there is a single input consonant /p/ and word minimality and head augmentation pressures result in [(pa̲t̲á̲ʔ̲)], where there are four epenthetic segments. Consequently, 4*n* seems a reasonable minimal restriction.

The same sort of restriction can be placed on prosodic recursion. A generous restriction is that a node n can be recursed only for as many segments as there are in the candidate. So, for [pat] there can be a candidate with three recursed PrWd nodes, but no more.

Arbitrary limits on the number of epenthetic segments per input and prosodic recursions may provoke a sense of unease. However, they provide the most effective way of avoiding the Infinity Problem. An alternative is to delete not only the uninterpretable winner but all candidates that share its uninterpretable property. The challenge in such an approach is in determining which part of a candidate is uninterpretable, translating this uninterpretability into phonological representation, and then deleting all candidates that contain this information. The Interpretive Loop mechanism is much simpler if segmental epenthesis and prosodification are restricted.

Restricting segmental epenthesis and prosodification is clearly a limitation on GEN, and therefore on Freedom of Analysis. However, it is a fundamentally different kind of restriction than preventing GEN from creating uninterpretable forms. 'Epenthesis' and 'prosodic structure' are defined solely in terms of phonological representation. Restricting epenthesis and prosodic structure in GEN therefore has a straightforward and well defined expression and formulation in phonological terms. In contrast, 'uninterpretability' is not a phonological concept – it is a phonetic one. Restrictions on phonological forms that are uninterpretable in another component therefore do not admit of a straightforward expression in phonological terms. Section 4 discusses this issue in more detail.

To summarize, restrictions on epenthesis and prosodic recursion mean that there are a finite number of candidates for any input. There is at least one candidate which is interpretable for any input: the contentless candidate $\emptyset$. Therefore, if the Loop re-applies enough times an interpretable form will be reached.

4. Alternatives examined

Section 3 presented a solution to the uninterpretable candidate problem that relied on a post-phonology mechanism. There are alternative methods of avoiding uninterpretable candidates. For example, restrictions could be placed on GEN (i.e. on Freedom of Analysis), preventing uninterpretable candidates from being generated. Alternatively, constraints against uninterpretable structures could be introduced in CON. Finally, the phonetic component could be allowed to repair uninterpretable forms. These options will be examined in turn and argued to be inadequate.[6]

4.1. GEN cannot ban uninterpretable structures

There are empirical and conceptual problems with claiming that GEN cannot produce uninterpretable candidates.

The empirical problem relates to underlying forms. GEN has two roles. It creates output candidates, and it creates underlying forms of lexical items. The process of learning an underlying form requires generation of potential underlying forms for a particular winner (e.g. Tesar & Smolensky 1998). As the only candidate-creation mechanism in the grammar, GEN must therefore create these underlying forms. So, if a lexical entry has an uninterpretable underlying form GEN must have created it, and so GEN must be capable of creating uninterpretable forms.

Lexical items certainly can have uninterpretable underlying phonological forms: the underlying form of lexical items can be interpretively incomplete. For example, underlying forms can contain only part of a segment (Akinlabi 1996 and references cited therein). Such 'featural morphemes' are interpretively incomplete on their own, as are morphemes that contain a single tone. If such lexical items were sent to the interpretive component alone, they would not be pronounced. As GEN can create such uninterpretable underlying forms, it therefore follows that there is no restriction on creating uninterpretable candidates.

Underlying forms also lack prosodic structure, accounting for why there are no languages that have contrastive syllabification, footing, and other prosodic constituents (of course, length and primary stress are exceptions). Forms without prosodic structure are uninterpretable, therefore GEN can create uninterpretable forms. An alternative is to say that all underlying forms have prosodic structure and are therefore interpretively complete in that respect, but that faithfulness constraints do not preserve such structure. The problem with this view is that McCarthy (2000) has argued that there are faithfulness constraints that preserve prosodic structure – they apply on the output-output dimension, and also between opaque candidates and outputs (McCarthy 1999, 2003). Clearly, faithfulness constraints that preserve prosodic structure do exist, but they are effective on every dimension except Input → Output. The most consistent way of accounting for this gap is to assume that underlying forms do not have prosodic structure.

Lexical items can also be interpretively incomplete by lacking segmental features. For example, Inkelas, Orgun, and Zoll (1997) argue that Turkish has three types of lexical items when it comes to the feature [voice]: those that are [−voice], those that are [+voice], and those that lack a [voice] feature. If an obstruent is underlyingly [+voice] or [−voice] it surfaces faithfully in all environments (e.g. /etyd/: [e.tyd] 'study', [e.tyd.-ler] 'study-plural'; /devlet/: [dev.let] 'state', [dev.le.t-i] 'state-accusative'). However, if an obstruent has no [voice] feature underlyingly, it alternates: [ka.nat] 'wing' cf. [ka.na.d-ɯ] 'wing-accusative'). Assuming that a segment is interpretively complete only if it contains all features, the underlying form for [ka.nat] is interpretively incomplete because its final obstruent lacks a [voice] feature.

It is also possible that GEN allows interpretive contradiction in underlying forms. This possibility is far less well documented, and finding evidence for it is difficult in any case. If, for example, an input segment is [+high] and [+low], it may act chameleonic, surfacing as a high vowel in some environments and low in others. The possibility of interpretive contradiction in lexical entries requires further exploration.

In short, if GEN creates underlying forms and (at least) some underlying forms are uninterpretable, GEN must be able to create uninterpretable forms. If GEN is capable of creating uninterpretable underlying forms, the least complex assumption is that it can create uninterpretable output candidates.

There are also conceptual reasons for thinking that GEN can create uninterpretable candidates. Suppose there are restrictions in GEN that prevent uninterpretable structures. These restrictions would include bans on [+high, +low] vowels, crossed association lines, lack of prosodic structure, lack of features, and a myriad of other restrictions. The restrictions are clearly not phonologically unified: there is no single phonological property that they all share. The only factor they have in common is that the structures they ban are not interpretable by the phonetic component. So, how did such restrictions come about?

A functionalist perspective provides a straightforward response: there must be a mechanism that allows construction of phonological restrictions based on performance factors, such as interpretability and other phonetic considerations. The mechanism would detect failures in interpretability, identify the source of the interpretability, and create a restriction in GEN to avoid such situations. Adopting such an approach effectively swings the door open to proposing that all phonological constraints and GEN restrictions are functionally motivated and not innate.

From a formalist/innatist perspective, the existence of a myriad of restrictions in GEN whose sole purpose is to avoid uninterpretable structures would be startling and a remarkably suspicious coincidence. The only justification for such restrictions would be that they are necessary to make the system work – i.e. 'minimally conceptually necessary' in Minimalist terminology (Chomsky 2001). In other words, the restrictions exist because without them there would be no way for the grammar to produce any output. This approach works for some phonological restrictions. Every output candidate needs prosodic structure to be interpretable. If forms without prosodic structure are harmonic bounds for all of those with prosodic structure, a GEN restriction that requires prosodic structure on candidates would have to evolve for there to be any output at all (if, of course, the Loop does not exist). So, it would be no surprise that such a GEN restriction exists: without it, there would be no speech.

However, many interpretability restrictions are not essential for the grammar to work. For example, if there was no GEN restriction on [+high, +low] vowels very few derivations would be doomed – only those inputs with underlying [+high, +low] vowels or those with a low and a high vowel that underwent [high] or [low] harmony would have no output. The same goes for a ban on crossed association lines: crossed lines only arise in specific situations, not in every derivation. To summarize, an innatist approach that appeals to GEN restrictions on uninterpretable structures has the problem of explaining why those restrictions exist. They are not all 'minimally conceptually necessary', so one would have to appeal to a massive amount of coincidental random mutation/exaptation to account for their existence.

If one wishes to avoid functionalism, the Loop avoids a solution to the problems identified above. With the Loop, GEN is free to create uninterpretable forms. Consequently, lexical items can have uninterpretable underlying forms. There are no restrictions on uninterpretable structures in GEN, so there is no need for a mechanism that 'looks ahead' to the phonetic component and constructs restrictions based on what can and cannot be interpreted.

4.2. CON cannot ban uninterpretable structures

Instead of putting restrictions on GEN, one could appeal to constraints in CON: i.e. there are constraints for every uninterpretable structure. For example, there would be a constraint *[+high, +low], a constraint against crossed lines, and constraints requiring full prosodification.

Constraints against uninterpretable forms make a range of incorrect predictions. If a constraint against an uninterpretable structure exists, it can be ranked so that for some input an uninterpretable output will always win. For example, if there is a constraint *[+high, +low] and it ranks below IDENT[low] and AGREE[high] in a language with [high] harmony, every input with both a [+high] and a [+low] vowel will have an uninterpretable winner with a [+high, +low] vowel. I am not aware of such pervasive effects in harmony systems where an underlying combination of low and high vowels spells doom for the derivation.

In short, having constraints against uninterpretable structures predicts pervasive types of ineffability that are not observed.

4.3. The Interpretive component cannot repair uninterpretable forms

An alternative is to restricting GEN or CON is to allow the phonetic component to repair uninterpretable forms. For example, if the phonetic component encounters a [+high, +low] vowel, it could interpret it as a low vowel. However, there are several problems in allowing the phonetic component such power.

A central problem is that the same uninterpretable form would have to be interpreted differently in different grammars. For example, in Pasiego Montañes Spanish [high] harmony, the winner is the uninterpretable [s❦lí:s]. The phonetic approach would be to invoke an interpretative principle whereby [❦] is interpreted as a low vowel – i.e. phonetically [salí:s]. However, in another language low vowels may participate in [high] harmony, becoming [silí:s] (cf Jingulu – Pensalfini 2002); the phonetic approach would have to allow another interpretive procedure whereby [❦] is interpreted as a high vowel. The learner would have to choose which interpretive rule to use. At this point the phonetic component effectively has the power to alter contrastive phonological specifications, and therefore to 'do' phonology.

Giving the phonetic component such power also means that it does not have to respect phonological restrictions. For example, [❦] could be interpreted phonetically as [a] even if the language's phonology bans this segment.

The ultimate result of permitting the phonetic component to repair uninterpretable structures is that it is thereby given power to alter, add, and delete phonological structures on a language-specific basis. In effect, it takes over the role of the phonological component.

With the Loop, the phonetic component can be restricted in terms of what it can and cannot do. While the phonetic component may introduce elements that have no phonological correlate (e.g. interpolation between tone specifications, intrusive segments through gestural overlap), it cannot eliminate features or segments or ignore phonological precedence relations. In short, the Loop allows the roles of the phonology as providing sound specifications and of the phonetics as an interpretive component to be maintained.

5. Implications of the Loop

The Loop allows Freedom of Analysis to have a great deal of freedom in positing different types of structure. It also allows GEN and CON to be simplified significantly: GEN and CON do not need to contain restrictions and constraints whose sole purpose is to ensure interpretability. Identifying all such constraints would require a complete theory of phonetic interpretation (i.e. what the phonetic component can and cannot interpret and what it requires in order to produce a complete set of motor commands), so I offer only a few suggestions and examples here.

Prosodic structure is converted to information about duration, loudness, intensity, pitch targets, and so on. Without such information, a candidate is uninterpretable. Since the Loop will eliminate unprosodified candidates, there is therefore no need for constraints or GEN restrictions to require prosodic structure. So, there is no need for GEN to require every segment to be syllabified, or for every candidate's prosodic structure to contain each member of the prosodic hierarchy. This contrasts with constraints like PARSE-σ "Every syllable must be part of a foot" – violation of PARSE-σ does not mean uninterpretability, and languages differ on the degree of syllable parsing. Therefore, the Loop does not affect the existence of PARSE-σ.

It is possible that the hierarchical order of prosodic nodes is also determined by interpretability. If a structure in which a σ node dominates a Ft node is uninterpretable, there is no need to have a restriction in GEN that imposes the correct dominance relations on the prosodic hierarchy.

The requirement that every prosodic constituent has a head is also essentially an interpretive requirement (Selkirk 1984, 1995; cf. Crowhurst 1996 cf. de Lacy 1999). A head is interpreted as the locus of duration enhancement, raised pitch, and so on. If a constituent's head is not marked in a candidate, it is therefore interpretively incomplete. The effects of the constraint HEADEDNESS (Selkirk 1995) are therefore an epiphenomenon of the Loop.

There is no need for constraints against interpretively incompatible features. For example, a constraint *[+high, +low] is unnecessary as all winners with [+high, +low] vowels are uninterpretable and will be eliminated as contenders for pronunciation. Similarly, there is no need for a ban on crossed association lines, either as a restriction on GEN or as a constraint. Crossed association lines are uninterpretable as they impose contradictory precedence requirements; consequently, any winner with crossed association lines will be eliminated, so there is no need for a ban on them in CON or GEN.

Assuming that the phonetic component cannot fill in missing feature values, only candidates with fully featurally specified segments will ever survive phonetic interpretation. Consequently, there is no need for constraints that require 'full specification' – i.e. that all segmental features be present for every segment in every candidate.

The Loop allows restrictions on the formal properties of phonological relations like precedence to be removed. For example, precedence (the ordering relation between nodes on a tier) is transitive, asymmetric, and irreflexive: i.e. [tak] is (informally abbreviating[7]) $\{t<a, t<k, a<k\}$. There is no need for GEN to require that precedence is irreflexive. GEN can generate precedence relations between any members; it could generate a candidate $\{t<a, a<a\}$. However, $\{t<a, a<a\}$ is uninterpretable: it requires that ‖a‖ temporally precede its realization. The same goes for the symmetry property: the precedence relations $\{t<a, a<t\}$ require that ‖t‖ precedes ‖a‖ and ‖a‖ precedes ‖t‖ – again an interpretive contradiction. Therefore, there is no need to specify that precedence is asymmetric – any candidates with a symmetric relation that is to be interpreted as temporal order are interpretively incompatible.

Of course, the Loop cannot eliminate every GEN restriction or CON constraint. It can only eliminate 'inviolable' properties or constraints; if a property differs from language to language it must be controlled by a constraint.

The Loop can also apply to syntax. Some proposals have already been made to eliminate syntactic restrictions and consider them as following from interpretive restrictions. For example, Heim & Kratzer (1998) propose that syntactic trees that are not binary branching are semantically uninterpretable. They define the interpretive mechanism of functional application so that it requires two daughters of every node (effectively one to provide the function and the other as its argument). Functional application therefore cannot apply to ternary- or unary- branching trees, so all such trees will be eliminated. With the Loop, there is therefore no need for a syntactic principle that demands binarity. Hale & Keyser (1993) and Heim & Kratzer (1998:51ff) make a similar argument for theta theory.

Apart from constraints, the Loop has implications for levels and 'intermediate forms'. Interpretability is only relevant for the winner. All other forms – underlying and intermediate forms – do not have to be interpretable. For underlying forms, this has no suprising implication – extant theories assume that underlying forms are uninterpretable in that they lack a great deal of structure (e.g. prosodic constituency).

However, problems arise for 'intermediate' forms – candidates that are not the winner but influence it, as in Sympathy or Cumulativity Theory (McCarthy 1999, 2003 resp.). The sympathetic candidate does not have to be interpretable because it does not pass through the interpretive component. For example, McCarthy (1999) discusses a case of opacity in Tiberian Hebrew where /deʃʔ/ is realized as [de.ʃe] – the winner's epenthetic vowel only appears in the winner because it aims to be faithful to the sympathetic (and unrealized) form [de.ʃeʔ]. The proposal works because [de.ʃeʔ] is the most harmonic of a designated set of losing candidates (all those that preserve input consonants). However, if having prosodic structure is not demanded by GEN, there is a more harmonic form: [deʃʔ] with no prosodic structure at all. This form is superior to [de.ʃeʔ] in faithfulness (it does not violate DEP) and in markedness because it has no syllable structure and so does not violate constraints like NOCODA. There is no way to eliminate unprosodified [deʃʔ] as the sympathetic winner because lack of prosody is only 'banned' by passing through the interpretive component, and the sympathetic candidate form never passes through it. The same applies to other theories such as Bye's (2001) 'virtual phonology'.

In a nutshell, the problem is that the intermediate form in analyses of opacity must be interpretable. However, if interpretability is an 'epiphenomenon' of the Loop and intermediate forms never pass through the interpretive component, there is no way to require that they be interpretable. The same goes for multi-level theories (e.g. Kiparsky 2006, McCarthy 2007): the output of every non-final level will be uninterpretable.

The Loop therefore implies that the grammar is strictly two-level with no reference to intermediate forms and losing candidates. I note this as a consequence of the Loop, and not self-evidently a good or bad thing. The majority of theories of opacity so far have used reference to a losing form (or at least uninterpreted form – e.g. McCarthy 1999, 2003, Jun 1999, Bye 2001), though some have not (McCarthy 1994, Goldrick 1999).

A final comment is that the Loop guarantees that for every input there will be an interpretable output. This raises the issue of what to do when an input has no (obvious) output (see Fanselow & Féry 2002 for discussion). An example is the English input *beautiful+er{comparative}*, which has no obvious output realization (**beautifuller*, **beauter*). In syntactic theories (e.g. Minimalism) it is common to ascribe ineffability to uninterpretability, but the Loop eliminates doing so as an option. In contrast, recent proposals about ineffability have not appealed to uninterpretability as a means of accounting for ineffability; instead, Prince & Smolensky (2004) argue that a 'null parse' candidate can account for some cases, and a variety of other proposals have also been made (see McCarthy 2002:198ff for an overview).

The Loop also has implications for pruning the candidate set. Samek-Lodovici & Prince (1999) observe that there are 'perpetual losers' in the candidate set – those that are singly or collectively harmonically bounded by interpretable candidates. If all perpetual losers could be eliminated, the candidate set would be finite. Riggle (2004) proposes an algorithm along these lines. However, uninterpretable candidates are harmonic bounds for interpretable ones. For example, from input /tak/ the candidate [tak] with no prosodic structure incurs a proper subset of markedness violations of those of any prosodified form. So, if all candidates that are harmonically bounded were eliminated from contention, there would be no interpretable competitors left. A straightforward solution to incorporate Riggle's proposal is to regenerate candidate sets for each loop: the Loop can ban a member of the candidate set from being generated, so GEN will be run again and again and eventually all uninterpretable candidates will be banned: then interpretable candidates will no longer be harmonically bounded.

6. Conclusions

The first aim of this article was to show that uninterpretable candidates pose significant problems in Optimality Theory. In many cases – perhaps all – uninterpretable candidates are harmonic bounds for interpretable ones. The standard belief that uninterpretable winners cause the derivation to stop ('crash') means that many inputs should have no pronounceable output, contrary to fact.

The second aim was to show that Freedom of Analysis can remain free: GEN can generate uninterpretable candidates as long as there is an Interpretive Loop. If the winning output is uninterpretable, it is eliminated from the candidate set and evaluation is run again, and so on until an interpretable winner emerges. This proposal requires limiting the candidate set to a finite number of candidates by restricting segmental epenthesis and prosodic recursion.

The Loop has implications for the constraint component CON. It eliminates all constraints whose sole purpose is to impose interpretability, such as the ban on crossed association lines, bans on interpretively incompatible feature combinations, and prosodic requirements such as exhaustive parsing and HEADEDNESS.

The Loop also has implications for the alphabet of phonological features and relations. The Loop means that not every phonological feature must be phonetically interpretable; the Loop will eliminate every winner that contains inherently uninterpretable features. Of course, for the phonological component to do its job (i.e. produce some interpretable form), it must contain enough features and relations to make some interpretable segment(s). However, the Loop makes it possible for the phonological component to function even if it contains uninterpretable features and relations and – of course – if GEN can create uninterpretable forms.

The next step is to determine how much the Loop is responsible for. The Loop provides the *potential* for Freedom of Analysis to be extremely free. However, the Loop has limitations: a putative GEN-restricting principle can only be ascribed to the Loop if it bans an uninterpretable structure; the Loop has nothing to say about violable restrictions.

Notes

*. My thanks to two anonymous reviewers, and to John McCarthy, Alan Prince, John Kingston, Lisa Selkirk, and Steve Parker for commenting on earlier versions of this work. The proposals in this chapter were presented in several different venues over the past several years. My thanks to the phonology group at the University of Massachusetts, Amherst in 1999, the MIT Phonology Circle in 2000, the audience at Concordia University in 2000 (in particular Charles Reiss), and the audience at Rutgers University in 2004.

1. I use 'motor commands' to refer to the output of the phonetic component. The output may be articulatory commands or acoustic targets – see Kingston (2006) for an overview. The relevant issue here is that the phonetic component converts the phonological output into a different representation, whatever that representation might be like.
2. Another approach is to deny that there is ordering on the tonal tier. However, the existence of floating tones in output forms shows that there must be precedence relations between tones (Pulleyblank 1986).
3. There is no consonant that perfectly satisfies all featural markedness constraints – see de Lacy (2006:sec.1.3.4).
4. Of course, this argument assumes that candidates should be fully prosodified (or only minimally unprosodified (e.g. only at edges)). This assumption sits comfortably with recent views that the effects of putative extraprosodicity are due to constraint interaction (e.g. Hung 1994). The only constraint an unprosodified form might violate is one that requires segments to belong to syllables ('PARSE-seg'). In more general terms, to get a prosodically complete form to beat a prosodically incomplete form, there would have to be constraints that required the presence of every prosodic level. The problem this approach raises is that it predicts languages that cannot pronounce words that violate some prosodic constraint. For example, the ranking ‖ MAX, DEP, NO-CODA ≫ PARSE-SEG ‖ means that the winning candidate for /pat/ is [pat] with no prosodic structure (prosodified [.pa.ti̠.], [.pa.], and [.pat.] are all ruled out). Therefore, /pat/ should not be pronouncable. In any case, the Loop allows constraints like PARSE-SEG to be eliminated from CON: note that PARSE-SEG's sole role is to ensure interpretability.
5. Coetzee (2004) proposes that all candidates are rank-ordered – winners as well as losers. This full rank-ordering is achieved by removing the winner from a candidate set and

running EVAL with the smaller candidate set (p.3). Despite its superficial similarity to the Loop, Coetzee's proposal is fundamentally different. The mechanism Coetzee describes does not literally eliminate candidates from the candidate set – it rank-orders losers. In contrast, the Loop really eliminates candidates, and losers are not rank-ordered. The Loop is compatible with Prince & Smolensky's (2004) proposal that EVAL picks the winner in a candidate set but does not distinguish between the winners. In short, Coetzee's proposal rank-orders losers and does not eliminate candidates, while the Loop does not rank-order losers and does eliminate candidates.

6. All of the discussion in this chapter assumes a uni-directional production model of OT. Some versions of bi-directional OT require every syntactic form to have a semantic interpretation, and this could in principle be applied to phonology phonetics: every phonological form could be required to have a phonetic interpretation, so ruling out uninterpretable candidates (see Beaver & Lee 2003 for discussion of various bi-directional models). The possibility that some version of a bi-directional model might be used to ban uninterpretable candidates is in itself uninteresting for the purposes of this article as the focus here is on uni-directional theories. It would be interesting if uni-directional models could not deal with the uninterpretability problem; however, as this chapter argues, they can if the Loop exists. This chapter eliminates the uninterpretability problem as a challenge for uni-directional OT; it provides no insight as to which of bi- and uni-directionalism is correct.
7. To be accurate, a tier is a string, and a string can be defined as a function from a finite set of elements S (drawn from a denumerably infinite set of elements (like natural numbers)) to phonological primes (features, nodes); precedence relations hold between the members of S, not between the phonological primes.

References

Akinlabi, Akinbiyi
1996 Featural affixation. *Journal of Linguistics* 32(2): 239–289.
to appear The prosodic organization of Lokaa tones. In *Proceedings of the 24th West African Languages Congress*.
Beaver, David I. and Hanjung Lee
2003 Input-output mismatches in OT. In *Optimality Theory and pragmatics*, Reinhard Blutner and Hank Zeevat (eds.), 112–153. Palgrave: MacMillan.
Bird, Steven and Ewan Klein
2000 Phonological events. *Journal of Linguistics* 26: 33–56.
Boersma, Paul
2003 Nasal harmony in functional phonology. In *The phonological spectrum*, Jeroen van de Weijer, Vincent J. van Heuven and Harry van der Hulst (eds.), pp. 3–35. (Vol. 1: Segmental structure) Amsterdam: John Benjamins.
Bye, Patrik
2001 Virtual Phonology: Rule sandwiching and multiple opacity in North Saami. Ph. D. diss, University of Tromsø. ROA 498.
Chomsky, Noam
1993 A minimalist program for linguistic theory. In *The view from Building 20:*

Essays in honor of Sylvain Bromberger, Kenneth Hale and Samuel J. Keyser (eds.), pp. 1–52. Cambridge, Mass.: MIT Press.

Chomsky, Noam
2001 Derivation by Phase. In *Ken Hale: A life in language*, Michael Kenstowicz (ed.), pp. 1–52. Cambridge, Mass..: MIT Press.

Chomsky, Noam and Morris Halle
1968 *The sound pattern of English*. New York: Harper & Row.

Coetzee, Andries
2004 What it means to be a loser: Non-optimal candidates in Optimality Theory. Ph. D. diss., University of Massachusetts, Amherst. ROA 687.

Crowhurst, Megan
1996 An optimal alternative to conflation. *Phonology* 13:409–424.

Fanselow, Gisbert and Caroline Féry
2002 Ineffability in grammar. In *Resolving Conflicts in Grammars*, Gisbert Fanselow & Caroline Fery (eds.), 265–307. (Sonderheft der Linguistischen Berichte) Hamburg: Buske.

Goldrick, Matthew
1999 Turbid output representations and the unity of opacity. In *Proceedings of NELS 30*, Ji-Yung Kim and Mako Hirotani (eds.), 231–245. Amherst, Mass.: GLSA Publications.

Goldsmith, John
1976 Autosegmental phonology. Ph. D. diss., MIT.

Gordon, Matthew
2006 Functionalism in phonology. In *The Cambridge Handbook of Phonology*, Paul de Lacy (ed.), ch.2. Cambridge: Cambridge University Press.

Hale, Kenneth, LaVerne Masayesva Jeanne, and Paul Platero
1977 Three cases of overgeneration. In *Formal Syntax*, Peter W. Culicover, Thomas Wasow, and Adrian Akmajian (eds.), 379–425. New York; London; San Francisco: Academic Press.

Hale, Kenneth and Samuel J. Keyser
1994 On argument structure and the lexical expression of syntactic relations. In *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger*, Kenneth Hale & Samuel Jay Keyser (eds.), 53–109. Cambridge, Mass.: MIT Press.

Hammond, Michael
1988 On deriving the Well-formedness Condition. *Linguistic Inquiry* 19: 319–325.

Hayes, Bruce
1981 A metrical theory of stress rules. Ph. D. diss., MIT.
1995 *Metrical Stress Theory: Principles And Case Studies*. Chicago: University of Chicago Press.

Heim, Irene and Angelika Kratzer
1998 *Semantics in Generative Grammar*. Malden, Mass.: Blackwell.

Inkelas, Sharon, Orhan Orgun, and Cheryl Zoll
1997 The implications of lexical exceptions for the nature of grammar. In *Derivation And Constraints in Phonology*, Iggy Roca (ed.), 393–418. Oxford: Oxford University Press.

Jun, Jongho
1999 Generalized sympathy. In *Proceedings of NELS 29*, Pius Tamanji, Masako Hirotani, and Nancy Hall (eds.), 121–135. Amherst, Mass.: GLSA.

Keating, Patricia
1988 The Phonology-Phonetics Interface. In *Linguistics: The Cambridge Survey*, Frederick Newmeyer (ed.), 281–302. (Volume I: Grammatical Theory) Cambridge, UK: Cambridge University Press.
1990 Phonetic representations in a generative grammar, *Journal of Phonetics* 18: 321-334.

Kingston, John
2006 The phonetics-phonology interface. In *The Cambridge Handbook of Phonology*, Paul de Lacy (ed.), ch. 17. Cambridge, UK: Cambridge University Press.

Kiparsky, Paul
2006 *Paradigmatic effects and opacity*. University of Chicago Press.

Łubowicz, Anna
2003 Contrast preservation in phonological mappings. Ph. D. diss., University of Massachusetts, Amherst. ROA 554.

de Lacy, Paul
1999 Sympathetic stress. ROA 294.
2006 *Markedness: Reduction and preservation in phonology*. Cambridge, UK: Cambridge University Press.

McCarthy, John
1984 Theoretical consequences of Montañes vowel harmony. *Lin-guistic Inquiry* 15(2): 291–328.
1994 Remarks on phonological opacity in Optimality Theory. ROA 79.
1999 Sympathy and Phonological Opacity. *Phonology* 16: 331-399.
2000 The Prosody of Phase in Rotuman. *Natural Language and Linguistic Theory* 18: 147-197.
2002 *Optimality Theory: A thematic guide*. Cambridge: Cambridge University Press.
2003 Sympathy, Cumulativity, and the Duke-of-York Gambit. In *The Optimal Syllable*, Caroline Féry and Ruben van de Vijver (eds.), 23–76. Cambridge: Cambridge University Press.
2007 Restraint of Analysis. This volume.

McCarthy, John and Alan Prince
1993 Prosodic morphology I: Constraint interaction and satisfaction. Rutgers Technical Report TR-3. New Brunswick, Rutgers University Center for Cognitive Science. ROA 482.
1995 Faithfulness and reduplicative identity. ROA 60.

Myers, Scott
1997 OCP effects in Optimality Theory. *Natural Language and Linguistic Theory* 15(4): 847-892.

Payne, David L.
1981 *The phonology and morphology of Axininca Campa*. Summer Institute of Linguistics Publications in Linguistics 66. Dallas, Texas.: SIL.

Pensalfini, Rob
2002 Vowel harmony in Jingulu. *Lingua* 112: 561-586.

Peperkamp, Sharon
1997 *Prosodic words*. HIL Dissertations 34. The Hague: Holland Academic Graphics.

Prince, Alan S., and Paul Smolensky
1993 Optimality theory. Constraint interaction in generative grammar. Technical Report #2, Rutgers University Center for Cognitive Science. ROA 537; Published in 2004 by Blackwell Publishers).

Rice, Keren
1989 *A Grammar of Slave*. Berlin/New York: Mouton de Gruyter.

Riggle, Jason
2004 Contenders and Learning. In *Proceedings of the 23 Annual Meeting of the North East Linguistic Society*, Benjamin Schmeiser, Vineeta Chand, Ann Kelleher and Angelo Rodriguez (eds.), 101-114. Somerville, Mass.: Cascadilla Press.

Sagey, Elisabeth
1988 On the ill-formedness of crossing association lines. *Linguistic Inquiry* 19: 109–118.

Samek-Lodovici, Vieri and Alan Prince
1999 Optima. ROA 363.

Selkirk, Elizabeth O.
1984 *Phonology and Syntax: The relation between sound and structure*. Cambridge, MA.: MIT Press.
1995 The prosodic structure of function words. In *Papers in Optimality Theory*, Jill Beckman, Laura Walsh Dickey and Suzanne Urbanczyk (eds.), 439–470. (University of Massachusetts Occasional Papers 18.) Amherst, Mass., GLSA.

Steriade, Donca
1995 Underspecification and markedness. In *The Handbook of Phonological Theory*, John Goldsmith (ed.), 114–174. Cambridge, Mass.: Blackwell.

Tesar, Bruce and Paul Smolensky
1998 Learnability in Optimality Theory. *Linguistic Inquiry* 29: 229–268.

Subject index

Language index

Author index