[Please cite as Report no. LRC-97-01, Linguistics Research Center, UCSC, Santa Cruz, CA.]

Markedness, Segment Realisation, and Locality in Spreading*

Máire Ní Chiosáin
University College, Dublin

Jaye Padgett
University of California, Santa Cruz

Think globally. Act locally.
Car Bumper Sticker

1 Introduction

One of the fundamental observations of linguistic theory is that linguistic elements do not interact over arbitrarily great distances. Instead, relations obtain locally. This conviction would seem trivial if all phonological interaction were obviously confined to adjacent segments, or all syntactic relations to adjacent words. In reality, the assumption of locality requires work for those assuming it; in phonology, it has for this reason been one of the important forces driving the elaboration of metrical and autosegmental representations.

Consider the case of long-distance feature spreading, or harmony. It is well known that segments within a spreading domain may appear to be non-participants, transparent to the harmony process. A few types of such transparency are listed in (1).

(1) a. Vowel harmony skips consonants
  b. Coronal harmony skips vowels and non-coronal consonants
  c. Nasal harmony skips glottal stops

Various strategies have been proposed to account for such cases of transparency. Within non-linear phonological frameworks, a property many approaches have in common is the preservation of locality by relativising it to what might very generally be called a *legitimate target*: some notion 'anchor', 'projection' or 'feature-bearing unit'. Locality is obeyed so long as spreading does not skip such a legitimate target. Notable examples of this line of thinking include Goldsmith (1976[79]) and Clements (1976a) on the notion 'feature-bearing unit', Halle and Vergnaud (1978) on 'projections' of features, Kiparsky (1981) on the notion 'harmonic vowel', and Archangeli and Pulleyblank (1987, 1994) and Anderson and Ewen (1987) on the relativisation of adjacency to prosodically or geometrically defined anchors. (See Odden 1994 for a recent example of this strategy and a review of the area.) The basic idea is depicted in (2), where a feature F is linked to the elements $\mathbf{T_1}$ and $\mathbf{T_2}$, legitimate targets in some respect. Locality is not violated by the skipping of intervening α, since α lacks whatever property it is that grants legitimacy (e.g. it is not F-bearing, has the wrong prosodic status, or lacks a certain feature geometric node, see the references cited above). Equivalently, the elements $\mathbf{T_1}$ and $\mathbf{T_2}$ are adjacent for the purposes of F-spreading.

(2) Local linkage (relativised to legitimate targets **T**)

```
T1  α  T2
  \    /
    F
```

The above approaches are sometimes combined with assumptions about underspecification: the intervening segment α might be transparent because it is unspecified for either F, or whatever feature/node makes α a legitimate target. (See for example Paradis and Prunet 1989 and Shaw 1991 on [coronal] transparency.)

In a series of recent works an alternative view of locality in spreading is considered, in which spreading is seen as strictly, segmentally local (Ní Chiosáin and Padgett 1993, McCarthy 1994, Itô, Mester and Padgett 1995, Flemming 1995a, Padgett 1995, Gafos 1996, Walker 1996). According to this view, all segments in a spreading domain are necessarily participants, as shown in the hypothetical vowel harmony example in (3)a. This contrasts with the common view whereby vowels alone are legitimate targets of such harmony, as in (3)b.

(3) Contrasting views of locality in spreading

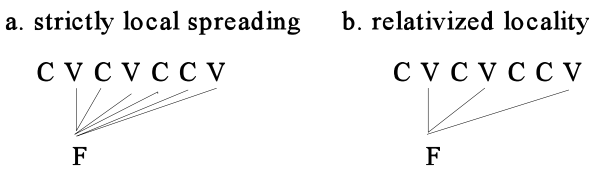

The first goal of this paper is to present an argument for local spreading that is strict in two ways. First, spreading respects strict *segmental* adjacency, as proposed by the references above and illustrated in (3)a. An essential result of this view is that segments are either blockers or participants in spreading; there is no transparency or skipping. Second, segmentally strict locality is inviolable; in Optimality Theoretic terms, Gen does not produce structures such as (3)b. As we will see, these assumptions are fundamental to an account of certain asymmetries in long-distance feature spreading. In particular, strict locality (along with other basic assumptions) explains why vocalic place features spread long-distance, while consonantal place features do not. More generally, strict locality is a likely key to resolving a basic problem: some features spread long-distance and others do not, a fact that follows from nothing in the theory otherwise. These points are the focus of section 3.

In an obvious sense, it is more restrictive to countenance only blockers and participants in spreading than it is to include a third class of transparent segments. This point can hardly constitute an argument in itself, however, in the absence of an alternative account of seemingly transparent segments. The second goal of this paper, pursued in section 4, is to flesh out and motivate such an account. To preview, consider the case of round harmony in a Turkish word [somun]$_{Rd}$ (section 4.2). The brackets indicate a span of [round]; our main concern is the intervocalic consonant, here *m*, which is in fact round, in the simple sense of coproduced with the lip rounding of a neighboring vowel. The reason consonants like the *m* here are thought to be skipped in round harmony is because the theory acknowledges only one kind of round consonant, notated C^w; yet the *m* in Turkish *somun* is not perceived to warrant the transcription m^w. A central point here is that phonological theory must entertain more segment types than has generally been the case; in particular we have the round *m* of *somun* in the sense given above, which we notate $[m]_{Rd}$ and distinguish from m^w. More important, recognising more segmental realisations does not lead to overgeneration of contrasts, given the right theory to control their occurrence and distribution. The necessary theory for this is an instantiation of the theory of adaptive dispersion (Liljencrants and Lindblom 1972, Lindblom 1986, 1990), implemented within Optimality Theory (Prince and Smolensky 1993), following Flemming (1995b, to appear). A satisfactory account of feature spreading, it turns out, depends on the essentially bidimensional understanding of markedness and contrast this theory provides: segment wellformedness depends on constraints phonetically grounded in articulatory complexity on the one hand, and others grounded in the needs of perceptual contrast on the other. Flemming (1995b) argues this point in considerable detail; here we show that the same ideas make sense of apparent transparency in a striking way.

The price of accepting transparency in phonology includes the relativisation of locality in various ways; the hope is that one theory of locality can bring coherence to the notion 'legitimate target', and unity of explanation to the various kinds of apparent transparency. This goal has in fact eluded the theory. The price of rejecting transparency, on the other hand, is more segment types for phonology, and a theory to control them, Dispersion Theory.[1] This is by far the lower price, it turns out. First, Dispersion Theory's bidimensional understanding of markedness has ample motivation independently; indeed it is absolutely necessary for an adequate understanding of sound patternings. In addition, the extra segment types entailed are phonetically real, and so must be generated by a full theory of sound patterning in any case. Second, it makes possible a genuine explanation for the cross-linguistic typology of 'transparency', in particular the asymmetry in the spreading behaviour of consonant and vowel place outlined above. Finally, this approach to locality brings an explanatory unity to diverse cases of 'transparency', and in fact eliminates the theory of locality as such. No theory of relativised locality fares as well in any of these areas.[2]

In the interest of keeping the discussion focussed and manageable, the claims made here are circumscribed in a few ways. First, we deal with assimilation or spreading only, remaining silent on the question of locality as it relates to dissimilations, 'floating features', and other phenomena. Though it is important in the end to consider how all of these facts relate (or fail to—see the conclusion), the topic of assimilation is both coherent as a point of departure, and worthy of its own exploration. Second, we leave open the possibility that there are constraints on locality in spreading that go beyond what is discussed here, though this issue certainly merits further scrutiny as well. For instance, Odden (1994) argues that phonological rules can be constrained so as to apply only under syllabic or segmental adjacency, extending no further even if no blockers intervene. Most importantly, for concreteness the discussion here focuses on the facts indicated in (1); in various places we suggest how the ideas could extend to other relevant phenomena. A brief discussion of more problematic cases of transparency can be found in the appendix.

Earlier works in phonology have excluded 'action at a distance'. A notable precedent is found in works in Firthian Prosodic Phonology, with its notion 'prosody'. A prosody is a notation generally respecting the strictest locality on the one hand; it is abstract on the other hand, sometimes having a discontinuous or varied phonetic exponence. (See for example Allen 1951, Waterson 1956, Ogden and Local 1994, Ogden 1995, among others.) Though the ideas we motivate are different, they could be viewed as an attempt to explore one understanding of the notion 'phonetic exponence'. From a very different theoretical perspective, Articulatory Phonology's notion 'gesture' by definition constitutes a continuous, uninterrupted articulatory event (Browman and Goldstein 1986 et seq.). We take up this latter point in the beginning of the next section, where we lay out some background assumptions.

2 Preliminary assumptions

2.1 The statement of strict locality

Consider the formulation of locality given below (Kiparsky 1981, Levergood 1984, Archangeli and Pulleyblank 1994, Pulleyblank 1994a; cf. Smolensky 1993). In contrast to the works cited, we take α,β,γ to be segments; that is, every segment is a legitimate target, so that locality is not achieved by relativising adjacency to any notion of tier, projection, anchor, or the like.

```
(4)   *α β γ
        \ /       where F is any feature, and α,β,γ are segments
         F
```

Various works have attempted to elucidate the intuition underlying (4) and to explicate more fully the formal assumptions behind it, often in relation to the No Crossing Constraint. (See Sagey 1988, Hammond 1988, Bird and Klein 1990, Scobbie 1991, Archangeli and Pulleyblank 1994; see Coleman and Local 1991 for more general relevant discussion.) As some of these authors note, it seems desirable to understand the restriction by reducing it to an issue of linear precedence relations. Assume, for instance, that the association lines in (4) mean that F overlaps α, while also overlapping γ (Sagey 1988). Since β is ordered between α and γ, F necessarily overlaps β as well. However, this conclusion follows only if F is taken to be a 'continuous, uninterrupted, unitary' entity (Scobbie 1991:64) and some formal rigor is given to this notion. To this end, Bird and Klein (1990:41) provide a definition of a *convex* phonological event (the term is borrowed from van Benthem 1983:68), which we adapt in (5) (cf. Scobbie 1991).

(5) A featural event F is convex iff it satisfies the following condition:

For all segments α, β, γ, if α precedes β, β precedes γ, α overlaps F and γ overlaps F, then β overlaps F.

This definition, understood in the context of the relevant formal assumptions about association and precedence (see the references above), is not actually reduced in comparison to (4). It is a necessary axiom, however, precisely a rigorous statement of what it means for a gesture to be continuous, uninterrupted and unitary. These properties are plausible attributes of phonetic gestures, as Scobbie notes. They are entailed by the conception of a gesture in Articulatory Phonology: gestures in this theory are by hypothesis implemented as dynamical systems having necessarily exactly these properties. We assume that convexity in the sense of (5) holds of phonological representations without exception, as stated earlier: in Optimality Theoretic terms, it constrains the candidate set that Gen produces.

Segmentally strict locality has the benefit of simplifying the question of blocking versus participation in feature spreading, in particular rendering line crossing a non-issue in phonology, another point Scobbie (1991) argues (cf. Archangeli and Pulleyblank 1994). To see this, consider the hypothetical case of [+back] spreading in (6), in which spreading from the initial vowel is threatened by the contrary [-back] secondary specification on the consonant. This case resembles a scenario of Turkish discussed in section 4.2. In (6)a spreading of [+back] is blocked (and the final vowel receives its [back] specification from the palatalised consonant). The only alternative to blocking is participation, as in (6)b, given strict locality. (6)c is an attempt at avoiding the problem by skipping the consonant, a logical possibility that can arise only if strict locality is not assumed. Since vowel harmony is in fact typically thought to skip consonants, eliminating this scenario has required the No Crossing Constraint (or other axioms that derive the constraint, see the references above). With locality segmentally strict, however, (6)c reduces to the issue of segment skipping, independently ruled out. Consider then (6)d. Once again there is no need to appeal to line crossing, since in this case the question reduces to one of feature incompatibility par excellence, since the *k* is specified for both values of [back] at once.

(6) Scenarios of blocking, participation and skipping

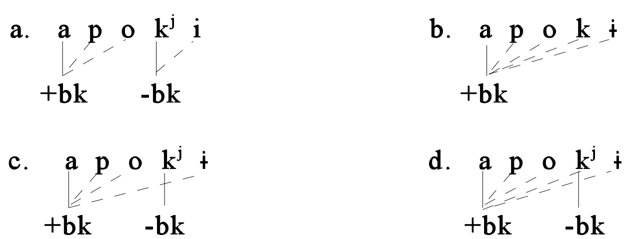

A comparison of (6)b and (6)d makes clear all that is needed in order to understand blocking. (6)d dramatises the worry of segment illformedness, familiarly understood in terms of phonetically grounded feature cooccurrence conditions in much work (see especially Kiparsky 1985, Archangeli and Pulleyblank 1986, 1994). On the other hand, (6)b represents the price paid of sacrificing the [-back] specification on the consonant. (Segment wellformedness is of course an issue here as well, though not to the degree seen in (6)d.) The undesirability of this sacrifice is implicitly acknowledged in autosegmental works assuming 'no feature changing', and explicitly as a faithfulness or featural identity violation in Optimality Theory (see Prince and Smolensky 1993 and McCarthy and Prince 1995). Recasting these ideas somewhat according to what follows in section 4, we understand blocking to result either from constraints on segment wellformedness (violated in (6)d), phonetically grounded specifically in articulatory simplicity, or from constraints requiring the preservation of contrast (violated in (6)b), a matter grounded in perceptual distinctiveness. To put it more generally, blocking is always a matter of substantive markedness concerns, where 'markedness' is broadly construed along these two dimensions. The issue of line crossing simply cannot arise.

2.2 Gestures

A further guiding assumption of this paper is that the basic articulatory unit of phonology is the *gesture*: in the case of major place, a unitary bundle of articulator and oral stricture features (Browman and Goldstein 1986, 1989, Padgett 1991[95], 1994). For our purposes, oral stricture can be characterised in terms of the feature values shown below (Chomsky and Halle 1968, Clements 1990); this particular choice of features for oral stricture is not crucial to what follows.

(7) Oral stricture: segment classes defined

	V(owel/glide)	**A**(pproximant)	**F**(ricative)	**S**(top/nasal)
[cons]	-	+	+	+
[approx]	+	+	-	-
[cont]	+	+	+	-

With the abbreviations V, A, F, S covering the feature value complexes defined above, partial representations for various labial segments are given in (8), where *u*, *f*, etc. stand for any appropriate representation of segmenthood (e.g. Root node, timing slot, etc.).

(8) Representations:

u/w	*β*	*f*	*p/m*
|	|	|	|
Lab_V	Lab_A	Lab_F	Lab_S

The important result here of adopting gestures as the basic articulatory unit follows when these units spread. By hypothesis, entities like $[\text{coronal}]_F$ and $[\text{labial}]_S$ form indissoluble units (i.e., 'feature bundles') of major place (by which we mean [labial], [coronal], [dorsal] or [pharyngeal]) and oral stricture, as the notation is intended to suggest. Spreading of major place gestures entails spreading of all of these properties. (It should be noted that we sometimes employ the terms [back], [round], etc., in what follows, for the sake of familiarity, and to put aside orthogonal questions of the feature inventory. In every case, though, these terms implicate gestures of vocalic stricture in the above sense.)

2.3 Segmental representations

The segmental representations assumed in this paper are minimal; specifically, there are segments, and they have associated features, as in (9). In terms of place, η^w consists of a stopped dorsal articulation and a vocalic labial articulation. Following Padgett (1995, to appear) feature classes like Place and Laryngeal are understood as set-theoretic postulates, instead of being encoded directly in the representation in the form of class nodes. This moves increases the explanatory depth of the basic feature class idea, and allows an account for facts outside the scope of feature geometry as usually construed (see the references for arguments).

(9) Representational parsimony within the segment

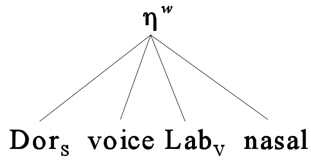

An important consequence of this view of segmental representation is that class nodes are not available as anchors for relativised locality purposes in spreading, nor as blockers (see below, and Padgett 1995 for discussion).

2.4 Optimality Theory

The analyses and general thinking of this paper are in good part propelled by the central tenets of constraint ranking and violability of Optimality Theory (henceforth OT, Prince and Smolensky 1993). A related property of this paper is the reduced role given to explanation in terms of representation. (On the status of representations as essentially inviolable constraints and the sometimes detrimental effect on explanation, see Padgett to appear.) Rather, locality issues are understood by means of the interplay of substantive, typically phonetically grounded constraints in an Optimality Theoretic hierarchy of violability. For arguments and analyses underpinning this general view of explanation in phonology, we refer the reader to Prince and Smolensky (1993), Smolensky (1993), McCarthy (1994), Cole and Kisseberth (1994, 1995a), Padgett (1995), Flemming (1995a,b), Gafos (1996), and Zoll (1996). In the case at hand, the relevant arguments appear below.

3 Strict locality and long-distance feature spreading

3.1 Spreading of vocalic vs. consonantal gestures

This sub-section takes up and elaborates on ideas of Ní Chiosáin and Padgett (1993). The point here is to show that the assumption of strict locality (as understood above) contributes necessarily to an understanding of the possibilities of long-distance place feature spreading in languages. The central observation is a simple one: while vocalic place gestures spread long-distance, as in vowel harmony patterns like the hypothetical examples in (10)a, consonantal place gestures do not, (10)b. Assimilations involving consonantal major place across vowels, as in (10)b, are simply not a feature of adult phonologies.[3] It is important here for the claim to be clear: as we will see below, the idea is not that no features distinguishing among consonants, even relevant to consonantal place, can spread long-distance (this does indeed happen in coronal harmonies, for example); rather, the idea is that what spreads cannot itself be a consonantal [labial], [coronal], [dorsal] or [pharyngeal] gesture. We adopt from Smolensky (1993) and Cole and Kisseberth (1994, 1995a) the convention of indicating a feature spreading domain by means of labeled bracketing; thus [kɔlɪmʊn]$_{\text{RTR}}$ indicates a word with the feature [RTR] spread throughout. Given the prohibition on segment skipping, this notation is unambiguous, and more convenient than the use of association lines.

(10) Asymmetry in spreading of vocalic and consonantal place gestures

a. vowel harmonies	b. non-existent consonantal gesture spread
[kolumun]$_{\text{Lab-V/round}}$	/pad/ → [pab]$_{\text{Lab-S}}$
[kɔlɪmʊn]$_{\text{RTR}}$	/kan/ → [kaη]$_{\text{Dor-S}}$

Assuming for the moment that spreading is strictly local as claimed here, the observation is that consonants (generally) participate in the spreading of vocalic place gestures, while vowels block consonantal place gesture spreading.

This fact has a rather well-known phonetic correlate. A large number of studies of articulatory dynamics support the general view that vocalic place gestures are articulatorily contiguous. Even when consonantal place gestures seem to intervene, the consonantal gestures are in fact superimposed on these vocalic gestures, with consonant and vowel being coproduced (Öhman 1966, Carney and Moll 1971, Fowler 1983, Browman and Goldstein 1990; see Gafos 1996 for further references, and an extensive discussion of the implications of this fact for phonology). While vocalic place gestures might be perturbed in various minor ways by superimposed consonantal place gestures, they essentially function as though contiguous on their own independent 'channel'. This point is best illustrated with a diagram as in (11), informally rendering a gestural score (Browman and Goldstein 1986, 1989, 1990) for the word 'coffee'. The important observation here is that consonants can be, and typically are, fully overlapped by vowels, but not vice versa. In section 4.2 we will make significant use of this coproduction of consonantal and vocalic place gestures.

(11) Simplified gestural score for 'coffee': contiguous vocalic gestures

```
                          a           i
Vowel 'channel'       |-----------|------------|
                        k           f
Consonants            |---|       |---|
```

Citing earlier studies, Gafos (1996) notes two reasons why vowel place pervades the articulatory space in this way. First, while consonantal constrictions are typically quick and highly localised, vocalic gestures involve a more global orchestration of tongue body, jaw height, tongue root, and so on, and tend to be slower. (Browman and Goldstein encode the difference in rate in terms of degrees of 'stiffness' for the respective gestures.) Second, the movements required for vowel place are largely independent of those required for consonant place, and it is possible to marshal both kinds of place constriction simultaneously (see the references cited). Still, though these points are consistent with the claim that vowel place features entirely overlap those of consonants, it remains unclear why the reverse scenario, in which consonant place features entirely overlap those of vowels, cannot occur. To explain both the phonological observation in (10) and its phonetic correlate, it is necessary to make reference to what we call the 'bottleneck effect' on the one hand, and to constraints on syllabification on the other; both notions are well motivated independently.

The bottleneck effect is not a principle of phonology or phonetics, but rather a general physical constraint on flowing systems, one that necessarily restricts the range of possible speech sounds. In relation to speech in particular, the result is the following: a segment's oral stricture is equal to that of its most constricted place component. (See the related discussion of 'tube geometry' in Browman and Goldstein 1989.) To illustrate the idea, various combinations of place gestures are considered below along with a representative example of a resultant segment. (12)a-b are trivial: segments with only one oral articulation display the oral stricture of that articulation. Combining vocalic gestures as in (12)c necessarily results in a vocalic segment; this would be true no matter how many such gestures were present. In the same way one consonantal gesture on top of another as in (12)d only makes a segment place-wise more complex; its basic (stop) stricture is not affected. The import of the bottleneck effect is clearest from example (12)e. There is only one possible interpretation of a segment combining consonantal and vocalic constrictions, namely that in which the greater constriction of the consonantal gesture determines the basic nature of the segment (i.e. for purposes of syllabification etc.): the result is a consonant (with a secondary vocalic articulation). This point is implicit in conventions of transcription: w^p (or ^pw) is not a possible transcription precisely because such a segment cannot be meaningfully distinguished from (12)e.

(12) The bottleneck effect

	gestural components	*resultant segment*
a.	Lab_V	*u*
b.	Lab_S	*p*
c.	Lab_V, Dor_V	*u*
d.	Lab_S, Dor_S	*kp*
e.	Lab_S, Lab_V	p^w

It should be clear at this point what goes wrong when consonantal place gestures spread across vowels. Given the bottleneck effect, and the assumption of strict locality, such spreading can only occur at the cost of making the vowel itself consonantal; spreading as in (13)a results in e.g. [p::], [pb:] or something similar. (We leave the precise result open, simply noting the problem with a question mark.)

(13) Bottleneck effect + strict locality ➜ syllabically impossible outcomes

```
a.   /p a d/  →  p ? b             b.   /k a n/  →  k ? η
       |           \ | /                  |            \ | /
      Lab_S        Lab_S                 Dor_S         Dor_S
```

Though syllabic consonants (and even obstruents) exist, what we are confronting in (13) is something worse, and possibly unattested. According to Bell's (1978) survey, syllabic consonants arise very generally through the loss of a vowel—often [ə]—in metrically weak syllables (as in German *haben* ~ *habm̩* 'to have' or English *harden* ~ *hardn̩*). They do not arise by assimilation to a consonantal gesture, nor by vowel 'hardening' of *any* sort. Thus, any theory employing the feature [cont], for example, must in some way rule out [-cont] spreading to a vowel, as well as vowel 'fortition' of any kind, even in languages allowing syllabic consonants. The issue therefore goes well beyond the markedness of syllabic consonants as it is usually conceived (though for many languages a constraint against syllabic consonants would be enough to rule out (13)). What matters for our purposes is that, independently of the claims made in this paper, vowel hardening must be excluded from the theory, or at least rendered extremely marked, by *some* means. However, in order to be concrete, we consider a specific line of thinking on this issue here.[4]

The following strong claim seems to have the desired consequences: consonants never occupy the syllable nucleus; rather, syllabic consonants are margins of syllables having a reduced vowel in the nucleus, typically schwa. (See also Coleman 1996 on this understanding of syllabic consonants.) Let us make use of constraints on syllable structure introduced by Prince and Smolensky (1993:134), specifically the family of constraints defined below.

(14) *PEAK/X : X must not be parsed as a syllable peak (i.e., associated to the nucleus)

In terms of (14), the proposal is that *PEAK/X is inviolable where X is consonantal, so that nuclei must be vowels. Assume, then, that *PEAK/CONSONANTAL constrains the candidate set produced by GEN, so that syllabic consonants are represented as in (15)a (the consonant can also occupy the onset, as in [pə]), and never as in (15)b. To see how this helps, it is useful to compare the German variant *habm̩* with a hypothetical derivation like /habĩ/ → [habm̩] resulting from spreading of Lab_S from *b* to nasalised *ĩ*. The latter derivation seems unattested, though it is just one of the ways that (15)b could be brought about in principle. Given the general prohibition on consonants in the nucleus, it is ruled out, and similarly spreading as in (13) is never possible.

(15) Syllabic consonants as margins of syllables with reduced vowel

a. Reduced vowel

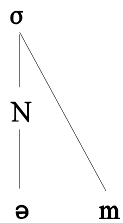

= [əm] or [m̩]

b. Violates *PEAK/CONSONANT

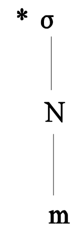

The key to this result lies in taking seriously Bell's observation: so called syllabic consonants are intrinsically connected to vowel reduction. Though real *habm̩* and the hypothetical case just considered are phonetically identical, they differ in the means by which they come about (i.e. real vowel reduction vs. impossible vowel hardening). The simplest way to make sense of this fact, and to eliminate the hardening example, is to *equate* syllabic consonants and syllables with reduced nuclei, regarding e.g. [əm]$_{\sigma}$, [mə]$_{\sigma}$ and [m̩]$_{\sigma}$ as phonetic variants of a phonological schwa in the nucleus. This nucleus occurs when vowel place distinctions are neutralised, as in metrically weak syllables. What such a structure violates is not e.g. *PEAK/C, but simply whatever constraints militate against this loss of contrast (see section 4.1). It is an obvious result of this view that the theory becomes incapable of representing a contrast between e.g. [əm] and [m̩], presumably a good result independent of our main concerns here. The ultimate implications of this idea cannot be explored further, since this discussion departs significantly from the paper's focus. For our purposes the important point is that derivations as in (13) must be at least extremely marked from any theoretical perspective.[5]

Recall the basic observation to be explained: vowels block the spreading of consonantal place gestures, while consonants do not block the spreading of vocalic place gestures. The contrast is schematically shown in (16) below, repeating (10) above. The problem with (16)b is that such spreading inevitably obliterates the CVC structure of a syllable, extending a consonantal constriction through the nucleus. Given an account for the severe markedness of such a possibility—such as the view of syllabic consonants outlined above—such a derivation is not possible.

(16) Asymmetry in spreading of vocalic and consonantal place gestures

a. vowel harmonies	b. non-existent consonant gesture spread
[kolumun]$_{\text{Lab-V/round}}$	/pad/ → [p?b]$_{\text{Lab-S}}$
[kɔlɪmʊn]$_{\text{RTR}}$	/kan/ → [k?η]$_{\text{Dor-S}}$

In contrast, vocalic place gestures as in (16)a can spread long-distance, because such spreading never alters the basic oral stricture of a segment: by the bottleneck effect, adding a vocalic place gesture to a segment can have no such effect. The result of superimposing a vocalic constriction on a sequence CVC, for example, is a sequence that remains CVC in basic stricture. In fact, coproduction of vowels during consonants is the normal case: recall the discussion around example (11), 'coffee', which we can notate as [kɔ]$_{ɔ}$[fi]$_{i}$ to indicate the actual span of the vowels, anticipating section 4.2. This explanation of the asymmetry in long-distance spreading between consonantal and vocalic gestures relies on the key assumptions of this paper: spreading is strictly local (such that all segments in the domain are participants), and articulators like [labial] and [coronal] are gestures, in the sense of being intrinsically specified for oral stricture. A question that remains to be addressed involves the precise nature of the realisation of the vocalic gesture on the participating consonants. This issue, and questions of realisation in general, are the focus of section 4 below.

It is important to note that the argument above does not entail that consonants can be permeated by vocalic place spreading in every case. For example, consonants specified themselves for secondary vocalic place features can block such spreading, for obvious reasons involving feature compatibility. The arguments above prohibit entirely the spreading of consonantal place gestures across vowels; they merely *allow* the spreading of vocalic place gestures across consonants. As always, other factors may intervene to prevent the latter kind of spreading. (See for instance the discussion of Turkish in section 4.2.)

3.2 Consonant harmonies

Ní Chiosáin and Padgett (1993) further argue that the assumption of strict locality, and the use of gestural primitives, sheds light on the facts of so-called consonant harmonies. The main result of the proposals in the preceding section is that consonantal major place gestures cannot spread long-distance (i.e. across a vowel). It follows that long-distance harmonies triggered by consonants necessarily spread something other than a consonantal major place gesture. This result turns out to provide a very natural characterisation of the attested consonant harmonies.

Emphasis or pharyngealisation harmonies provide a simple illustration of this point. A secondary pharyngeal articulation is a distinctive property of certain coronal consonants in Palestinian Arabic (as it is in other dialects). This articulation spreads long-distance; leftward spreading is illustrated in (17). (See Herzallah 1990 for extensive discussion of the spreading processes and their phonetic effects on the participating vowels and consonants.)

(17) Emphasis/pharyngealisation spreading in Palestinian Arabic

/ballaas^{ʕ}a/	→	[bɑllɑɑs]$_{ʕ}$a	'thief (fem)'
/ʕabiit^{ʕ}/	→	[ʕɑbiit]$_{ʕ}$	'unruly (masc)'

Though the spreading originates from consonants, what matters for our purposes is what spreads. Emphasis is in fact a place gesture of vocalic constriction, similar in this property to palatalisation or labialisation on consonants (Herzallah 1990, McCarthy 1991); this is therefore just the kind of 'consonant harmony' we expect.

The best known consonant harmonies, perhaps, are those involving coronal consonants. Relevant examples include Chumash (Poser 1982), Sanskrit (Steriade 1986, Schein and Steriade

1986) and Tahltan (Shaw 1991). Consider the well-known facts of the retroflex harmony in Sanskrit known as 'Ṇati'. Retroflexion spreads rightward from continuants (*ʂ* and *r*) to *n*. As shown in (18)a-b (respectively), spreading occurs both under segmental adjacency and long-distance. For the moment we defer the question of what exactly spreads, calling it *x*. (18)c shows that intervening coronal consonants block this spreading.[6]

(18) Retroflexion harmony in Sanskrit

a.	i[ʂɳ]$_x$a:	'seek (present)'
	p[rɳ]$_x$a:	'fill (present)'
b.	pu[ra:ɳ]$_x$a	'fill (middle participle)'
	k[ʂub^{h}a:ɳ]$_x$a	'quake (middle participle)'
	čak[ʂa:ɳ]$_x$a	'see (middle participle)'
	k[rpama:ɳ]$_x$a	'lament (middle participle)'
c.	ma[r]$_x$ǰa:na	'wipe (middle participle)'
	k[ʂ]$_x$veda:na	'hum (middle participle)'

Steriade (1986) characterises Ṇati as a spreading of [coronal]; Shaw (1991) makes a similar claim for Tahltan. In our terms this is not possible: spreading a fricative coronal gesture Cor_F long-distance would make fricatives of the intervening vowels. It follows from this that something other than [coronal] spreads, something that does not entail consonantal constriction. The obvious candidate for such spreading would be whatever 'minor' gesture(s) determine the tongue tip orientation of retroflexion itself, assuming that retroflexion itself implies nothing about segment stricture, and so is articulatorily compatible with the intervening segments in the spreading domains of (18). That retroflexion itself has no strictural implications seems clear from the existence of retroflex stops, liquids and vowels across languages. Though a retroflexion gesture has an inherent *degree* of displacement, as a major place gesture does, the displacement itself is basically in the 'horizontal' dimension, not in the 'vertical' dimension; it is only the latter that translates into stricture. This is the content of the major/minor place terminological distinction we have been assuming. (See section 4.3 for further discussion.)

Previous accounts of other coronal harmonies similarly assume that what spreads is some minor place feature distinguishing among coronals. Poser (1982), for example, analyzes Chumash harmony as the spreading of [anterior]. The argument from Steriade (1986) and Shaw (1991) for the spreading of [coronal] itself rests on cases—'Nati', and Tahltan coronal harmony respectively—where it is claimed that more than one feature distinguishing among coronals spreads; this necessitates the spreading of some superordinate node in the now familiar feature geometric reasoning, and [coronal] is the likely choice for this node. In our terms two alternative conclusions can be considered. First, the relevant tongue blade/tip orientation features are grouped together under a node, or within a feature class, that is *not* [coronal], as suggested by Ní Chiosáin and Padgett (1993). Second, these cases actually involve single feature spreading, e.g., the spreading of a privative feature of retroflexion instead of the familiar [-anterior, -distributed] complex of SPE argued for by Steriade (1986). Something like the latter view is likely correct for retroflexion, though the status of this property in feature theory has been unclear. In recent work, Gnanadesikan (1994) proposes to dispense with the feature [anterior], and (following Selkirk 1991) to treat retroflexion as a secondary vocalic [dorsal] gesture for coronals, in part to explain the patterning of retroflex consonants with back vowels. Flemming (1995b) argues that this latter patterning is acoustic at heart rather than articulatory, but similarly posits a single feature for retroflexion. Finally in an extensive survey of known coronal harmony systems, Gafos (1996) motivates two new features subclassifying coronals (meant to supersede [anterior] and [distributed]), and argues that all coronal harmony systems involve the spreading of one of these features. Gafos' analysis of Sanskrit relies on a single feature value indicating retroflexion.

Two benefits immediately obtain from the view that what spreads in coronal harmonies is not [coronal] but features further specifying coronals for tongue blade/tip orientation. First, we can now explain why coronal harmonies are not mirrored by labial or dorsal consonant harmonies, or any kind of long-distance spreading of these features. The spreading of consonantal major place gestures is unattested, the point of the preceding section; if consonantal [coronal] can spread, why not also consonantal [labial] and [dorsal]? Second, we explain why spreading in coronal harmonies always originates from only a subset of a language's coronal segments, in particular segments that are contrastively specified for e.g. apicality, laminality, or retroflexion, and always targets other coronal segments. This is exactly what is entailed by the view that what spreads is whatever feature(s) distinguish the *relevant coronal gestures.* Again, this fact remains a mystery if what spreads is [coronal] itself. In fact, [coronal] spreading formulations tend to be unnecessarily complicated because they ignore this generalisation; thus, in Sanskrit we could not simply posit [coronal] spreading, but must anyway restrict spreading to retroflex coronals. The general point is that the assumptions argued for in this paper lead naturally to conclusions about coronal harmonies that are required for independent reasons. We return to the analysis of Sanskrit in section 4.3.

3.3 Section summary and implications

The point of the preceding discussion has been to argue for segmentally strict locality, by showing that this assumption predicts just the right range of long-distance place spreading phenomena. In particular, long-distance place spreading never involves consonantal major place gestures; instead what spreads are vocalic place gestures, or gestures having no bearing at all on segment stricture, including those of tongue tip orientation. This account relies on gestures as articulatory primes, and so provides an additional argument for them. In addition, the basic question of blocking versus participation is decided solely by substantive markedness constraints. While such constraints are a well-motivated element of phonological explanation (see e.g., Kiparsky 1985 and Archangeli and Pulleyblank 1986, 1994 on feature cooccurrence constraints), they are often used in conjunction with notions of relativised locality and concomitant structural devices. A basic point here is that the substantive explanations can and must do all of the work.

Consider an alternative explanation for the asymmetry between consonant and vowel place spreading, one that relies on a feature geometric difference between consonant and vowel place (Clements 1985, 1991, Clements and Hume 1995). Suppose that vowel place features are organised under a feature geometric class node that consonants lack, called VPlace in (19) (Clements 1985, 1991, Steriade 1987, Prince 1987, Odden 1991, Ní Chiosáin 1994, Clements and Hume 1995 and references therein). Vowels, on the other hand, share the superordinate organising node Place. Then the asymmetry might follow from line crossing. Consonants will be transparent to VPlace spreading by virtue of lacking the VPlace node, as shown in (19)a. (Consonants specified for secondary vocalic features would have the VPlace node and so block such spreading.) In addition, individual vowel place features can skip consonants while still respecting locality, if locality is relativised to the VPlace anchor/projection. In contrast, spreading of a consonant's Place node across a vowel is impossible, since (by assumption) both consonants and vowels bear this node; such spreading is prevented by line crossing, as shown in (19)b.

(19) Transparency and blocking through feature geometric projections

a. Consonants are transparent to vocalic place spreading

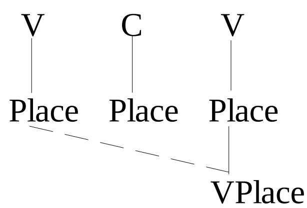

b. Vowels block consonantal place spreading

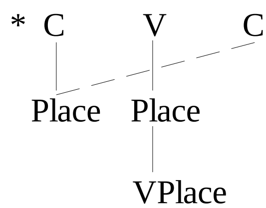

As Ní Chiosáin and Padgett (1993) note, this explanation for the spreading asymmetry fails in an important way. The relevant empirical observation concerns not simply total place assimilation across vowels, but spreading of any consonantal gesture across vowels, even a single, terminal, feature, say [labial] or [dorsal]. Yet the geometry of (19) says nothing about the failure of single articulator spreading as in /pad/ → **pab*. A separate problem concerns the spreading in (19)a: while spreading of individual vowel place features across consonants might be considered local if VPlace nodes are the only legitimate targets, VPlace spreading itself is not local in this sense, since both consonants and vowels bear its anchor node Place. A Place node is skipped, raising the question of what determines 'local' spreading in this case. Next to line crossing in (19)b an entirely different kind of locality statement is therefore required. This geometrical account therefore suffers on grounds of both empirical adequacy and theoretical unity.

One might attempt to construct a different or more elaborate geometry that makes the correct predictions. However, this general strategy seems dubious for several reasons. First, on independent grounds the role of class nodes in explaining basic feature class generalisations is doubtful (Padgett 1995, to appear); that is, class nodes fail to adequately explain the very facts they were originally intended for. Though they might be called on all the same to explain the logically independent facts of blocking and transparency, they should be judged according to their success in this latter area alone. On that note, however, a second point arises: class nodes have not in general provided a successful account of blocking and transparency effects. While we defer the main discussion of transparency to section 4, a broad perspective on the problem can be achieved by considering the competition between two explanations for the blocking effect just considered, one relying on class nodes, line crossing and further required statements as in (19), and one calling on substantive markedness constraints (implicating segment wellformedness for instance) in conjunction with locality as pursued here. It seems likely that we should strive in principle to eliminate one of these two kinds of explanation rather than accept a hybrid theory; class nodes like VPlace after all always mirror a substantive property of the segments that bear them (e.g., [-consonantal]). We see no hope in reducing all known explanations to feature geometric ones. At the same time, the latter face difficulties and apparent contradictions in the most circumscribed empirical domains they are applied to, as in the case of (19) above. (For another argument to this effect, see Walker's 1996 reanalysis of the 'variable dependency' of [nasal] required by Piggott's 1992 geometric approach to nasal harmony facts.)

To pursue the point further, one might ask in what sense a geometrical account as in (19) above constitutes an explanation at all, apart from the drawbacks already noted. Though it succeeds to some extent in encoding the asymmetry in consonant versus vowel place spreading, it cannot itself explain why things are this way. It would be just as simple to posit a class node that only consonants have, and not vowels, wrongly entailing a world in which consonantal place features spread long-distance, while vocalic place features remain confined. (Van der Hulst and van de Weijer 1995 also note the essentially diacritical use of class nodes in this explanation.) The only hope of rescuing the theory from this arbitrariness lies in an appeal to the substantive basis of the geometry: we can say there is a node VPlace because of the phonetic fact that vowels are articulated in a separate 'channel' as noted earlier. In fact, the works cited above arguing for something like VPlace have made this connection to the work of Öhman (1966) and others. This phonetic fact is indeed very likely related to the phonological generalisations, as we noted earlier—but is itself in need of explanation. As we saw, the explanation is simply rooted in the substantive properties of the segments themselves. Given strict locality, these substantive properties provide an explanation for the phonological spreading asymmetry with no further ado.

Another kind of explanation for the asymmetry above, again exploiting a relativised locality, relies on the idea that only vowels constitute syllable heads or (nuclear) moras. This basic idea is pursued in various ways. A very direct approach treats prosodic nodes such as the syllable as the anchor position of the spreading feature, as shown below (e.g., Anderson and Ewen 1987, van der Hulst and van de Weijer 1995; cf. Pierrehumbert and Beckman 1988 on tonal association). Assuming that vowels are syllable heads and that only heads implement a feature associated in this way, only vowels will implement the feature. Alternatively, spreading occurs subsegmentally (as in a feature geometric hierarchy), but only segments affiliated with some prosodic unit (a mora, for instance) count as legitimate targets (e.g., Archangeli and Pulleyblank 1994). What such accounts have in common is that the spreading feature is made relevant only to vowels by virtue of a vowel's prosodic status.

(20) Transparency by prosodic anchor

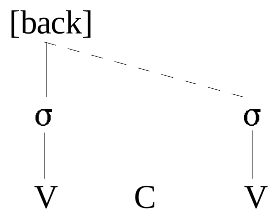

While such accounts differ in their precise claims and explanatory appeal, they share the drawback that long-distance vowel place spreading is rooted in a mechanism of relativised locality otherwise not needed; this mechanism cannot extend in any obvious way to other kinds of long-distance spreading, including that of coronal harmonies. In effect, vowel harmony is stipulated to skip consonants, and is viewed as fundamentally different from other kinds of spreading. In contrast, the account rooted in the bottleneck effect builds directly on substantive properties of segments and syllables that are independently motivated; more generally, we seek a single and unified approach that addresses *all* facts of blocking, participation and 'transparency'.

The line of thinking pursued here may well shed light on the more fundamental question of what features can spread long-distance at all, where by long-distance we mean encompassing CVC or VCV sequences at least. Given strict locality and the consequent reliance on substantive markedness constraints to explain blocking versus participation, the following answer suggests itself: features that spread long-distance must be compatible at least to some degree with both consonants and vowels.[7] We list in (21)a features that spread long-distance, and in (21)b features that do not. This list does not exhaust the inventory of possible features, but covers perhaps the clearest cases.[8]

(21) a. Features that spread long-distance

Vocalic place gestures	Vowel harmonies
Tongue shape/orientation	Coronal harmonies; *r*/*l* assimilation (Cohn 1992)
[nasal]	Nasal harmonies
[aspiration]	Register harmonies
Tone	Tone association, spreading

b. Features that do not spread long-distance

Consonantal place gestures
[voice]

Vowel place features and those of tongue shape or orientation are generally compatible with both consonants and vowels, in the sense pursued above. Consider the other features in (21)a. None of these features contributes in any way to oral constriction degree. More generally, [nasal] and [aspiration] are clearly compatible with both consonants and vowels, since they can appear contrastively on either in some languages. This is of course not to say that all segments will be equally permissive of spreading. In nasal harmony systems there is a clear hierarchy of segments in terms of their likelihood to participate (see section 4.4). Similarly for [aspiration], Thurgood (1996) finds in a discussion of Chamic register harmonies that *s* and *h* are more likely to allow register spreading than are other consonants. (Register is manifested as a complex of breathiness, tone and vowel height.) This is not surprising if consonants are participants, since these are just the segments having a significant spread glottis component themselves. The behaviour of tone is less clear in these terms. The conclusion we are led to by our assumptions is that tone is generally compatible with consonants (as it obviously is with vowels). We take up this point again briefly in section 4.4.

The failure of consonantal gestures to spread long distance, in (21)b, is now a familiar point. As for [voice], the question is why there are no processes like /bat/ → *bad* if this feature is compatible with both consonants and vowels. Though we offer no answer to this last question, it should be borne in mind that the prediction being considered concerns what *can* spread long-distance, not what *must* spread in this way. The hypothesis that long-distance spreading is allowed only for features that are compatible with both consonants and vowels constrains the possibilities, in fact providing a roughly accurate characterisation of the facts. This finding, though in need of further thought and refinement, seems highly significant, and depends on the assumption of strict locality pursued here. So far as we know, the features that spread in this way make up an arbitrary list in the theory otherwise.

4 Markedness, segment realisation, and permeability in spreading

The evidence for segment skipping in spreading might seem compelling: in many cases of long-distance spreading, there are intervening segments that in one way or another do not seem to realise the spreading feature. In the remainder of this paper, our goal is to bring this fact of seemingly pervasive transparency in line with the proposal of strict locality. In a sense, segment skipping is a phonologist's means of controlling segment realisation. We argue here for a much more direct means of effecting this control. A central concern governing our thinking is that any explanation of apparent transparency must preserve the results of section 3 above; this ensures that the account is truly different from those relying on relativised locality. The sections that follow focus in particular on the kinds of apparent transparency listed in (22), with the goal of characterising the core determinants of segment realisation across the cases.

(22) a. Consonants are permeated by vowel place harmony
b. Vowels and non-coronal consonants are permeated by coronal harmony
c. Glottal stops are permeated by nasal harmony

We argue that in all of these cases the relevant segments actually participate in the spreading, *overtly* bearing the spreading feature. That is, we deny the basic fact of transparency, and therefore call such segments *permeable* instead. This stance is possible once we acknowledge bona fide segment types not usually entertained by phonologists. Consider as an illustrative example rounding harmony in Turkish *somun*, which is discussed in greater detail in section 4.2 below. Harmony might appear to involve only the vowels: the intervening labial consonant is not labialised, i.e.*$som^w un$. However, this segment is indeed round, since it overlaps the vocalic round gesture of a neighboring vowel. We have here a segment that will be denoted by $[m]_{Rd}$, phonetically distinct from m^w. The perennial worry in positing new segment types is that we will predict nonexistent contrasts. This does not occur, however, given the right theory of contrast. The theory that correctly distributes these segment realisations, and accounts for patterns of markedness and contrast generally, incorporates insights of the theory of adaptive dispersion (Liljencrants and Lindblom 1972, Lindblom 1986, 1990) into Optimality Theory (Prince and Smolensky 1993), a view developed and motivated by Flemming (1995b, to appear). The general failure to acknowledge segments like $[m]_{Rd}$, and hence the positing of transparency in cases such as Turkish, turns out to be an artifact of an understanding of markedness and segment realisation, long held in phonology, that wrongly equates representational distinctions with *contrastive* distinctions. In Dispersion Theory, on the other hand, potential contrasts must satisfy output constraints that directly regulate the wellformedness of contrast itself; hence the number of actual contrasts is divorced from the issue of how many phonetic distinctions we recognise. The larger claim of this section, therefore, is that an adequate understanding of locality depends fundamentally on a new conception of markedness in phonology.

In explaining permeability by reference to notions of segment realisation, we are inspired by recent work of Smolensky (1993), Cole and Kisseberth (1994, 1995a), Flemming (1995a), Gafos (1996) and Walker (1996). These works understand the realisational idea in very different ways, however. The following sections build on ideas of Flemming (1995b, to appear), extending them to questions of locality. In addition, some similarities to ideas of Browman and Goldstein (1986 et seq.) and work in Firthian Prosodic Phonology emerge. For discussion of other treatments of permeability, and another kind of apparent transparency not treated here (including vowel transparency in harmonies), the reader is referred to the appendix.

4.1 Dispersion Theory in OT

Underlying the problem of segment realisation are very general issues of markedness and contrast. Within one line of work in generative phonology, markedness and contrast are intrinsically linked via constraints on feature cooccurrence: a segment is marked to some degree if it violates a constraint of the form *[F, G, ..., Z], and the activity of such a constraint in the phonology can suppress a potential contrast (Kiparsky 1985, Archangeli and Pulleyblank 1986, 1994). In Optimality Theoretic terms, Prince and Smolensky (1993) recast this idea in the following way: if a constraint *[F, G, ..., Z], which for convenience can be abbreviated *S, where S is the relevant (class of) segment(s), dominates the relevant faithfulness constraints, then a contrast will be suppressed; otherwise, it can emerge. Take as an example the high vowels *i*, *u* and *ɨ*. On typological grounds we might posit the following universal ranking of constraints.

(23) Universal ranking of markedness constraints: high vowels

*ɨ >> *u >> *i

The hypothesised ranking, immutable, is intended to express a typological generalisation: the central unrounded vowel is the least frequent high vowel, while the front unrounded vowel is perhaps the most frequent. Further, a language that has *ɨ* must have *u* and *i* according to this hierarchy. To see this, we must consider what happens when faithfulness is included in the hierarchy. To simplify the discussion we assume one general constraint IDENT requiring identity of feature content between input and output (McCarthy and Prince 1995); this constraint ensures that a posited contrast will surface if the relevant markedness constraints are lower ranked. As Prince and Smolensky (1993) argue, markedness hierarchies coupled with faithfulness provide an appealingly direct and elegant account of markedness implications and contrast. An important result of the Optimality Theoretic account, in addition, is that contrast is an emergent output property, following entirely from the ranking of constraints in the grammar:

(24) Typological predictions of markedness and faithfulness

	Ranking	*Result*
a.	IDENT >> *ɨ >> *u >> *i	ɨ, u, i surface
b.	*ɨ >> IDENT >> *u >> *i	u, i surface
c.	*ɨ >> *u >> IDENT >> *i	i surfaces
d.	*ɨ >> *u >> *i >> IDENT	No high vowels surface

This Optimality Theoretic account for the high vowel typology shares one important drawback with all phonological approaches to markedness known to us. While systems such as (24)a and (24)b are uncontroversially attested (and we put aside (24)d as not germane here), the prediction entailed by (24)c is simply incorrect: while languages with two vowels overwhelmingly prefer *i* and *u*, those with only one high vowel select *ɨ*. (The 'linear' vowel systems of the Caucasus such as Abkhaz and Kabardian provide well-known examples, Trubetzkoy 1939; see Ladefoged and Maddieson 1996:286-7 and Flemming 1995b for other examples.) The hierarchy in (24) specifying *ɨ* as the worst vowel is fatally incapable of capturing this fact; instead it predicts a consistent preference for the front vowel *i*, and saving that, *u*. This problem is obviously general to any theory that uniformly favours *i* and *u* over *ɨ*. Flemming (1995b) makes this point, and uses it to argue for Dispersion Theory (henceforth DT). DT adapts a well-known phonetic explanation for this markedness dichotomy from the theory of adaptive dispersion. According to Lindblom, inventories strike a balance between two often contradictory needs. There is a tendency to maximise the perceptual distinctiveness (dispersion) of contrasts; however, there is also a need for articulations to be minimally complex. Given two high vowels in a language, perceptual considerations demand that they be *i* and *u,* since these are maximally distinct; achieving this distinction comes at the cost of articulating these particular vowels. On the other hand, a language with no high vowel contrast by definition makes no demands of distinctiveness or dispersion in the front-back dimension. In such a language, articulatory concerns carry the day, and the result is *ɨ*: This vowel is articulatorily simpler, involving the least displacement of the vocal tract configuration from the neutral position of *ə*.

A crucial property distinguishing Dispersion Theory from other approaches to markedness is what we call its *bidimensionality*. A basic claim of the theory is that there are separate families of constraints regulating articulatory simplicity on the one hand and perceptual distinctiveness on the other, and that these often conflict. Contrast this with the prevalent *unidimensional* understanding of markedness exemplified in (24): here segments are ranked along only one universal scale. The problem for any unidimensional theory is that it faces paradoxes like the one discussed above. On the one hand, *ɨ* seems highly marked with respect to *i* and *u*, since it occurs in fewer languages, and since the appearance of *ɨ* implies the appearance of the other two—assuming there is a contrast among high vowels. On the other hand, *ɨ* seems *unmarked* with respect to the other two vowels, since systems with no contrast among high vowels always prefer *ɨ*.

As Flemming notes, since the dispersion idea posits intrinsically conflicting tendencies on inventory structure, it receives a potentially natural and rigorous implementation in OT. Consider then the two dimensions of markedness given in (25) and (26). The first is articulatory markedness, which we will largely represent in terms of a universal ranking among segments, following Prince and Smolensky (1993). However, the criterion of articulatory simplicity requires a different ranking from what was seen above, with *ɨ* at the bottom. From the perspective of articulation, this is the least marked high vowel.

(25) Articulatory markedness

a. *u >> *i >> *ɨ

More striking are constraints of the perceptual dimension. The essential insight of DT is that markedness cannot be reckoned on a segment-internal basis alone (as (25) reckons it), but must evaluate *contrasts* also. *ɨ* is disfavoured in comparison to *i* or *u* precisely when there is contrast among high vowels, because a contrast between *i* and *u* is more perceptually distinct. It is not that either *i* or *u* is individually less marked than *ɨ*; rather, the *contrast* between *i* and *u* is unmarked. Hence the constraint shown in (26)a, which as stated requires two things: that there be a contrast in vowel colour (backness and/or rounding), and that it be a good one in the sense of respecting a *minimal distance* requirement. Suppose that by 'sufficient' minimal distance we mean larger than the distance between *ɨ* and either of the other vowels, as informally illustrated in (26)b. Then a two-way contrast must be between *i* and *u* to satisfy this constraint. It should be noted that the CONTRAST constraints we formulate conflate two separate families of constraints employed in Flemming (1995b, to appear), one requiring that there be contrasts (of varying numbers) and another imposing a range of minimal distance requirements. For example, since there are vowel systems with all three high vowels, there cannot be one universal minimal distance demand in the case of vowel colour. (Moreover, minimal distance demands themselves conflict with constraints preferring more contrasts.) Since every case we examine below in the discussion of locality involves at most a two-way contrast, however, and hence a single universal minimal distance threshold for contrast, we can conflate the two independent constraint types in this way. This is intended as an expository simplification only, one that does not sacrifice any explanatory depth, or alter the basic predictions of the theory.

(26) Perceptual markedness

a. CONTRAST(COLOUR): Maintain a colour contrast with sufficient perceptual distance.

b. Minimal distance: i........ɨ........u
|________|

Before considering larger implications of these notions, let us see how they work in the case of this high vowel example. The interaction of the two classes of constraints—the articulatory markedness constraints and the perceptual markedness constraints—results in two possible scenarios. On the one hand, high ranking CONTRAST can favour articulatory complexity, as in (27). Note here how CONTRAST evaluates sets of segments, that is inventories, along some dimension of contrast, not single forms. We use the notation i • u, for example, to denote a contrast between these three segments. This manner of evaluation is dictated by the central generalisation that markedness on the perceptual side evaluates contrast, and not merely segments.

(27) Contrast can force articulatorily marked realisations

	Contrast(Colour)	*u	*i	*ɨ
a. i • ɨ	*!		*	*
b. ɨ • u	*!	*		*
c. ☞ i • u		*	*	
d. i	*!		*	
e. ɨ	*!			*
f. u	*!	*		

CONTRAST(COLOUR) is violated by inventories attempting no contrast at all, (27)d-f, and by inventories that fail to satisfy the minimal distance requirement, (27)a-b. The only remaining candidate, and the winner, is the two-member inventory i • u. DT shares an interesting property with Prince and Smolensky's faithfulness-based understanding of contrast: contrasts are a property of the output.

On the other hand, articulatory markedness can suppress contrast, resulting in articulatory simplicity. In (28), the least marked vowel, articulatorily, is favoured. Note that the same result is obtained by any ranking in which CONTRAST(COLOUR) is dominated by an articulatory markedness constraint.

(28) Articulatory markedness can suppress contrast

	*u	Contrast(Colour)	*i	*ɨ
a. i •ɨ		*	*!	*
b. ɨ •u	*!	*		*
c. i • u	*!		*	
d. i		*	*!	
e. ☞ ɨ		*		*
f. u	*!	*		

By positing two often conflicting dimensions of markedness, DT resolves the paradox facing unidimensional markedness discussed above. It does this by acknowledging the wellformedness of contrast as an issue. The case of the high vowels illustrates what we might call the *dispersion syndrome*, in which a segment that is marked under conditions of contrast becomes unmarked in the absence of contrast. This example is far from being isolated. Factoring in the height dimension for vowels, for example, we note the same paradox confronting unidimensional theories in accounting for ə. Though no language eliminates all vowel place contrasts, having only this phonemic vowel, ə is the reduced vowel par excellence across languages in metrically weak positions. Assuming that neutralisation diagnoses what is least marked, ə would seem to be the best vowel. It can only be seen as an embarrassment to unidimensional markedness, therefore, that ə is so highly marked when contrast is at stake (i.e., in phonemic inventories). The usual response to this contradiction is to treat ə as special in some way, either as 'placeless', 'cold', or by some other means. These ad hoc attributions only obscure the central markedness generalisation: ə is bad for contrast because it is not peripheral (minimal distance); when contrast is not at stake, ə is the best vowel possible, because it is articulatorily least marked. Flemming (1995b) provides many more examples of the dispersion syndrome at play, to which we refer the reader.

We briefly touch on several more points in this introduction to DT. First, the preceding discussion of the contrast-based dimension of markedness assumes that it is perceptual contrast that matters. In principle one might suppose that contrast in articulation is the goal. The former view grounds contrast evaluation in the needs of the listener, while the latter seems speaker directed. In fact all evidence supports the view that contrast evaluation is grounded in perception. Though sounds can be made articulatorily distinct by many means, the means actually selected in language are just those that enhance perceptual differences. (See Kingston and Diehl 1994 for a similar point with reference to phonetic implementation.) The contrast between *i* and *u* illustrates this point. To have two high vowels that are articulatorily distinct, one might just as well have a rounded front vowel and an unrounded back one (spread lips). The fact that languages select *i* and *u*, rather than *y* and *ɯ*, receives a perceptual explanation (see Stevens, Keyser and Kawasaki 1986, Flemming 1995b, and references therein). Hence the term *colour*, denoting the auditory dimension implicated by changes in [back] and [round], essentially variation in the second and third formant frequencies. (For discussion of vowel colour with further references, and a feature class characterisation uniting [back] and [round], see Padgett 1995.) Other cases explored in the following sections also make this point; the failure of the nasalised glottal stop seen in section 4.4, for example, only makes sense as a perceptual failure.

A second point concerns the status of faithfulness, and of underlying representations, in DT. As was noted above, one of the fundamental roles of faithfulness constraints in OT is to ensure (or dispel) contrast. Yet DT posits constraints that demand contrast directly in the output, entirely taking over this role for faithfulness. Flemming (1995b) in fact argues that faithfulness, and underlying representations themselves should be eliminated from the theory. To address the other major role of underlying forms, the encoding of morphological relatedness, he suggests that similarity among morpheme alternants be dealt with entirely by constraints governing the similarity of *surface* forms. Similar ideas to this latter one resonate in an increasing number of works arguing for constraints on the identity of output forms, or other means of maintaining surface similarities, e.g. Benua (1995), Buckley (1995), Burzio (1994, 1996), Itô and Mester (to appear-a), Kager (1996), Kenstowicz (1996), McCarthy (1995), cf. Orgun (1994, 1996), Kenstowicz (1995). For the sake of discussion in this paper, we follow Flemming, as well as Burzio (1996), in requiring no underlying representations for the forms we consider, at least as crucial determinants of output wellformedness. Given the fact that in OT contrast is an emergent output property, and given alternative means of handling morpheme alternants, this move is not as subversive as it might seem. For the same reason we employ no faithfulness constraints, instead letting the CONTRAST constraints do the relevant work.

Finally, an important research goal for DT will be making more precise the statement of constraints on contrast and minimal distance, and the manner in which such constraints are evaluated. Though the direct evaluation of contrast requires a kind of globality that might seem daunting at first—since candidates are not simply forms, but sets of forms—we take the view eloquently expressed by Prince and Smolensky (1993) that theoretical developments cannot be constrained by a priori computational assumptions.[9] This point is all the more forceful when there is clear and wide-ranging empirical support for the relevant ideas; this is the case with DT. The issues raised here and in the preceding paragraph are too far-reaching to receive any definitive attention within the perimeters of this article. Our goal in what follows is more modest: we argue that DT essentially resolves the problem of locality for many cases in phonology, bringing an explanatory coherence to the area not possible before. This argument brings a new range of facts into the purview of DT, and therefore provides more motivation for it. In doing so, it reduces apparently parochial issues of locality and transparency to more basic issues of markedness, independently required, essentially obviating altogether a 'theory of locality'. Our hope, then, is to convince the reader that issues like those raised above are worth taking on.

4.2 Vowel harmony and consonant permeability

We turn now to what is probably the most commonly assumed case of transparency in a spreading domain—consonant transparency in vowel harmony. Since the account pursued focuses on questions of realisation in such domains, it is necessary to explore first the question of the realisation of consonants bearing vowel features more generally. An important observation is that even where the vowel features on consonants are contrastive, the realisation of such consonants can be variable, reflecting the interaction of constraints on contrast and constraints on articulatory markedness. Contrasts, in other words, have variable realisations predicted by DT.

4.2.1 Vocalic place features on consonants

Consider the three-way phonetic contrast shown in (29) below, between CV sequences containing a palatalised consonant, a velarised consonant, and a 'plain' consonant. Given the fact of coproduction of consonants and vowels discussed in section 3, the designation 'plain' generally means that the consonant shares (or overlaps) vowel place properties of the nuclear vowel. This coproduction is represented as $[CV]_V$. The other two cases are notated as $[C]_jV$ and $[C]_ɣV$ respectively—in just such cases coproduction does *not* obtain—though we could have retained the familiar superscript notation C^j and $C^ɣ$ instead.

(29) $[C]_j$a $[C]_ɣ$a $[Ca]_a$

The first point of interest here is that this three-way phonetic contrast does not translate into a three-way *phonemic* contrast; backness contrasts are at most binary, hence the traditional single, binary feature [back]. A similar point might be made for other traditional binary features. For example, though the phonetic degree of nasalisation can in principle be varied continuously, phonological theory has recognised only a single, binary [nasal] contrast, with good empirical reason.[10] On the other hand, vowel height contrasts are more numerous, requiring more features (or feature values) in the theory. The notion *minimal distance*, introduced above, makes explicit what is implicit in familiar feature systems (Flemming 1995b, to appear): in some phonetic dimensions, such as nasality or backness in consonants, there is room for only one contrast; in others, such as vowel height, there is room for more. The notion minimal distance is potentially more informative, however. Consider another fact about backness in consonants: it is well known that languages vary in the precise implementation of this contrast, drawing on a small variety of phonetic realisations. Thus, some contrasts are described as plain vs. palatalised, others as plain vs. velarised. Consider therefore (30), which is meant to visually clarify these two properties of backness contrasts in consonants: on the one hand, the minimal distance required for such a contrast is large enough to exclude a three-way contrast ('only one binary/unary feature'); on the other hand, it is not so large as to require always that the contrast be realised via the extremes of palatalised versus velarised consonants. This minimal distance is indicated with respect to six illustrative CV sequences: palatalised, velarised and plain (coproduced) consonants, each before either a front or back vowel. These sequences are arranged in a manner to suggest their relative similarities; hence the coproduced and palatalised C's before *i* are perceptually very close together, while both are quite distant from a velarised C in this context. The facts are different when the vowel is *u*; now it is the coproduced and velarised consonants that are very similar (indeed, given the rounding of *u* added in, virtually indistinguishable).[11]

(30) Minimal distance (MD) for a backness contrast in consonants

'plain vs. velarised'

| MD |

[C]$_j$i....[Ci]$_i$..........................[C]$_ɣ$i
[C]$_j$u.......................[Cu]$_u$....[C]$_ɣ$u

|______________|

MD

'palatalised vs. plain'

As the diagram suggests, we posit that the minimal distance for contrast cannot be met whenever the opposing consonants share the same backness specification. That is, neither [C]$_j$i versus [Ci]$_i$, nor [C]$_ɣ$u versus [Cu]$_u$ is possible. Let us assume that the same is true before mid vowels: neither [C]$_j$e versus [Ce]$_e$, nor [C]$_ɣ$o versus [Co]$_o$ occurs.[12] Given this state of affairs, the array of realisations implied by the placement of the MD spans above is an optimal one for such a contrast, from the bidimensional viewpoint of DT. Taking up the perceptual dimension first: before *i*, the palatalised vs. coproduced distinction is too slight to support a contrast; this context dictates that velarisation be brought in. Before *u*, on the other hand, palatalisation is required to maintain a contrast. With these needs of contrast met, articulatory markedness can have its say, allowing the plain, coproduced realisations to oppose the more complex ones. In this light, consider the distributional pattern of the (distinctive) backness contrast in consonants in Irish, best observed in environments that contain long vowels, since short vowels in Irish acquire their backness specification from neighbouring consonants. The contrast is realised as follows in open syllables containing long vowels: a secondary articulation is pronounced when the consonant is followed by a long vowel that carries the opposite specification for the feature [back], i.e. a [-back] labial consonant is realised with a palatal off-glide preceding a long high back vowel (31)a, while a [+back] labial consonant is realised with a velar off-glide preceding a long high front vowel (31)c. Off-glides are not produced when these consonants are followed by long high vowels of the same backness specification (31)b,d. Similar facts involving consonant-mid vowel sequences are given in (31)e-h.

(31) C-high vowel sequences in Irish: C-mid vowel sequences in Irish:

a.	f^{j}u:	'worth'	e.	b^{j}o:	'alive'
b.	fu:ə	'hate'	f.	bo:	'cow'
c.	b^{ɣ}i:	'yellow'	g.	b^{ɣ}e:l/b^{ɣ}i:l	'danger'
d.	bi:	'be' IMPER	h.	be:l	'mouth'

With contrast regulated by output constraints of the CONTRAST family, a backness contrast in consonants is the result of a constraint CONTRAST(C-BK) (we generalise over all consonants for the sake of discussion).

(32) CONTRAST(C-BACK): Maintain a front/back contrast in consonants, with sufficient distance

Consider then the facts of articulatory complexity. We take the standard view that $[C]_j$ and $[C]_ɣ$ are articulatorily more complex than coproduced ('plain') consonants. Considering the Irish CV sequences in (31) above, relevant constraints might be as in (33).

(33) $*[C]_j$ >> $*[C]_i$, $*[C]_e$
$*[C]_ɣ$ >> $*[C]_u$, $*[C]_o$

However, given the ubiquitousness of $[C]_V$ (coproduction), it is not clear that constraints such as those dominated in (33) should be distinguished at all (at least for most consonant-vowel combinations). That is, coproduction causes little or no articulatory difficulty, and constraints such as $*[C]_j$ and $*[C]_ɣ$ are enough. It is also convenient to generalise over these latter constraints; we will recast them in terms of a contextual markedness constraint COPRODUCE, requiring coproduction of consonants and tautosyllabic vowels. (The following is essentially an alignment constraint, i.e., Align-L/R(vplace, σ); see section 4.2.2.)

(34) COPRODUCE: Every vowel place feature is associated with the leftmost/rightmost segment of the syllable

Syllables like *b^jo* 'alive' violate COPRODUCE, since the features of *o* do not extend to the left edge of the syllable. Since we are phonetically distinguishing $[C]_j$i and $[Ci]_i$ (see note 11), the former violates COPRODUCE also. Entertaining this distinction might appear contrary to the widely accepted view that high vowels and the corresponding glides are identical in featural make-up, e.g. Kaye & Lowenstamm (1984), Selkirk (1984), Levin (1985). The basis of the latter view, however, is that the relevant segments cannot *contrast*. In our terms, this fact should fall out from a minimal distance requirement similar to the one posited above.

Returning then to the case at hand, in order for a backness contrast in consonants to be possible, the constraint CONTRAST(C-BK) must dominate the constraint requiring coproduction.

(35) Before back vowels: 'plain' vs. palatalised

	Contrast(C-bk)	Coproduce
a. $[f]_j$u: • $[f]_ɣ$u:		**!
b. ☞ $[f]_j$u: • $[fu:]_u$		*
c. $[f]_ɣ$u: • $[fu:]_u$	*!	*
d. $[f]_j$u:	*!	*
e. $[f]_ɣ$u:	*!	*
f. $[fu:]_u$	*!	

CONTRAST(C-BK) is trivially violated in (35)d-f, which lack any contrast. The contrast in (35)c falls short of the minimal distance required to satisfy CONTRAST(C-BK) (as discussed above). With contrast satisfied, the two remaining candidates are evaluated as to articulatory difficulty. Of these, (35)b maintains contrast while involving least articulatory effort. Thus preceding a back vowel, the optimal contrast is a palatalised-plain one.[13] Preceding a front vowel, on the other hand, the optimal contrast is a velarised-plain one, as seen in (36) where the forms contain a front vowel. Single form candidates as in (35)d-f are omitted here.

(36) Before front vowels: 'plain' vs. velarised

	Contrast(C-bk)	Coproduce
a. $[b]_j$i: • $[b]_ɣ$i:		**!
b. $[b]_j$i: • $[bi:]_i$	*!	*
c. ☞ $[b]_ɣ$i: • $[bi:]_i$		*

As these tableaux show, DT's bidimensional approach to markedness predicts exactly the kind of variability in realisation seen in the Irish case above: articulatory complexity is forced where necessary to fulfill the minimal distance requirement on contrast, giving one of C^j or $C^ɣ$, depending on the vocalic context. Articulatory simplicity determines the remaining realisation. The result is that coproduced $[C]_V$ ('plain' C) can enter into some contrasts. The emergence of the coproduced realisation of consonants in these circumstances mirrors what is found in vowel harmony domains, as discussed below for Turkish: once the needs of contrast are either met or otherwise rendered irrelevant, segment realisation is determined by articulatory markedness constraints.

4.2.2 Consonants in vowel harmony domains

In the discussion of apparent consonant transparency in vowel harmony domains that follows, the idea of segment realisation as determined by constraint interaction is once more pivotal. In this instance realisation is determined by the interaction of constraints requiring feature spreading, and articulatory markedness constraints; significantly, the requirements of contrast are overridden in spreading domains and so do not play a role in determining the realisation of the segments in question. The empirical focus of the discussion are the well known facts of [back] and [round] harmony in Turkish.

Vowels in Turkish agree in backness with the preceding vowel. A high vowel in addition agrees in roundness. In most analyses the features spread from the first syllable of the root (which is also the first syllable of the word, since the language is suffixing). In what follows we confine our attention to those cases where both [back] and [round] spread, that is, to forms containing high target vowels.

(37) Turkish vowel harmony: [back] and [round] spread to target [+high] vowels

somun	'loaf'	iji	'good'
son-un	'end, gen.'	ip-in	'rope, gen.'
öküz	'ox'	adɨm	'step'
jüzün	'face, gen.'	sap-ɨn	'stalk, gen.'

For the facts and basic analysis of Turkish we rely on Lees (1961), Clements and Sezer (1982), van der Hulst and van de Weijer (1991), Kirchner (1993), and references therein.
We assume that spreading is compelled by constraints that have the general form shown below. These constraints require that a feature align with either a left or a right word edge; in conjunction with locality constraints, satisfaction of alignment often results in long-distance feature spreading. (See for example Kirchner 1993, Smolensky 1993, Pulleyblank 1994b and Cole and Kisseberth 1994 on alignment for this purpose, and McCarthy and Prince 1993a on the general notion of alignment.) Following the references cited, alignment is taken here to be gradiently violable, such that one violation is tallied for every segment intervening between the anchor segment and the word edge. (See Zoll 1996 for extensive discussion of this aspect of alignment and its formalisation.)[14]

(38) Featural alignment constraint schema for harmony

ALIGN-L/R(F, PWD): Every feature F is associated with the leftmost/rightmost segment of the prosodic word

To effect harmony in Turkish, we assume the constraint shown in (39). The right-edge alignment mimics to a large extent the rightward spreading posited in many analyses (see Anderson 1980 for argumentation in favour of rightward spreading). Following Padgett (1995) a unified statement of

[back] and [round], or colour, harmony is assumed, though this point is orthogonal to the issues explored here.[15]

(39) [back]/[round] harmony in Turkish, assuming *Colour* = {back, round}

ALIGN-R(COLOUR, PWD): Every feature F, F ∈ *Colour*, is associated with the rightmost segment of the prosodic word

An obvious consequence of harmony is the neutralisation of contrast. In Turkish, it is colour contrasts in non-initial syllables that suffer. The loss of contrast, generally understood as faithfulness or correspondence constraint violations in OT, instead violates constraints directly demanding output contrast in DT, as we have seen. Recall the constraint CONTRAST(COLOUR) from section 4.1, requiring a contrast in colour for vowels. This constraint is necessarily violated as a result of harmony, requiring the ranking shown in (40). Candidates, as before, are sets of potentially contrasting forms. Here we limit our attention to forms with unrounded high vowels. Brackets in the tableau indicate [back] harmony spans.

(40) Basic [back] harmony

	Align-R	Contrast(Colour)
a. ☞ [ɨp-ɨn] • [ip-in]		*
b. [ɨ]p-in • [i]p-ɨn	*!**	
c. [ɨ]p-ɨn • [i]p-ɨn	*!**	*

Candidate (40)a respects alignment and therefore violates the contrast constraint: though the non-initial vowels are different, they are not in contrast. (40)b violates alignment while respecting contrast in non-initial syllables. Given the ranking of constraints, this form loses. (40)c violates alignment, while also violating the contrast constraint by means other than harmony: the non-initial vowel is reduced to *ɨ*. The grammar therefore selects (40)a. This account of the general harmony facts is hardly complete, and leaves open some questions about the precise workings of constraints on contrast. (See the relevant discussion at the end of section 4.1.) The object is simply to outline an understanding of harmony from which the discussion of consonant permeability can proceed, and to make the connection between harmony and contrast neutralisation.[16]

Turning then to the locality question, we note that strict locality directly entails segments we can notate as $[p]_{-bk}$, and $[p]_{+bk}$, occurring in (40)a. Assuming nothing opposes coproduction more generally, we have, more specifically, $[p]_{i}$ and $[p]_{ɨ}$. (Coproduction also dictates that the *n* in these examples be a participant; a separate question which we leave open is whether or not alignment should have this effect.) Yet phonology has generally concluded that consonants are *not* participants in vowel harmony. For any unidimensional understanding of markedness, to conclude the opposite simply leads to a paradox, since representational distinctions are generally identified with *contrastive* distinctions. This paradox can be approached from either of two directions.

First, suppose consonants are indeed participants; then segments like $[p]_i$ must be rather unmarked, since consonants are almost *always* participants (barring cases where they have their own secondary vocalic specifications). Yet if $[p]_i$ is unmarked (or even exists), why does it not occur contrastively, e.g., $[p]_i$ versus $[p]_j$? From the perspective of OT, for example, the problem can be set out as in (41). *p* will participate in palatal harmony given a ranking of alignment over the relevant markedness constraint, as shown in (41)a. To predict that $[p]_i$ could surface contrastively, all we need assume is that faithfulness or identity could also outrank this markedness constraint, (41)b. It is difficult to imagine what could prevent the latter ranking: given the pervasive occurrence of $[p]_i$ in languages via coproduction, the constraint $*[p]_i$ is surely low-ranked on any universal scale of markedness. This problem is of course not particular to OT, but arises under all current conceptions of markedness, which share unidimensionality as a hallmark.

(41) Unmarked → contrastive for unidimensional markedness

a.	ALIGN-R	>> $*[p]_i$	*$[p]_i$ is generated by harmony*
b.	IDENT	>> $*[p]_i$	*$[p]_i$ is contrastive*

The second way to approach the paradox is through languages like Irish or Russian, in which secondary vocalic articulations on consonants are indeed contrastive. From such cases it is frequently concluded that contrastive palatalisation on a consonant involves a *j*-like off-glide (putting aside possible further consequences, such as a change to palato-alveolar place for coronals); similarly, contrastive rounding is taken to involve a *w*-like off-glide. Note that these expectations are oversimplified, given the manifestation of the contrast in Irish seen above, in which 'plain' consonants can enter into the contrast so long as minimal distance is maintained. Ignoring this issue, the expectation derived from such observations of realisations in contrast is that [-back] harmony in a form like *ipin* 'rope (gen.)', together with strict locality, should entail [ip^jin^j], since there is no other notion of 'frontness' in a consonant available to the theory. Similarly, [round] harmony in *somun* 'rope' should give something like [$s^wom^wun^w$]. Since the Turkish forms do not actually warrant transcriptions of this sort, the conclusion this reasoning leads to is that consonants do not acquire these spreading features. This is of course the conclusion reached, usually implicitly, for vowel harmonies in general in both the SPE and autosegmental traditions. This conclusion is at the heart of the problem addressed in section 3: it forces a formulation of locality that is relativised (in this case to vowels), geometrically or otherwise, and it thereby undermines a genuine understanding of the spreading asymmetry between consonant and vowel place, which rests on strict locality.

Approached from either direction, the source of the paradox above lies in the assumption that potentially contrastive distinctions are all and only the distinctions known to phonology. The way out is to cease this identification, and to acknowledge the extra realisational possibilities this move provides. The consonants of Turkish participate in harmony if we understand this statement in a somewhat obvious way: the consonants of *somun* are coproduced with vocalic lip rounding, and those of *ipin* with tongue fronting. That is, the consonants have exactly the properties we would expect of them if lip rounding were maintained throughout *somun* and tongue fronting throughout *ipin* respectively, and nothing more. Indeed, the most straightforward interpretation of a harmony span is as a gesture maintained, uniformly and uninterrupted, for the relevant duration (see section 2.1 for discussion), and Turkish vowel gestures seem to bear this assumption out. Working within the Firthian framework, Waterson (1956) posited uninterrupted vowel prosodies, e.g. $\underline{\text{somun}}_u$, $\underline{\text{ipin}}_i$ (though Waterson used superscripted lines to notate them). Taking this to have implications for the consonants as well as the vowels, she took palatograms of various Turkish consonants showing that they are articulated with a fronted tongue body in words such as *ipin*. More recently, Boyce (1990) finds that lip rounding gestures in nonsense words like *utku* are maintained in a more or less uniform 'plateau' across the consonants in Turkish, while in English there are two gestures clearly distinguished with a period of diminished rounding during the consonants. (Making a similar argument to ours, Gafos 1996 also notes the significance of these studies.) These phonetic observations are perhaps not particularly surprising. The point is that they qualify as phonological observations as well, once we broaden the range of realisational distinctions available to the theory, a move requiring a deeper understanding of contrast.

Consider again the case of [back] in Turkish. Since there is no backness contrast in Turkish consonants (with a limited exception, see below), we might conclude that the ranking COPRODUCE >> CONTRAST(C-BK) holds. (Notice that this is the opposite ranking to the Irish case discussed in the previous section.) However, there is another reason why a backness contrast cannot obtain in Turkish: harmony itself conflicts with it. Just as harmony necessarily neutralises colour contrasts in vowels, as we saw above, it forces neutralisations for secondary colour in consonants, as shown in (42). Here and throughout we show only potential contrasts relevant to the contrast constraint at hand; hence the contrasts displayed here are based on consonantal backness. Candidate (42)a attempts a contrast between plain and velarised *p*. The velarised consonant blocks spreading of [-back], since [-back] and [+back] specifications on a single segment are incompatible (see section 2.1); we assume that $p^ɣ$ spreads its own [+back] value rightwards. (See the analogous spreading of [-back] from palatalised velars below.) Alignment is nevertheless violated due to the failure of [-back] to reach the right word edge in forms with velarisation. Since this candidate loses, alignment must dominate the consonantal contrast constraint. (42)b is a more interesting attempt at a contrast, this time between p^j and $[p]_i$; that is, a palatalisation contrast is attempted between words that both respect [-back] alignment fully. This attempt is possible since palatalisation does not conflict with the [-back] harmony requirement. Yet contrast again fails here, since the minimal distance requirement posited in the previous section is not met. Candidates (42)c and (42)d have no backness contrast in consonants, while respecting alignment. They are distinguished by a gratuitous coproduction violation in the case of (42)c. Summing up, we see here the dispersion syndrome: a backness contrast between consonants is impossible (due to both alignment and minimal distance requirements); what determines the consonantal realisations then is articulatory complexity alone, and coproduction holds.

(42) Harmony neutralises consonantal backness contrasts

	Align-R	Contrast(C-bk)	Coproduce
a. $[i]_{-bk}[p^{ɣ}ɨn]_{+bk}$ • $[ipin]_{-bk}$	*!**		*
b. $[ip^{j}in]_{-bk}$ • $[ipin]_{-bk}$		*	*!
c. $[ip^{j}in]_{-bk}$		*	*!
d. ☞ $[ipin]_{-bk}$		*	

From DT's bidimensional perspective, every posited segment must be evaluated with respect to two separate markedness dimensions, grounded respectively in articulatory and perceptual facts, as we have seen. For the case at hand this analysis was largely carried out in the last section. On the articulatory side, segments like $[p]_i$ are very unmarked, in fact occurring by coproduction in languages routinely. Since harmony is exactly a scenario in which only articulatory complexity can matter, as shown above, we predict correctly that these unmarked coproduced segments will result. On the side of perceptual distinctiveness, what is *marked* are contrasts such as *pi* versus *p*^{j}*i*. Since DT treats this fact separately, it avoids the paradox seen above. The full parallel here with the facts of *ɨ* discussed in section 4.1 should be clear: a unidimensional markedness paradox resolves itself once the two genuine dimensions of markedness are recognised. Stepping back, we have the larger conclusion: segment skipping is an artifact of an incorrect approach to markedness. Given a better markedness theory, the issue of locality actually evaporates.

It is true that DT entails some cost: the theory must recognise more segment types than is customary and it must posit constraints that regulate them. Why is this preferable to some version of relativised locality and segment skipping? First, DT is required independently of locality issues, simply to achieve an adequate explanation of markedness facts in phonology. This was the point of section 4.1, and our belief is that further research will bear this claim out. Second, several considerations mitigate concern about introducing more segment types. Most important, they do not in themselves entail nonexistent contrasts, since contrast must itself be regulated, as we have seen. In addition, while they might be new segments to phonology, they are phonetically true distinctions (categorically expressed, as usual for phonology), and in the larger picture of a full theory of sound structure, every theory must generate them. Third, the account here preserves the explanation for the spreading asymmetry between consonant and vowel place features seen in section 3. It does this because it preserves segmentally strict locality. Relativised locality does not genuinely succeed on this critical point, as we saw. Finally, the DT approach truly unifies seemingly diverse cases of apparent transparency. On this point also, relativised locality fails badly. This difference between the two approaches becomes clear as we look at further cases.

The constraint ranking ALIGN >> CONTRAST thus results in consonants participating in the spreading, rendering the notion 'transparency' obsolete. The reverse ranking—CONTRAST >> ALIGN—on the other hand, has the result that contrastiveness in consonants can take precedence over spreading, yielding instances of blocking. For completeness this point is illustrated below with data involving palatalisation contrasts in velars in Turkish. Turkish distinguishes plain and palatalised velars; the distribution of these variants is largely predictable, but is contrastive in a limited way also. Most occurrences of k^j and g^j (which may also be transcribed as *c* and *ɟ*) are contextually determined in a simple manner: velars are palatalised when released in a syllable containing a front vowel (Clements and Sezer 1982; see Orgun 1993 on this statement). In addition, however, palatalisation is contrastive for (disharmonic) stems ending in *aK*, where *K* is a velar. This contrastive palatalisation manifests itself when a vowel-initial suffix is attached; otherwise it is neutralised, as shown below. These palatalised velars block the spreading of [+back] from the preceding vowel and instead trigger front harmony themselves, also shown.[17]

(43) Contrastive palatalisation

idrak	idra:k^ji	'perception' nom/acc sg
emlak	eml^ja:k^ji	'real estate' nom/acc sg

The forms above contrast with velar-final stems in which the final velar does not alternate and where backness spreading is not blocked. A constraint sanctioning this limited contrast in velars is stated for expedience in (44), leaving a genuine understanding of the stem-final condition to other work. Also put aside is an analysis of the neutralisation facts seen above; this would suggest an appeal to preferential contrast preservation in released positions (see Beckman to appear on positional contrast, and Padgett to appear on the relevant release notion).

(44) STEMFINALCONSTRAST : Maintain a backness contrast in stem-final velars

Since this constraint requires palatalised velars, it must outrank the articulatory markedness constraint: STEMFINALCONSTRAST (SFC) >> COPRODUCE. Further, since the palatalised velar blocks the spreading of [back] from the final stem vowel, SFC must also outrank the alignment constraint requiring harmony. As we saw earlier, the contrast could not exist otherwise.

(45) C^j blocks [back] harmony: SFC >> ALIGN

	SFC	Align	Coproduce
a. ☞ a:][k^ji] • a:kɨ]		**	*
b. a:kɨ]	*!		
c. a:][k^ji]	*!	**	*

In (45) we consider only partial hypothetical forms, focussing on the stem-final palatalisation facts and the potential spreading of [back] from the final stem vowel. Candidate (45)a above is one in which a contrasting pair of velar-final stems is considered. Were spreading successful in *all* velar-final stems, the palatalisation contrast would in all cases be neutralised, as in (45)b, violating SFC.

Thus maintaining contrast, and therefore the more marked realisation, takes precedence in this instance over harmony and yields an instance of blocking. More generally, this further exemplifies points made in section 2.1: given strict locality, blocking is always simply a matter of markedness considerations—in the shape of articulatory- and contrast-based constraints— with no appeal to line crossing necessary.

4.3 Coronal harmonies and vowel/non-coronal permeability

We are now in a better position to complete the discussion of coronal harmonies, focusing on Sanskrit, begun in section 3. The Sanskrit Nati data examined there are repeated here for convenience. Recall that Nati is the rightward spreading of retroflexion from a continuant (*r* or *ʂ*) to *n*; it applies both under adjacency and long-distance, as shown in (46)a-b; it is blocked by intervening coronals, (46)c. A subscript R denotes retroflexion.

(46) Retroflexion harmony in Sanskrit

a.	i$[ʂɳ]_R$a:	'seek (present)'
	p$[rɳ]_R$a:	'fill (present)'
b.	pu$[ra:ɳ]_R$a	'fill (middle participle)'
	k$[ʂub^ha:ɳ]_R$a	'quake (middle participle)'
	čak$[ʂa:ɳ]_R$a	'see (middle participle)'
	k$[rpama:ɳ]_R$a	'lament (middle participle)'
c.	ma$[r]_R$ǰa:na	'wipe (middle participle)'
	k$[ʂ]_R$veda:na	'hum (middle participle)'

There are several challenges to the formulation of the spreading constraint. First, there are conditions on both the triggers (continuants) and the target (dental nasal). Second, in words with two or more dental nasals to the right of *r/ʂ*, retroflexion spreads only to the first one. Thus the form /varn+ana:na:m/ surfaces as *va$[rɳ]_R$ana:na:m*, and not **va$[rɳaɳa:ɳ]_R$a:m*. In fact, Nati is a rather celebrated case of non-iterativity, often seen in earlier works as support for the view that long-distance spreading occurs by the serial reapplication, or iteration, of a rule. Spreading to only the first nasal in Nati follows under this view, because the target *n* cannot in turn act as a trigger for a second application of the rule; only retroflex continuants trigger Nati. The appeal of this explanation is to derive the limitation on spreading from the trigger condition, the latter independently needed. (See Johnson 1972, Howard 1973, Anderson 1974, Jensen and Strong-Jensen 1976, Ringen 1976, Kenstowicz and Kisseberth 1977, Battistella 1979, Kiparsky 1985, and references therein for general discussion of this issue.) As stated, however, this explanation presupposes serial rule ordering.

Consider first the trigger condition itself. As Gafos (1996) notes, there may be a principled basis on which to distinguish continuant from stop retroflex triggers. Citing Steriade (1995b), he points out that the acoustic cues for retroflexion are present throughout a continuant; in a stop, however, they are most robust preceding closure (primarily in a lowering of F3), absent during closure, and relatively weak afterwards. If acoustic salience implies a greater tendency to trigger spreading, then continuants should be inherently better triggers for this reason. Gafos consequently posits a constraint requiring retroflex harmony specifically from continuant triggers. We follow Gafos in this, while making a further assumption in order to reconcile the 'non-iterative' nature of the process with OT's parallel evaluation of surface forms: we assume that [retroflex] is aligned not to the right word edge, but to a consonant. The constraint is informally stated in (47). The condition on the ordering of the trigger and target segments guarantees that spreading occurs to the right; were we to align to the right word edge, we would wrongly encompass segments beyond the first dental nasal.[18]

(47) ALIGN-R([RETROFLEX], C): Align any [retroflex] feature contained in a [+continuant] segment S_m to a consonant S_n, where $n > m$

Turning to the restriction on the target (dental nasal), it seems likely that this fact is related to the greater tendency of nasals over obstruents to assimilate in consonant clusters. Jun (1995) and Padgett (to appear), building on ideas of Ohala and Ohala (1993) and Mohanan (1993), argue for a decomposition of faithfulness constraints, such that faithfulness to nasal place is intrinsically lower ranked than faithfulness to obstruent place. This intrinsic ordering predicts that nasals can sometimes assimilate in place while obstruents do not, explaining at the same time why the reverse never holds. While preserving the central idea of an obstruent versus nasal place asymmetry, we implement it here by means of contrast constraints. Let us focus on the contrast in Sanskrit between the retroflex and dental sounds, shown in (48). This minimal contrast involves a difference in the orientation of the tongue tip/blade on the articulatory side, and a difference in higher formant values on the acoustic side (lower F3-4 for retroflexes). To implement this contrast we assume a single feature [retroflex], following on the discussion of section 3.2.

(48) The dental vs. retroflex contrast in Sanskrit

Dental	t	tʰ	d	dʰ	s	n	l
Retroflex	ʈ	ʈʰ	ɖ	ɖʰ	ʂ	ɳ	r

This contrast is forced by a constraint CONTRAST(RET). Assume in particular a family of such constraints along the lines of (49)a. Given the discussion above, the important distinction here is that between contrasts in nasals and contrasts in other consonants. We therefore group the latter together for expository purposes under the name $\text{CONTRAST}_{\text{CONS}}$(RET), and posit the ranking shown in (49)b.

(49) Constraints on retroflexion contrast, and inherent weakness of nasal place

a. $\text{CONTRAST}_{\text{STOP}}$(RET), $\text{CONTRAST}_{\text{FRIC}}$(RET), $\text{CONTRAST}_{\text{NAS}}$(RET), $\text{CONTRAST}_{\text{LIQ}}$(RET)
b. $\text{CONTRAST}_{\text{CONS}}$(RET) >> $\text{CONTRAST}_{\text{NAS}}$(RET)

The target condition of Nati simply reflects a more general cross-linguistic fact: nasal place is more prone to neutralisation than is obstruent place. Since Nati forces a loss in nasal place contrasts, alignment must outrank the relevant contrast constraint; however, since obstruents do not undergo Nati, alignment itself is dominated by the constraint on obstruent place contrasts. This overall ranking and its effects are shown in tableaux (50) and (51). Candidate (50)a fails to spread retroflexion at all, allowing for the possibility of a contrast between *n* and *ɲ* after *r/ʂ*, as shown. Hence it satisfies CONTRAST$_{NAS}$(RET) but violates alignment. The competition goes to (50)b-c. Both of these candidates satisfy alignment as formulated; this constraint crucially demands that only some consonant (not all consonants, or a word edge) be targeted. Given this formulation, however, candidate (50)c goes too far, entertaining unnecessary violations of CONTRAST$_{NAS}$(RET).

(50) Nati: ALIGN-R(RETRO) >> CONTRAST$_{NAS}$(RET)

	Contrast$_{Cons}$(Ret)	Align-R(retro)	Contrast$_{Nas}$(Ret)
a. va[r]$_R$nana:na:m • va[r]$_R$ɲana:na:m		*!	
b. ☞ va[rɳ]$_R$ana:na:m			*
c. va[rɳaɳa:ɳ]$_R$a:m			**!*

In tableau (51) we see how spreading fails to target other dentals, or to spread across them. The problem with both (51)b-c is that targeting other dentals leads to a fatal violation of the high-ranking constraints enforcing the dental-retroflex contrast, in this case a *d* vs. *ɖ* contrast. Spreading is therefore impossible in such forms. We do not consider a candidate in which the coronal obstruent is both dental and retroflex; these values are as incompatible as [+/- back], implying an articulatory impossibility suitable for Gen. Once again we see that the combined reliance on strict locality and markedness constraints obviates any appeal to line crossing.

(51) Blocking of Nati: CONTRAST$_{CONS}$(RET) >> ALIGN-R(RETRO)

	Contrast$_{Cons}$(Ret)	Align-R(retro)	Contrast$_{Nas}$(Ret)
a. ☞ k[ʂ]$_R$veda:na • k[ʂ]$_R$veɖa:na		*	
b. k[ʂveɖ]$_R$a:na	*!		
c. k[ʂveɖa:ɳ]$_R$a	*!		*

The account of blocking by palatals, the other series of coronals, is essentially the same, with the exception that even the palatal nasal blocks (and does not undergo), so that all of the relevant contrast constraints distinguishing palatals from other coronals dominate alignment.

With a basic analysis of Nati at hand, we can now turn to our central concerns of locality and feature realisation. In an earlier discussion of the relevance of coronal harmonies for a theory of locality, Ní Chiosáin and Padgett (1993) posit a spreading as in (52), in which the spreading feature is [coronal]-dependent, and spreading is local relativised to the [coronal] tier, assuming a feature geometry familiar from Steriade (1986), Sagey (1986) and McCarthy (1988).

(52) Relativised locality in coronal harmony; *r* = [retroflex]

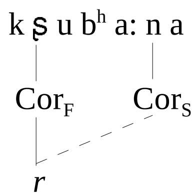

It was a liability of this account that, while the spreading of major place gestures like [coronal] was assumed to be subject to strict (segmental) locality (for reasons seen in section 3), the spreading of [coronal]-dependent features was not. Relying in part on recent work on coronal harmonies by Flemming (1995a) and Gafos (1996), we take the more uniform position here that all spreading is strictly local as defined in section 2.1. Therefore, all segments intervening between the triggers of retroflexion and the target *n* are participants in the Nati spreading domain, e.g. *k[ʂub*^{h}*a:ɳ]*$_{r}$*a*. Consistent with the general arguments made earlier, we seek an understanding of both blocking and permeability effects in coronal harmonies, and the related issue of segment realisations, that relies on substantive constraints on contrast and articulatory complexity. We dealt with the cases of blocking above; what remains is to examine the behaviour of vowels and non-coronal consonants, all permeable.

Consider first the explanation for this permeability offered by the relativised locality view in (52). The reason vowels and non-coronal consonants do not block Nati is because they lack [coronal] nodes; definitionally dependent on the feature [coronal], [retroflexion] seeks out this mother node as its only legitimate anchor. Given these assumptions, non-coronals cannot block; in fact, they cannot be relevant in any way, not even realising the retroflexion that spans them. One problem for this account as stated arises over its implication that vowels cannot be retroflexed. As we noted in section 3.2, retroflexion is a feature that occurs, sometimes contrastively, on segments of all constriction degrees across languages: vowels, liquids and consonants (see for example Ladefoged and Maddieson 1996). The vowel in English words like *bird* [bɝd] (in some dialects) is retroflex, for instance. The purported non-participation of vowels in Sanskrit could not then be definitional in any sense. Our assumption of strict, segmental locality of course entails that vowels, and all intervening sounds, are participants in Nati. The immediate result is that there are segments such as *[a]*$_{R}$, *[b*^{h}*]*$_{R}$, and so on. From the perspective of DT, we need to inquire into both the articulatory and perceptual markedness of such segments. Let us consider each in turn.

Retroflexion is a tongue blade/tip articulation (a curling back in the roughly horizontal dimension) that makes no demands itself on a segment's stricture (a matter of the roughly vertical orientation for coronals). In other words, a retroflexion gesture is not like e.g., [coronal]$_{F}$; its 'degree' of implementation essentially means the degree of *retroflexion.* (Compare the front/back tongue body dimension in vowels, basically stricture neutral.) It follows from this understanding of retroflexion that this feature can spread across vowels, and segments of any constriction degree, without implying any change in the strictural status of these participating segments; that is, in Nati a retroflex $[\text{coronal}]_F$ spreads only retroflexion, and not coronality, or frication. More generally, retroflexion is articulatorily compatible with most segment types. Since retroflexion imposed on a gesture involves an extra degree of articulatory complexity, however, the general ranking depicted below holds across segment types.

(53) Articulatory markedness of retroflexion (X is any set of features)

*[X, retroflex] >> *[X]

On the perceptual side, the existence of contrastively retroflexed vowels in some languages implies a constraint $\text{C}{\scriptsize\text{ONTRAST}}_{\text{V}{\tiny\text{OWEL}}}$(Ret) (or perhaps a family of such constraints). Since retroflexion is not contrastive in Sanskrit vowels, though, we can conclude that the articulatory complexity constraint dominates the contrast constraint for this class of sounds, as shown below. (The opposite ranking holds for Sanskrit consonants, of course.) The analysis correctly predicts that languages lacking such a contrast have non-retroflexed vowels, the simpler segment type.

(54) Retroflexion is not contrastive in Sanskrit vowels: *[V, retro] >> CONTRAST_V(RET)

	*[V, retro]	Contrast_V(Ret)
a. V • $[V]_R$	*!	
b. ☞ V		*
c. $[V]_R$	*!	*

On the other hand, such vowels are forced by the Nati alignment constraint; alignment therefore dominates the constraint on articulatory complexity:

(55) Nati forces retroflex vowels: ALIGN-R(RETRO) >> *[V, retroflex]

	Align-R(retro)	*[V, retro]
a. ☞ k[ʂub^{h}a:ɳ]$_R$a		**
b. k[ʂ]$_R$ub^{h}a:na	*!	

The analysis of non-coronal consonants like $[b^h]_R$ is even more straightforward. Retroflexion is not contrastive for such consonants in any language, implying a basic failure of minimal perceptual distance between $[b^h]_R$ and b^h and other analogous pairs. Even if we suppose that constraints demanding a retroflexion contrast in consonants are so generally stated as to demand contrast of labials (for example) as well as coronals, non-coronal pairings cannot in principle satisfy such constraints. For this reason candidate (56)a necessarily fails in contrast just as surely as do candidates (56)b-c. Given this state of affairs, the ranking of the contrast constraint cannot make a difference in the case of non-coronals; outside of Nati domains, articulatory simplicity will favor (56)b no matter what.[19] Since segments like $[b^h]_R$ do occur in Nati domains, ALIGN-R(RETRO) must dominate the articulatory constraint, however. The argument is analogous to that of (55).

(56) Permanent failure of retroflexion contrast for non-coronals

	$Contrast_{Cons}$(Ret)	*[C, retro]
a. b^h • $[b^h]_R$	*	*
b. ☞ b^h	*	
c. $[b^h]_R$	*	*

This general approach to coronal harmonies (without an explicit DT analysis), beyond (and including) Sanskrit is forcefully argued for in recent works of Flemming (1995a) and Gafos (1996). An important claim of these works concerns the nature of the spreading features in coronal harmony systems. Consistent with the arguments above, what spreads are gestures that define the orientation of the tongue tip/blade, having no strictural implications. It is argued that these features are articulatorily compatible with a broad range of sounds, including vowels and non-coronal consonants. For articulations other than retroflexion involved in coronal harmonies (articulations involving degrees of apicality/laminality or surface area contact, for instance), the perceptual effects on vowels, non-coronal consonants, and even coronal stops in a spreading domain are minimal or undetectable. (The fact that only coronal fricatives can fruitfully contrast for these articulations leads directly to an account of the transparency of coronal stops in systems such as Tahltan.) Both authors (citing Bladon and Nolan 1977) note that tongue tip postures coarticulate across vowels in English, for example, a phonetic correlate of some coronal harmony scenarios. Other phonetic studies similarly suggest that the tongue tip (unlike the tongue body) is free to assume various positions without interfering with simultaneous vowel articulations (e.g. Farnetani, Vagges, Magno-Caldognetto 1985, and references in Gafos 1996). See these works for further discussion of these issues (and especially Gafos 1996 for an exhaustive survey of coronal harmonies) coming to conclusions about locality and realisation broadly consistent with ours.

Did the intervening sounds truly realise retroflexion in a Nati domain? Some earlier non-generative treatments of Nati assumed explicitly that intervening segments indeed participate. In an account couched in Firthian Prosodic Phonology, for example, Allen (1951) argued specifically against 'action at a distance' in phonology, and maintained that all segments intervening between the continuant retroflex trigger and nasal target of Nati were participants, within an extended and continuous span of retroflexion such as we assume here.[20] Allen further suggested that the (incorrect) view that Nati skips intervening segments stems from the representation of the process in Sanskrit orthography: since retroflexion in Sanskrit is contrastive only for consonants, any effect on other segments (such as vowels) are not indicated by the orthography. It should not be concluded from this orthographic representation that intervening vowels did not realise retroflexion. (Compare the frequent transcription of [m], both in orthographies and IPA, for the non-contrastive phonetic sound [ɱ], even in languages where the latter is the invariant outcome before labiodental fricatives.) In languages subject to phonetic scrutiny today one finds examples of such effects on intervening vowels by consonant-to-consonant retroflex spread (see Steriade 1995b). We cannot be certain of the original Sanskrit pronunciations, but the analysis here is empirically plausible.

4.4 Nasal harmony and glottal permeability

Though the focus of this paper is on facts of place of articulation, the ideas concerning contrast, markedness, and locality in spreading are meant to apply generally. With this in mind, we consider here some other cases of apparent transparency, and argue that they too can be profitably recast, once again by trading transparency for a deeper understanding of markedness and segment realisation. Our attention will be devoted largely to the case of glottal stops in nasal spreading domains, because this case is relatively simple; we then speculate on how the notions discussed could extend to other specific cases.

A well-studied case of glottal permeability to nasal spreading is that of Sundanese (Robins 1957, Cohn 1990, 1993). In this language nasality spreads rightward from a nasal consonant throughout a word. Spreading is blocked by a (non-nasal) consonant with a supraglottal constriction of any degree, including a glide; spreading is not blocked by laryngeal segments, however. The forms in (57)a demonstrate that nasal spreading permeates vowels and the laryngeal segments *ʔ* and *h*; those in (57)b show blocking by glides, liquids and obstruents. Examples are drawn from both Robins and Cohn; morphological boundaries are not indicated, and a subscript N denotes nasality of the indicated span.[21]

(57) Sundanese nasal spreading

a.	[ɲãĩãn]$_N$	'wet (active)'
	[mĩʔã]$_N$sih	'love (active)'
	[nìʔĩ]$_N$s	'relax in a cool place (active)'
	ku[mãɦ̃ã]$_N$	'how?'
	[ɲãɦ̃õ]$_N$	'know (active)'
b.	[ηã]$_N$wih	'sing (active)'
	[ηã]$_N$jak	'sift (active)'
	[mã]$_N$rios	'examine (active)'
	[ηũ]$_N$liat	'stretch (active)'
	[ηũ]$_N$dag	'pursue (active)'
	[ηĩ]$_N$sər	'displace (active)'

Our analysis of the Sundanese data largely follows that of Walker (1995, 1996), who builds on the references cited above. First, rightward nasal spreading is forced by a nasal-specific instantiation of alignment to the prosodic word edge called ALIGN-R(NAS) here. Since nasalised vowels surface due to this constraint, alignment must dominate any constraints that militate against them. This is illustrated in tableau (58), where the subordinate articulatory markedness constraint *NAS-VOWEL punishes nasalised vowels for their extra articulatory complexity. This constraint is violated (three times) by the grammatical form *ɲãĩãn*. The alternative (58)b fares better by articulatory markedness, but fatally violates alignment (by four segments).

(58) Spreading forces markedness violations: ALIGN-R(NAS) >> *NAS-VOWEL

	Align-R(nas)	*Nas-Vowel
a. ☞ $[ɲãĩãn]_N$		***
b. $[ɲ]_N aia[n]_N$	*!***	

On the other hand, other constraints outranking ALIGN must be responsible for the failure of spreading to proceed across the glides, liquids and obstruents. For this purpose we call on further articulatory markedness constraints that discourage nasalised versions of these segment types, as Walker does. Such segments, like nasalised vowels, are known to be marked in comparison to their oral counterparts; in addition, a hierarchy of markedness among all of these nasalised segments can be discerned from a cross-linguistic survey of facts of inventories, feature spreading and so on, with nasalised segments becoming worse in the order vowel, glide, liquid, obstruent. (See Schourup 1972, Pulleyblank 1989, Cohn 1990, 1993, Padgett 1994, 1991[95], and Walker 1995, 1996.) Consider then the hierarchy of articulatory markedness constraints below. (We take up the glottal segment below.)[22]

(59) Articulatory markedness in nasalised segments

*Nas-Obst >> *Nas-Liquid >> *Nas-Glide >> *Nas-Vowel >> $*[ʔ]_N$

For our purposes it is sufficient to group together everything above *NAS-VOWEL, and refer to them as *NAS-CONS. These constraints dominate *NAS-VOWEL universally, and must further dominate alignment in order to explain the failure of nasality to reach the right word-edge; this explains the ill-formedness of (60)b. However, alignment is not violated in an all-or-nothing way; the fact that (60)a surfaces, rather than (60)c, shows that alignment is gradiently violable. Since violations must be forced, partial spreading necessarily wins out over no spreading at all, as shown.

(60) Spreading and nasalised markedness: *Nas-Cons >> Align >> *Nas-Vowel

	*Nas-Cons	Align-R(nas)	*Nas-Vowel
a. ☞ [mã]$_N$rios		****	*
b. [mãr̃ĩõs̃]$_N$	*!*		***
c. [m]$_N$arios		*****!	

We are now in a position to consider the behaviour of the laryngeal segments, focussing in particular on the glottal stop. This segment is permeable to the harmony process, as seen in examples like *[nìʔĩ]$_N$s* from (57)a above.[23] This is not a fact particular to Sundanese; other languages in which glottal stops are permeable to nasal spreading are not hard to find, and include Capanahua, Guaraní, and Seneca (see Schourup 1972). We are not aware of any cases in which glottal stops block nasal spreading. (Cohn 1993 seems to imply such cases in a brief mention of sources, but Pullum and Walker 1996 find no evidence of blocking by glottal stops in a careful review of those sources.)

Yet consider the now familiar paradox presented by this fact for any unidimensional markedness theory. The fact that glottal stops rarely or never block nasal spreading suggests that a constraint *[ʔ]$_N$ penalising nasalised glottal stops, if such a constraint exists at all, must be universally low-ranked on the hierarchy of nasalised segments described above. This is Walker's (1996) conclusion, and she groups these segments roughly with nasalised vowels. The problem is that nasalised glottal stops are contrastive in no language, setting them far apart from nasalised vowels for example. In Optimality Theoretic terms again, if *[ʔ]$_N$ is so eminently violable due to higher ranking alignment, we should expect it also to be dominated occasionally by faithfulness constraints. This apparent contradiction, whereby nasalised glottal stops are never contrastive, yet glottal stops do not block nasal harmony, is resolved by Cohn (1990) by positing that glottal stops are skipped by nasal harmony. Since they are never targeted, the theory can maintain the consistent position that nasalised glottal stops are impossible. In order to effect this skipping Cohn suggests that the feature [nasal] resides under a feature geometric class node Supralaryngeal, a node that laryngeal segments definitionally lack (Clements 1985, Sagey 1986). Locality is obeyed on the assumption that it is relativised to this mother node anchor, as shown in (61). Since glottal stops lack this mother node, they are also incapable of bearing nasality for the purpose of contrast; in this way the facts of harmony and contrast are unified.

(61) Locality relativised to the Supralaryngeal node

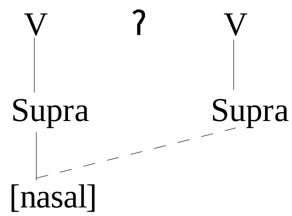

The resolution of this problem presented by DT is instead the following. First, glottal stops are indeed participants in harmony (as the notation *[nìʔĩ]*$_N$*s* suggests), a conclusion following from segmentally strict locality. Second, there is therefore a segment *[ʔ]*$_N$ that phonology must reckon with. Further, since glottal stops are so permeable by nasal spreading, this derived segment must present little or no articulatory difficulties; in this sense, the conclusion reached earlier that such a segment is unmarked is entirely correct. In all of the above we are in agreement with Walker (1996) and Pullum and Walker (1996). Finally, since *[ʔ]*$_N$ never contrasts with *ʔ*, this potential contrast must fatally fail the test of minimal perceptual distance. We unpack the latter markedness statements in the following discussion.

Spreading of nasality onto *ʔ* plausibly causes little difficulty in the area of articulation, as Cohn (1990, 1993) notes. Glottal constriction is in no significant way antagonistic to a lowered velum; these articulators are independent. Similarly, a lowered velum superimposed on *ʔ* has no aerodynamic consequences (good or bad). In this latter respect glottal stops differ from other nasalised segments, such as liquids or fricatives, in which oral air flow is compromised by any nasal 'leak'. In the case of fricatives for example, the resultant loss of oral air flow makes the realisation of frication difficult (Ohala 1975). On a hierarchy of articulatory markedness reflecting articulatory difficulty, *[ʔ]*$_N$ belongs at or near the bottom, as Walker (1996) suggests; we assume for discussion that *[ʔ]*$_N$ occupies the bottom of this scale, as shown in (59) above. Finally, we posit the universal ranking *[ʔ]$_N$ >> *ʔ, since the former involves more articulatory complexity.

From the perspective of perception there are again no consequences of velum lowering in the case of glottal stops; there is of course no nasal airflow during a glottal stop with a lowered velum, since there is no airflow at all during a glottal constriction. Hence there can be no perceived nasality during glottal closure: the acoustic result of [nasal] on *ʔ* is simply *ʔ*, hence our use of the notation *[ʔ]*$_N$, rather than ʔ (we take only the latter to imply audible nasality). Here is a quintessential case of failed minimal perceptual distance, the distance being zero. Assuming a general constraint CONTRAST(NAS)—which a pairing such as *d* • *n* would satisfy, for instance—the pairing *[ʔ]*$_N$ • *ʔ* would always violate this constraint. This is shown in tableau (62). Candidate (62)a is a valiant attempt at contrast by decree, but since satisfaction of CONTRAST(NAS) implies providing a minimal perceptual distance for contrast—impossible here—this attempt fails. Candidates (62)b-c equally fail in the area of contrast, but differ in their implications for articulatory markedness. Assuming that plain *ʔ* is indeed less marked in this respect, this segment surfaces (all else being equal).

(62) *[ʔ]$_N$* • *ʔ*: failed contrast

	Contrast(nas)	*[ʔ]$_N$	*ʔ
a. [ʔ]$_N$ • ʔ	*	*!	*
b. [ʔ]$_N$	*	*!	
c. ☞ ʔ	*		*

It should be evident from this analysis that the ranking of CONTRAST(NAS) is inconsequential to the outcome in the case of glottal stop. This is not true for all segment types, of course; where the contrast constraint can be met in principle, it can force a contrast in spite of articulatory markedness, as tableau (63) makes clear. (We assume for discussion that voiced obstruents are articulatorily worse than nasals.)[24]

(63) *d* • *n* contrast

	Contrast(nas)	*d	*n
a. ☞ d • n		*	*
b. d	*!	*	
c. n	*!		*

Let us return now to the Sundanese forms. Adding the constraint **[ʔ]$_N$* to the hierarchy for Sundanese shown in (60) above, we derive the participation of the glottal stop, as shown below. The permeability of glottal stop to spreading follows from its inherent low ranking on the scale penalising nasalised segments. Though as a contrastive segment a nasalised glottal stop is highly marked—in fact, impossible—this fact is doubly irrelevant in nasal harmony systems. First, as we noted earlier in the discussion of place features, the existence of harmony entails that contrasts along the relevant dimension are neutralised in the spreading domain; with contrast factored out, only the dimension of articulatory markedness comes into play. Second, a contrast between plain and nasalised glottal stops is in principle impossible, simply because such a contrast can never meet the criterion of minimal perceptual distance. As the reader can verify, glottal permeability will therefore result no matter where a constraint like CONTRAST(NAS) appears in the Sundanese hierarchy.[25]

(64) *ʔ* participates in nasal harmony

	*Nas-Cons	Align-R(nas)	*Nas-Vowel	*[ʔ]$_N$
a. [n]$_N$iʔis		**!**		
b. [nĩ]$_N$ʔis		**!*	*	
c. ☞ [nĩʔĩ]$_N$s		*	**	*
d. [nĩʔĩ s̃]$_N$	*!		**	*

Once more we have the dispersion syndrome: for the purposes of contrast with *ʔ*, [ʔ]$_N$ is deeply marked. In terms of articulation, though, it could hardly be better. DT's bidimensional view makes it possible to embrace these two sides of this segment; once they are separated out, the paradox seen above resolves in a way that requires no appeal to relativised locality. With this explanation we simply capitalise on the properties of a markedness theory (DT) that is independently necessary, and a segment ([ʔ]$_N$) that exists (as Cohn 1993 agrees; see also Ohala 1990, note 10, and Walker and Pullum 1996). From this broad perspective, explanations in terms of geometric representations and locality assumptions turn out to be at best redundant. More specific to this case, the feature geometry required for relativising locality in nasal harmonies, in particular a node Supralaryngeal, is very poorly motivated on independent grounds (see McCarthy 1988, Padgett 1991[95] and references therein).

The permeability of *h* can be understood in essentially the same way, with the only difference following from the fact that this segment (unlike *ʔ*) requires air flow. This means that a lowered velum has aerodynamic consequences for *h*, and that the perceptual distance between *h* and *[h]$_N$* is not precisely zero. The distance is small enough, however, that these segments cannot easily meet the demands of contrast (Ohala 1975, Cohn 1993).

We conclude this section by briefly discussing other well known cases of apparent transparency, in particular the permeability of obstruents to nasal spreading in some languages, and the frequent permeability of obstruents to tone spreading. Representative examples of each are given in (65) and (66). (Subscript H indicates a high tone spreading domain in (66).)

(65) Nasality and obstruents: Southern Barasano (examples from Piggott 1992)

[w̃ãtĩ]$_N$	'demon'
[mãsã]$_N$	'people'
[kãmõkã]$_N$	'a rattle'

(66) Tone and obstruents: Ciyao (examples from Hyman and Ngunga 1994)

nga-ni-ju-di[vádí]$_H$ l-a	'he didn't forget'
[jú-kú]$_H$-ci-dy-á	'he is eating it'

Our view of permeability once again implies segment realisations not usually considered, e.g. *[k]*$_N$ and *[k]*$_H$, etc. That such segment types are not usually considered can be attributed to two facts, among others. First, nasality and tone are not audible in the obstruents (at least during the closure phase of voiceless obstruents); second, there is the prevalent unidimensional understanding of markedness: distinctions must be contrastive. Yet viewed from the current perspective of DT, these points are not compelling. Indeed, we know from the case of glottal stop, and other cases in this paper, that lack of audibility, and the related lack of contrast, do not entail the lack of a segment type. Therefore the existence of *[k]*$_N$ and *[k]*$_H$, etc., is a live issue. These segments cannot exist in contrast with their plain counterparts, presumably for the same reason that *ʔ* and *[ʔ]*$_N$ cannot contrast: they cannot bear nasality or tone in any perceptible way. (This point is overstated somewhat; see the remarks involving tone below.) Since contrast is not at issue here, the important question is whether obstruents are *articulatorily* compatible with nasality and tone for the purposes of spreading.

Consider nasality first. Walker (1996) argues on typological grounds that 'transparent' obstruents are in fact participants in nasal harmony. This argument follows from the observation that obstruents are the only segment class known to exhibit (apparent) transparency to nasal spreading; when not transparent, obstruents block nasal spreading (or if voiced, they may undergo and become nasals). Vowels, glides, and liquids, on the other hand, are never skipped by nasal harmony; rather, they either block or *participate*. This is a very odd asymmetry in behaviour, and one that receives essentially ad hoc accounts in previous analyses of nasal harmony patterns. If locality is segmentally strict, then all segments are either blockers or participants, removing the anomaly. From the perspective of typology, at least, this is a clear improvement. The burden of explanation for the behaviour of obstruents must now be borne by the theory of segment realisation.

The clear challenge is that obstruents by definition require significant obstruction of the vocal tract; since a lowered velum presents a nasal leak, the danger is that a segment like *[k]*$_N$ will simply (and incorrectly) amount to a nasal. That nasality is antagonistic to obstruents is confirmed by the typology of blocking: as Walker notes, when obstruents participate in nasal harmony (or are 'skipped'), so do all other segment types; on the other hand, vowels, glides and liquids can be participants while obstruents block. Obstruents are therefore at the top of the hierarchy seen earlier disfavoring nasalised segments, and represent the segment class least compatible with nasality. Nevertheless, given our assumptions, a constraint *NAS-OBST must be violable (i.e., we have *[k]*$_N$ in marked cases like Southern Barasano). Suppose we grant that such a segment would have a raised velum at the point of release, since the crucial obstruent property is a characteristic burst. A possibility is that the velum is lowered to some degree for some or all of the segment up to this release. This might predict a reduction in burst amplitudes or durations in comparison to those of plain obstruents, since with a nasal leak there would be less time for the build-up of pressure in the oral cavity. It might also predict partial voicing into the closure. So far as we know, there are no relevant studies bearing on such stops. However, Gerfen (1996) argues for something akin to this scenario for fricatives in nasal spreading domains of Coatzospan Mixtec, in a study of nasal air pressure tracings. The data lead him to conclude that fricatives have a closed velum phonetic target towards their right edges; up to this point, the fricatives are argued to present nasal air flow. If this interpretation is accurate, such facts would suggest that a spreading gesture (in this case an open velum) need not be entirely continuous in implementation (see section 2.1); perhaps this need not be too surprising, since velum position is known to vary according to segment type anyway. However, if voiceless segments in nasal spans could be systematically even partially nasalised, we would have some support for our posited $[k]_N$, etc. On a more subtle note, it is possible that the articulatory correlate of [nasal] should be sought at a more abstract level than that of velum position, such as one of gestural planning. Just as articulatory evidence supports the idea of a continuous vocalic gestural stream with consonantal gestures—even velars—'superimposed' (see section 3 and references there), it is conceivable that the raised velum required to maintain obstruency is superimposed on a continuous [nasal] gestural span. Putting these speculations aside, we simply note that further work is called for.

The case of tone is also complex. First, though tone is inhibited during the closure phase of an obstruent, obstruents, and other consonants, do have intrinsic tonal properties which they contribute audibly to their immediate vicinity. The observation is that voiceless consonants cause a small rise in F_0, while voiced consonants cause it to fall; both effects are strongest in a following vowel. In some languages this phonetic effect has been elevated to a phonological constraint whereby vowels bear high tones after voiceless consonants and low tones after voiced ones. This can lead to the existence of phonemic tones, if the laryngeal distinctions among obstruents are subsequently neutralised. (On these claims see Ohala 1978, Hombert, Ohala and Ewan 1979, Svantesson 1989 and references therein.) In this sense, even obstruents must already be viewed as bearing tonal specifications, though these specifications may not be audible during closure, and are often small enough as to be rather inconsequential. Further, these tonal specifications can interact with tone in vowels. In Nupe, for example, there is a process of low tone spreading from a prefix to a high-toned root, creating a rising contour; it applies through voiced consonants only, giving forms like those in (67)a next to those in (67)b. That only voiced consonants allow the spreading to proceed implies that they are indeed participants in the spreading, and not skipped: the low tone is compatible with intrinsic low tone of voiced consonants, but is blocked by the more antagonistic high tone of voiceless ones. (The examples are from Kenstowicz and Kisseberth 1979:267-8).

(67) Obstruent participation in tone spread: Nupe

a.	ègbǎ	'a border on a garment'	b.	èkpá	'length'
	èbě	'pumpkin'		èfú	'honey'
	èdzǎ	'sash'		ètú	'parasite'

It is much more common for tone to spread through obstruents, and all consonants, without any interference at all. From our perspective, the conclusion that should be drawn from this fact is that consonants are in the worst case only mildly antagonistic to tonal specifications in terms of articulation. We can say no more given our current level of understanding of the matter.

5 Conclusion

As we have shown, once markedness is factored into independent components which have articulatory and perceptual underpinnings respectively, we can go a long way toward resolving the long-standing problem of locality in spreading. The same principles have been applied to each case examined in this paper, giving them an explanatory coherence that has been notably absent in theories positing relativisation to prosodic anchors and feature geometric nodes, along with separate process-specific stipulations for computing locality. (See Gafos 1996 for further argumentation to this effect.) Moreover, we find an explanatory unity between these locality facts and the fundamental domain of markedness more generally: there are no locality-specific devices at all. As we have noted, the segment types entailed by the account proposed amount to empirical claims that can be tested; further, Dispersion Theory is already independently well-motivated—indeed, it was not conceived of to explain locality facts, and the analyses above provide independent support for it. There are of course as many questions as answers at this point. An important goal for future work is to understand the precise nature and workings of contrast constraints better. Further, there is more to markedness than the two dimensions featured above. For instance, a countervailing force to pure dispersion is symmetry of inventories (or contrastive feature minimisation), another phonological hallmark. Our primary goal here has simply been to make the case for strict locality, and to show how DT can make it possible.

A worry for some will be the explicit reference to functional grounding seen throughout. On this issue we make two brief points. First, DT is a theory of contrast and markedness, subjects which are at the very heart of phonology. If markedness and contrast are best understood by reference to phonetic underpinnings, then phonology must explicate those underpinnings, among other things. Second, phonetic grounding is not new to phonology. In fact, significant advances in phonology have often had the property of bringing phonology more in line with phonetics in some way: feature theory, autosegmentalism, feature geometry, and grounded feature cooccurrence conditions come to mind. These connections with phonetics have sometimes been treated as a matter of happy convergence between phonetics and phonology, but given the history it seems more than reasonable to revise this view, and to treat phonetic facts themselves as a motivating factor in phonological explanations.

What of other phenomena not examined here that seem to require 'action at a distance'? We note that recent work by Gafos (1996) makes a persuasive case that templatic effects of Semitic, and other similar cases, require neither planar C-V segregation nor the concomitant cross-vowel spreading of consonants in words like *samam* (see McCarthy 1979); instead these cases fall into the independently necessary realm of reduplicative effects. Obligatory Contour (dissimilatory) effects more generally can involve apparent action at a distance, as when consonants in a CVC form must not be of identical place (see McCarthy 1986, Mester 1986, Yip 1989, Padgett 1991[95], and for new perspectives Pierrehumbert 1993, Frisch 1996, Itô and Mester to appear-b and Alderete to appear). Apart from whether this is action at a distance in an interesting sense, it is clear that the locality requirements on spreading and dissimilatory effects are not the same—compare the facts of spreading surveyed in section 3 with OCP place effects. Rather than hoping for a unified theory of locality, it seems more promising to reduce 'locality' effects to independent properties of the phenomena involved. This has been our goal in the case of spreading. Finally, there are cases of apparent transparency in spreading that are not obviously amenable to the view explored in this paper; transparent vowels in vowel harmony present such a challenge (see the appendix). This raises an important point. For any account of apparent transparency to be worth considering, it should preserve the generalisations discussed in section 3. Put most broadly, those generalisations amount to the following: segments in a spreading domain are always compatible with the spreading feature (even if the result is marked). Indeed, this best sums up what features spread long distance in phonology, and why. We believe that this compatibility requirement holds of all segments in a spreading domain, including 'transparent' ones. If this claim is right, segment skipping is *in principle* incapable of explaining it, since skipped segments by hypothesis make no contact with the spreading feature.

Appendix on transparent vowels: segment transparency vs. constraint opacity

This appendix has two goals. The first is to consider cases of apparent transparency in spreading that do not appear to be readily understood in terms of our claims above. Examples involving vowel transparency in vowel harmonies are perhaps the clearest kind of such case. Our central claim about transparency above is that it is only apparent: for most alleged examples, it turns out that the spreading feature is indeed phonetically realised on the relevant segment. Vowel transparencies (in at least some cases) do not seem likely to succumb to this understanding of permeability: no amount of close listening reveals the spreading feature. In fact, such vowels might be specified in a manner that is directly antagonistic or contradictory to the spreading feature. Our second goal is to suggest an understanding of these cases that nevertheless fully reconciles them with segmentally strict locality. The basic idea is not novel, but is in fact quite retrogressive, resurrecting an old view in generative accounts of vowel transparency: transparent vowels actually undergo harmony at an abstract level of representation; however, the offending feature cannot appear at the surface. Vowel transparency, in this view, is really a kind of constraint *opacity*; in this case, what is opaque is the locality generalisation, which in fact holds firm. A basic argument for this understanding of vowel transparency is the very fact that it retains intact the claims made in this paper. This is a serious issue: other accounts of vowel transparency basically renege on the promise of a real understanding of locality, and of the asymmetries in place spreading discussed in section 3. For reasons of space the discussion here cannot be definitive, or go beyond briefly making these points.

Our example for discussion comes from a process of lax (or centralisation) harmony in Pasiego Montañes Spanish (Penny 1969a-b, 1970, McCarthy 1984, Vago 1988, Hualde 1989, Flemming 1993).[26] We note the existence of an equally interesting height harmony in the language, also involving transparency, though we do not focus on it here. The forms below are taken from McCarthy (1984). The inventory of vowels occurring at the surface in Pasiego are given in (68). The capitals denote realisations that are generally described as lax or centralised in comparison to their counterparts. We will return to details of Pasiego vowel realisations later, and use the terms 'tense' and 'lax' throughout. Note, in particular, the absence of a lax counterpart to *e*.

(68) Pasiego tense/lax distinction

i/I		u/U
e		o/O
	A/a	

The lax variants occur under very restricted conditions: only in masculine singular count nouns or adjectives ending in a high back stressless vowel. In such words, all of the vowels (excepting *e*) are lax. In all other words, all vowels are instead tense. We assume a feature [tense] (following McCarthy 1984) to characterise this distinction between lax and tense words, illustrated below.

(69) Tense vs. lax distinction in Pasiego words

[+tense] words		[-tense] words (masc. sg. count)	
abiʎánus	'hazels'	AbIʎÁnU	'hazel'
soldáus	'soldiers'	sOldÁU	'soldier'
kastáɲus	'chestnut trees'	kAstÁɲU	'chestnut tree'
simpátikus	'congenial' (pl)	sImpÁtIkU	'congenial'
málu	'evil' (sg mass)	mÁlU	'evil'

We also follow McCarthy (1984) in assuming that [-tense] is morphologically introduced at the right edge, and that leftward vowel harmony dictates the word-level effect. Suppose in particular that [-tense] indicates 'masculine singular count' in nouns/adjectives, and that an undominated constraint MORPHCONTRAST requires that this morphological contrast be phonologically expressed.[27] (We make the somewhat common assumption that constraints can force the realisation of morphemes, while employing a statement in line with the view of contrast we assume.) Harmony is effected by ALIGN-L([-TENSE], PWD), while the very general dispreference for the lax vowels requires a constraint, here called *[-tense]. This latter constraint is intended as a cover for a genuine account of the preference for the tense variants. We insert this caveat because it is not plausible to assume that lax or centralised variants are more complex articulatorily than the tense or peripheral vowels; rather, the preference for the peripheral vowels is likely to be rooted in minimal perceptual distance as dispersion theory suggests (see Flemming 1995b for discussion of cases like this, and see further below). Candidate (70)c fails because it entails a lack of contrast between this form and its mass noun counterpart *malu* seen above; (70)b falls short of full harmony.

(70) Appearance and harmony of [-tense]: MORPHCONTRAST, ALIGN >> *[-tense]

	MorphContrast	Align-L	*[-tense]
a. ☞ [mAlU]$_{\text{-T}}$			**
b. ma[lU]$_{\text{-T}}$		**	*
c. malu	*!		

The point of interest here is the behaviour of the mid front vowel *e*. This vowel fails to undergo any change in lax forms; more important, it is transparent to the transmission of laxness. The italicised forms below illustrate this transparency.

(71) Mid front vowel inertness, and transparency, to laxness harmony

ermánus	**e**rmÁnU	'brother' pl/sg
bedánus	b**e**dÁnU	'wood-chisel' pl/sg
komfesonárjus	*kOmf**e**sOnÁrjU*	'confessional' pl/sg
kampečánus	*kAmp**e**čÁnU*	'noble' pl/sg

The account as presented so far predicts instead that *e* should undergo harmony. Suppose that a constraint *E rules out this non-occurring vowel (though see the remarks above concerning *[-tense]). Since harmony actually makes it past the mid front vowel to the left edge of a word, segmentally strict locality entails that *e* is a participant, violating this constraint. Alignment must therefore dominate *E. Given locality, the reverse ranking could only lead to blocking of harmony. This is illustrated in the following tableau.

(72) *e* undergoes lax harmony (assuming strict locality)

	Align-L	*E	*[-tense]
a. ☞ [kOmfEsOnArjU]$_{\text{-T}}$		*	*****
b. komfe[sOnArjU]$_{\text{-T}}$	*!****		***

We suggest that this is the right analysis of Pasiego. Before proceeding, let us consider what the nature of the problem is for our general account of spreading and locality. First, cases of apparent transparency such as this are only a problem assuming that harmony indeed proceeds past the relevant vowel and to the word edge. The alternative is that the feature value that seems to arrive by spreading occurs on the far side of the vowel as a default value. This latter strategy has a well-traveled history in vowel harmony analyses, and is a means by which previous analysts have attempted to avoid the conclusion that vowel transparency exists. (See van der Hulst and van de Weijer 1995 for some discussion and references.) In OT terms, the idea would be that

[-tense] 'emerges as unmarked' on the initial vowel in *kOmfesOnArjU*. Though such cases might exist, this general approach runs into trouble whenever the spreading feature is not plausibly viewed as unmarked. Pasiego seems to be such a case, given the highly restricted occurrence of laxness in the language. Other problematic cases exist (involving e.g., transparency to rounding harmony, or to ATR harmony). We must conclude that the relevant lax vowels in Pasiego are lax by harmony, so that the issue of transparency is real.

A more subtle move to avoid this conclusion involves probing more seriously the nature of the 'tense/lax' or peripheral/central distinction in Pasiego. For instance, suppose that what Penny (the source of data for subsequent works) transcribes as [e] in lax words is in fact lax. This attempt cannot hope to rest on a carelessly missed distinction between [e] and its supposed lax counterpart, however. Penny's comprehensive studies of Pasiego include transcriptions that are unusually intricate in detail, with accompanying discussion. He is very clear in the assertion that *e* does not alternate. Perhaps a viable move, though, would be to assume that the sets of segments denoted by 'tense' and 'lax' in fact overlap, both containing phonetic *e*. Penny (1969b) treats *e* as a member of each class, for instance, and the statement that transparent vowels 'occur in both harmonic sets' is often made, as a descriptive generalisation. The idea here, though, is to take it seriously as a theoretical position also. This is of course inconsistent with the usual understanding of a binary feature such as [tense], where [+tense] and [-tense] are defined as contradictory. However, the Pasiego vowel alternations are not trivially characterised by [tense] (or [ATR]) in any case. A more precise description of each class is as follows. Penny (1969b) states that the 'tense' variants are essentially equivalent to the vowels of Castilian. The 'lax' variants of *i* and *u* are centralised and slightly open in comparison; the variant of *o* is 'similar to the French [ø] and [œ], but without any great protrusion of the lips'; the counterpart of *a* is 'very palatal (i.e. fronted and slightly raised, almost [ɛ])'. Our attempt to render these descriptions in IPA can be seen below. (Penny seems to emphasise centralisation more than lowering in the case of the high vowels, hence the symbols chosen; *ɵ* is rounded schwa.) The vowels occurring in 'lax' words are italicised and enclosed in the box, with lines connecting them to their tense counterparts. As can be seen, even ignoring *e*, it is not clear how to characterise the tense ~ lax alternation in a unitary way; the fronting and raising of *a* in particular is not obviously a centralisation or a laxing. Yet the lax subsystem itself might be nicely rationalised as a rather good (if reduced) five-vowel system, with the vowels more or less evenly distributed in the limited space. Note that if *e* alternated with a lax variant such as *ɛ*, this ideal minimal distance between vowels would be compromised, since the variant of *a* is nearly the same.

(73) Pasiego 'tense' vs. 'lax' in more detail

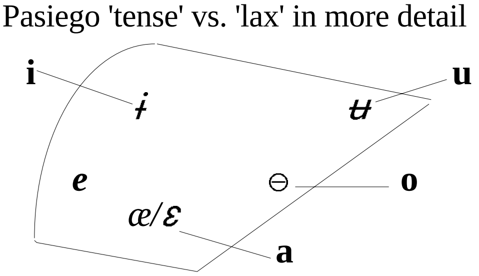

In fact, Hualde (1989) notes that in Tudanca Montañes, a nearby dialect, the tense/lax inventories are almost the same as in Pasiego, but that *e* does distinguish a lax variant, though only slightly.

Significantly, this variant is actually higher than *e*. This fact is again hard to rationalise as a laxing or centralisation from *e*, but it makes much sense in the overall system, given the crowding from below. These considerations suggest strongly that laxing cannot be understood as a featural value in the usual sense of requiring some articulatory or acoustic adjustment that is similar for each segment (i.e., each vowel is merely centralised, or articulated more laxly). Perhaps therefore laxing should be reconceptualised as a process dictating that the vowel *system* itself be reduced (a smaller minimal distance between vowels), while making no demands about the realisation of any particular vowel. Of course, other constraints would have to put these vowels in the right places, balancing the separate requirements of perceptual distance, articulatory simplicity, and identity to the tense correspondents (perhaps in related output forms). Assuming that *e* is the least complex tense vowel articulatorily (closest to the rest position), this vowel might even be the favored one to retain its identity in the new system, giving us an explanation for its behaviour. This view of the tense/lax distinction, could it be maintained, would eliminate any contradiction in stating that *e* is a member of both subsystems, and it seems a promising understanding of Pasiego. It remains to be seen whether applying this more global or dispersion theoretic thinking to this or other vowel systems might truly cut the Gordian knot of transparency, allowing 'transparent' vowels to be meaningfully analyzed as participants in both relevant domains. Therefore, for the sake of discussion, we continue below to assume that Pasiego is a problematic case, in which *e* in e.g., *kOmfesOnArjU*, is contradictorily specified for a feature that spreads through it.

In that case the form derived in (72), *kOmfEsOnArjU*, is not the surface form. Thus, another mapping to a final output is required. We consider a simplified picture of such a mapping in the tableau below, exploiting the ideas about the nature of laxness immediately above. In particular, assume that the constraint CONTRAST(LAX) regulates the minimal distance required for contrast among the vowels in lax words, clearly smaller than what is required for tense words, but still making its own demands. (A full account of such contrasts requires separate constraints determining both the number of contrasting vowels and the distance between them, Flemming 1995b.) Since *E* (*ɛ*) is not sufficiently distinct from *A* ('almost' *ɛ*, according to Penny), candidate (74)b is not optimal. In principle either vowel might give way; *A* might revert to *a* instead of *E* to *e*. Assuming that *e* is indeed the less complex vowel itself, the latter resolution is forced. Identity to the input (abstract) *E* is sacrificed.[28]

(74) Realisational mapping: Input = $[\text{kOmfEsOnArjU}]_{\text{-T}}$

$[\text{kOmfEsOnArjU}]_{\text{-T}}$	Contrast(Lax)	*a	*e	Ident-IO(lax)
a. ☞ kOmf**e**sOnArjU			*	*
b. kOmfEsOnArjU	*!			
c. kOmfEsOnarjU		*		*

A bird's-eye view of the entire representation for *kOmfesOnArjU* is given in (75). We use the word 'representation' at this juncture, rather than derivation, to emphasise the point that an account appealing to abstract subrepresentations (a helpful alternative to 'level') does not thereby invoke a serial derivation. The idea instead is that the two subrepresentations are both part of a larger overall representation of this form, while the realisational mapping defines the conditions under which these two subrepresentations can be licitly put together. See especially Orgun (1994, 1996) for an explicit statement of such a view of representations in Optimality Theory, and for its potential in accounting for cyclic effects.

(75) Representation of *kOmfesOnArjU*

[kOmfEsOnArjU]$_{-T}$	*abstracted subrepresentation where harmony is enforced*
⇕	*(realisational mapping)*
[kOm]$_{-T}$[fe]$_{+T}$[sOnArjU]$_{-T}$	*phonetically realised subrepresentation*

Though we cannot articulate a theory of such mappings here, there are important reasons for taking this view of transparency seriously—assuming the problem is genuine as discussed above—and favoring it as a hypothesis over the alternative, namely the skipping of transparent segments.

Consider two versions of the skipping idea, shown below. The 'traditional' conception of skipping is depicted in (76)a and should be familiar (e.g., Kiparsky 1981, many others, for vowel harmony). That in (76)b is found in more recent works within OT, notably Smolensky (1993), and Cole and Kisseberth (1994, 1995a). The traditional view of skipping is in fact on equal ground with the account above in requiring a realisational mapping; the sort of representation entertained—with a skipped vowel lacking any specification for the relevant feature—is not the surface form. To achieve the surface form, a feature ([+tense] here) must be inserted onto the skipped vowel, presumably splitting the multiply linked feature.

(76) Segment skipping approaches

a. Traditional:

```
kOmfesOnArjU
  \   |  / /
   \  | //
     -T
```

Also implies a realisational mapping

b. Recent: [kOm[fe]$_{+T}$sOnArjU]$_{-T}$ *No extra mapping*

In contrast, the more recent understanding of transparency seeks to account for all of the facts in one input-output mapping. Representations like the one in (76)b are taken to violate a constraint penalising the embedding of an αF domain in a -αF domain (*EMBED in Smolensky) or the failed featural expression of an F domain, where an F domain and the *realisation* of feature F are independent notions (EXPRESS in Cole and Kisseberth 1994, 1995a).

The worry for any of these views of skipping is that they sacrifice the result from section 3 above: a serious argument for strictly local spreading comes from the deep asymmetry in the long distance spreading potential of consonant and vowel place features. The idea, recall, is that consonant place gestures cannot extend across vowels without destroying syllable wellformedness. This simple account presupposes that vowels could not be skipped, but must participate in the spreading. More generally, the best hope we see for an explanation for which features spread long distance and which do not involves the same sort of appeal: the question reduces to one of whether intervening segments are compatible with the spreading feature. From this perspective, an appeal to segment skipping in vowel harmony (or anywhere) could only be attractive given a principled means of allowing only the right kind of skipping in phonology, and never the wrong kind; but we are not aware of any proposals to this effect, nor do we see much hope for them. Next to (76)b, for instance, we must exclude false long distance consonant place assimilations such as /pan/ → *pam*, as it stands easily represented by skipping, $[p[a]m]_{Lab\text{-}S}$. It will not do to suppose that such forms are ruled out because a vowel is illformed within the domain of a (stop) consonantal gesture, even if failing to 'express' this gesture. Such a move would only duplicate within a theory of embedding what we already know must be true about segment wellformedness itself, and so fails to genuinely explain the central fact, that even 'transparency' always implies compatibility.[29]

In contrast, an appeal to extra mappings and abstract subrepresentations actually preserves the basic spreading asymmetry. An attempt to construct the same false example (/pan/ → *pam*) with this kind of analysis is shown below. Assume that with the spreading of a stop Labial gesture through a syllable we have a form transcribed as shown, with an extended *p* consuming the vowel. We know already that this subrepresentation cannot be entertained by Gen, a basic claim of section 3. Since this subrepresentation is not possible, no larger representation containing it (involving any mapping between them) is possible. Note that the model we are considering assumes that forms entertained at any level, since they are regulated by constraints of OT in the usual way, must be possible by the principles of universal grammar, certainly a natural assumption, and likely the only plausible one. Crucially, at no level is locality in spreading disobeyed; instead, apparent skipping results from the imposition of markedness demands on a participating vowel (through the realisational mapping) for the purposes of surfacing.

(77) Realisational mapping cannot derive **pam*

$[p{:}m]_{Lab}$	*abstracted subrepresentation is impossible (Gen won't consider it)*
⇕	*(realisational mapping)*
$[p]_{Lab}[a][m]_{Lab}$	

We emphasise the general prediction this view of transparency makes: since 'transparent' segments are actually participants (albeit at an abstract level), they will always be compatible with the spreading feature in principle (even if the result is a relatively marked segment). If this view is correct—and the facts of place feature spreading seem to require it—then any skipping theory is at a serious intrinsic disadvantage: why should a skipped segment have to be compatible with the feature that is not associated to it?

Another reason to seriously consider abstract subrepresentations, and allow some notion of constraint opacity to preserve a coherent theory of locality, is that such things are probably required anyway, apart from any discussion of locality. (The same can hardly be said of segment skipping.) There are well known examples in phonology of constraint opacity that will not obviously succumb to any other treatment, celebrated cases including that of Yawelmani Yokuts (see Archangeli and Suzuki 1996 for recent discussion), or the behaviour of Russian *v* as though it were *w* for the purposes of voicing assimilation (Hayes 1984, Kiparsky 1985 and references therein), or Japanese *ŋ* (in some cases) as though it were *g* for the purposes of Lyman's Law (Itô and Mester to appear-a).[30] More generally, the positive advantages of parallelism demonstrated by Prince and Smolensky (1993), McCarthy and Prince (1993b) and further works in OT do not in principle preclude the possibility of extra levels of representation where these confer their own advantages. (Smolensky 1996 makes a more recent statement to this effect.) This possibility is in fact crucially assumed in various works in OT, including McCarthy and Prince (1993b), Itô, Mester and Padgett (1995), Walker (1996), and Orgun (1994, 1996), and it was seriously entertained by Prince and Smolensky (1993).

Returning to the understanding of vowel transparency proposed, we note that McCarthy (1984) argued that lax harmony in Pasiego crucially generates lax *E* at an intermediate stage of a (serial) derivation, based on facts involving the interaction of lax harmony and height harmony. Hualde (1989) presents a convincing alternative account of the height facts that undermines this particular argument, however. Hualde himself (note 21) assumes as we do that *e* participates in lax harmony, with 'phonetic realisation' accounting for its apparently tense appearance. In a discussion of transparency of *a* to [+ATR] spreading in Kinande, Archangeli and Pulleyblank (1994) similarly argue following Hyman (1989) that this vowel participates, with the phonetic component failing to distinguish a [+ATR] and [-ATR] *a*. (See also Hulst 1988, and Schlindwein 1987.) In fact the distinction *is* realised on long, low-toned *a*, giving *ə* in [+ATR] contexts; it fails to be realised otherwise. Various earlier generative accounts likewise accounted for transparency in terms of participation and 'subsequent repair' (see for instance Clements 1976b, and for more general discussion of abstractness in vowel haromony Vago 1973, 1976).

Finally, we note that this type of transparency is not in principle restricted to vowels. Given other scenarios in which a segment participating in spreading might be disfavored for reasons of markedness, we predict the same possibility. In fact, it is conceivable that some cases similar to those explored in section 4, analyzed as involving full participation in spreading, could be reanalyzed by speakers along the lines suggested here. (For example, coronal harmonies in which a learner fails to maintain the tongue tip/blade gesture throughout a domain of spreading, because its effect isn't audible in non-sibilant segments). The general prediction of the approach remains firm, however: even cases of transparency along these lines require a level of representation at which locality must be obeyed; hence 'transparent' segments must be compatible with the spreading feature.

Notes

*.We would like to thank Jill Beckman, Bruce Hayes, René Kager, Geoffrey Pullum, Rachel Walker, Laura Walsh, Mark Verhijde, and members of a class at the LOT Winter School in

Nijmegen (January 1996), whose ideas and feedback made a significant contribution to this paper. We thank Junko Itô and Armin Mester in particular, who provided incisive commentary on an earlier draft, leading to many improvements. Our debt to ideas of Edward Flemming, and also Adamantios Gafos, will be especially obvious. We also thank audiences of the Trilateral Phonology Weekend (Stanford, 1996), HILP3 (Amsterdam, 1997) and a Stanford phonology workshop. This work was supported by faculty research funds granted by the University of California, Santa Cruz, and University College Dublin.

1. Following Flemming, we use this name, despite important conceptual differences separating these ideas from the work of Lindblom.

2. Some other recent works arguably advocating a kind of strict (non-relativised) locality are Smolensky (1993) and Cole and Kisseberth (1994, 1995a). Yet these works pursue the locality and segment realisation issues in an importantly different way. See the appendix.

3. Pronunciations like [gʌk] for 'duck' and similar consonantal place 'harmonies' do occur in child production. The fact that they never surface as patterns in adult phonologies makes it all the more clear that they must be ruled out of the theory. (See Levelt 1994, Gafos 1996, and references therein on this phenomenon.)

4. An alternative explanation to what follows might posit that viable onset-nucleus or nucleus-coda sequences must have audible or salient transitions. Thus long *b: or *n: (possible results of the spreading considered) cannot compose the onset-nucleus. The explanation below is more general, but further consideration of the issue is warranted.

5. Some important issues that arise: first, if syllabic consonants occupy the syllable margin, why are they never moraic, contributing weight to a syllable? Itô (p.c.) suggests an explanation: such a moraic consonant violates the inherent prominence relation of moras in a syllable, in which the first must be 'strong' and the second 'weak' (see Zec 1988). With the 'strong' mora empty, a moraic margin would be unlawfully more prominent. Alternatively, vowel reduction in at least some cases results from lack of weight itself, hence there is no moraic consonant. Second, how do we explain the differential harmonic status of syllabic consonants of differing sonority, a fact used to motivate a family of violable *PEAK/CONSONANT constraints in Prince and Smolensky (1993) for the well known case of Berber? One possible answer is that a syllable inherits the sonority of its coda (assuming the consonant occupies this position) given a reduced nucleus. Nasals are better syllabic consonants than obstruents, in such a view, because the material in the rime is more sonorous in the former case. See Coleman (1996) for a different approach to Berber without nuclear consonants and further discussion of this issue.

6. Schein and Steriade (1986) provide a detailed account of the restrictions on the application of Nati, most of which are orthogonal to our concerns (e.g. Nati doesn't apply when there are overriding constraints on the nasal place, as when the nasal is word-final or followed by an obstruent stop).

7. This idea is reminiscent of Ohala's (1981) generalisation concerning which features dissimilate among consonants, namely those that 'spill over' into neighboring vowels. The features singled out are not quite the same. On a different note, Zoll (p.c.) informs us that (21) also correctly distinguishes the features that 'float' (the same as those that spread long-distance) from those that do not. This fact suggests that floating might receive the same account suggested of some transparencies in the appendix.

8. Absent from (21) are stricture or manner features like [continuant], [consonantal], [approximant] and [sonorant]. Though the former three spread along with place features as part of the gesture complex, it is far from clear that any of these features ever spreads independently at all, let alone long distance. (See e.g., McCarthy 1988, Padgett 1991[95] and Kaisse 1992 for discussion.) Hence they do not bear readily on the discussion at hand.

9. Compare the fundamental globality inherent in OT's selection of an optimal candidate, which precedes by elimination of competing candidate possibilities. The latter are in principle infinite in number. DT adds to this the consideration of sets of forms or segments as candidates; the number of possible sets (i.e., systems of contrast) is standardly assumed to be finite.

10. Though we refer to *binary* features here, the point being made is independent of questions of binarity vs. privativity: a single, privative feature [nasal] or [front] likewise implies a two-way contrast in the relevant dimension.

11. The relevant perceptual dimension in this case is the second formant frequency at the release of the consonant. In addition, the phonetic distinction between $[Ci]_i$ and $[C]_j i$ can be signaled by some degree of frication in the latter. In terms of F2 alone, the diagram may understate the closeness of these two forms: both have F2 origins in the neighborhood of 2 KHz for one of the authors pronouncing *bi* and *b*^{j}*i*, while that of *b*^{ɣ}*i* is approximately 1 KHz. However, though the plain and palatalised forms are too close for contrast, they clearly are phonetically distinguishable.

12. It is not a requirement of the account that such contrasts be universally ruled out; what is important is that such contrasts are at best marked, more so for the high vowels. To allow for variation, we would posit a family of varying minimal distance requirements, as in Flemming (1995b).

13. A further candidate $[f]_j$u: • $[fu:]_u$ •$[f]_ɣ$u: will lose due to violation of the minimal distance requirement. The expository conflation here of CONTRAST and MINDIST constraints (see Flemming 1995b on their separate use) obscures this point.

14. An implicit claim behind alignment is that facts of spreading require some kind of spreading imperative in the theory. For arguments that spreading cannot be due to demands of licensing or segmental markedness alone, see Padgett (to appear). This cited work also follows Steriade (1995a) and Beckman (to appear) in its attempt to reduce directionality effects to facts of phonological and morphological prominence. It is not clear that all directionality in spreading can be eliminated in this way, and for our purposes here it is convenient to stipulate directionality

within the alignment constraints in the usual way.

15. There is some ambiguity in our use of *colour* in this paper. For the purposes of determining minimal distance requirements, the perceptually-relevant second and third formants are at issue. Here we use *colour* to denote the set of [back] and [round], articulatory features. Though the latter together directly determine the formant locations, they apparently must be referred to separately in order to distinguish their behaviour in Turkish, in which [back], but not [round] spreads to non-high vowels.

16. A full account of harmony would, among other things, explain the rightward directionality. Following Beckman (to appear), we take this fact to follow from the preference given to contrast preservation in initial syllables. Implemented in terms of positional faithfulness by Beckman, the idea implies here the relativisation of CONTRAST constraints to position: an undominated CONTRAST-σ_1(COLOUR) protects the word-initial contrasts. (More precisely, a four-way colour contrast is maintained in this position, requiring constraints more demanding than the two-way contrast constraints we focus on. See Flemming 1995b.)

17. The more pervasive allophonic palatalisation of velars, seen in [-back] forms like *g^jür* 'abundant' and *k^jek^jik^j* 'thyme', makes it clear that discernibly palatalised segments can exist in front harmonic spans. Assuming some constraint that forces such allophonic palatalisation in front syllables, this constraint conspires with alignment of [-back] to ensure that velars palatalise in these words. We note here that palatalisation is contrastive in other very limited contexts in Turkish. See the references cited.

18. Failure of 'iteration' might find a genuine explanation, if we take more seriously the apparent contour nature of retroflex stops. The reason the acoustic cues to retroflexion in stops are more salient before closure lies in the manner of their articulation, involving a flap gesture starting from the most retroflex position and ending further forward, towards the site of alveolars (Steriade 1995b, Flemming 1995b and references therein). Suppose therefore that the retroflex stops are in fact necessarily contours, or perhaps bisegmental in the sense of Steriade (1994), basically [retroflex][dental] (or [low F3][high F3]). Each retroflex stop therefore contains its own rightward bound on [retroflex] spreading, namely [dental]. This idea could explain simultaneously the trigger condition, and the failure to spread beyond the first nasal, matching or exceeding the insight of the older serial account. In lieu of pursuing such ideas in our confined space, we maintain the constraint formulation in the text.

19. Retroflexion could conceivably be viewed as contrastive for non-coronal consonants, were we to analyze the *br* sequence in English *brake*, for example, as in fact a retroflex labial, i.e. *$[b]_R$ake*. In that case the analysis of Sanskrit non-coronals would more closely parallel that of the vowels. (See Steriade 1994 for an analysis of *Cr* clusters in Mazateco somewhat resembling this view.)

20. Thanks to Armin Mester for first drawing our attention to this article. Gafos (1996), also citing Allen, makes these same points, and also quotes Whitney (1889) as maintaining a similar view of locality in Nati.

21. An interesting complication of the description here that is not of direct concern involves the patterning of nasalisation in infixed forms, a patterning requiring fairly intricate phonology-morphology interaction (see the references cited).

22. Cohn (1990, 1993) notes the problem for feature theory presented by the behaviour of the glides. Though glides are held by many to be featurally identical to vowels (see section 4.2.1), they do not behave so in Sundanese nasal harmony. Cohn assumes that glides are specified [+cons] in Sundanese (cf. Piggott 1992). We refer the reader to Walker (1995, 1996) for an Optimality-Theoretic analysis of the range of nasal harmony types, which posits the nasalised markedness hierarchy shown here. This hierarchy of constraints predicts the right cross-linguistic patterning of nasal harmony.

23. The appearance of *ʔ* is predictable in Sundanese; it serves to break hiatus between like vowels or to provide an onset for certain vowel-initial morphemes (see Robins 1957 and Cohn 1990 for details). This point might suggest a sequential derivational means of deriving transparency, in which glottal stop is transparent because it is inserted after nasal spreading occurs. This approach to the permeability is unsatisfactory, for two reasons. First, permeability must be reckoned with for underlyingly present segments in other languages anyway—the account here is meant to extend to all cases. Second, this view predicts that glottal stops might just as well be blockers, since the rule ordering could be reversed. Cohn (1990) likewise rejects an ordering stipulation in favor of independently motivated assumptions concerning markedness and possible contrast (cast in feature geometric terms, see below).

24. As an intriguing aside, there are languages lacking a nasal vs. voiced obstruent contrast at the segmental level, but opposing nasal vs. oral words or harmonic spans (Anderson 1976, Piggott 1992), e.g., Guaraní, or Southern Barasano (see below). In such languages, it is common for the voiced obstruents, which occur in oral words, to be realised as prenasalised, e.g., *ⁿd*. Why should this segment emerge under these circumstances, since prenasalised segments are *more* marked than plain nasals or stops in other familiar respects? This is likely another unidimensional markedness paradox. In DT, when *d* and *n* directly contrast, adding a contrast with *ⁿd* can only be more marked, since this segment lies perceptually between the two. In the absence of a contrast however (or if the contrast is sustained by entire words), *ⁿd* should indeed be preferred to *d*, since the nasal leak facilitates obstruent voicing.

25. Since spreading through other consonants is ruled out by the ranking *NAS-CONS >> ALIGN, the ranking of CONTRAST(NAS) will not affect the patterning of these segments either. A possibility to consider, however, is that blocking vs. participation in nasal harmony could indeed be determined by the ranking of constraints on contrast, at least for some segment types. In the case of voiced obstruents, assuming that the only alternative to blocking as in *nũdag* 'pursue' is the formation of nasals as in **nũnãŋ*, contrast between voiced obstruents and nasals is clearly at stake. The facts of vowels are more intriguing. It is not clear that nasalised vowels should be ruled out on grounds of articulatory difficulty, contrary to our assumptions here. Unlike fricatives, these segments do not require oral turbulence, and it is not difficult to produce them with audible nasality. On the other hand, *place* distinctions among vowels are perceptually reduced under

nasalisation (see Beddor 1993); perhaps constraints demanding perceptual distance in place between vowels can therefore cause blocking of nasal harmony.

26. We are indebted to Jill Beckman for much enlightening discussion of the Pasiego facts and possible analyses.

27. Hualde (1989) argues instead that laxing, and the subsequent harmony, are purely phonologically induced: final high back unstressed vowels lax; only masculine singular count words end in this vowel, so that the morphological connection is incidental. Whether this is correct or not, it seems fair to note that laxing in words like *mAlU* bears the major burden of distinguishing such masculine singular count forms from their counterparts like *malu* (the related mass noun). Hence the appeal to morphological contrast seems justified.

28. The use of an input, implied by the notion of an abstract level of representation, might seem contradictory to our tentative avoidance of underlying representations earlier. This is not the case. The point earlier was that the existence of output constraints regulating contrast, as well as similarity among related morphemes, undermines the motivation for underlying representations as determinants of phonological wellformedness in any systematic way. What we are considering here is what would have been called an 'intermediate' representation, the motivation for which is something different: the resolution of opacity.

29. These same criticisms can be applied to Pulleyblank's (1994a) proposed understanding of alignment, effectively generating segment skipping, and to the related notion of 'contextual transparency' in Archangeli and Pulleyblank (1994, section 4.7).

30. Many cases of apparent constraint opacity can be effectively treated in terms of correspondence or identity between morphologically related output forms (see references in section 4.1). The cases we are contemplating, however, cannot be treated in this way. Another often workable alternative to subrepresentations is to grant constraints the power to state conditions or environments on both input and output forms, (Koskeniemmi 1983, Lakoff 1993, McCarthy to appear, Cole and Kisseberth 1995b, compare the role of containment in early versions of OT, Prince and Smolensky 1993). Here we see an old trade-off in phonology: either the number of representational levels can be minimised, or the number and kinds of constraints/rules—and their potential unnaturalness—can be, but not both, since the two demands conflict. See Archangeli and Suzuki (1996) for an argument that even empowering constraints in this way cannot obviate 'intermediate' levels of representation.

References

Alderete, John. (to appear). Dissimilation as local conjunction. *NELS 27*.

Allen, W. (1951). Some prosodic aspects of retroflexion and aspiration in Sanskrit. *Bulletin of the School of Oriental and African Studies* 13. 939-946.

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

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

Anderson, Stephen R. (1976). Nasal consonants and the internal structure of segments. *Language* 52. 326-44.

Anderson, Stephen R. (1980). Problems and perspectives in the description of vowel harmony. In R.M. Vago (ed.) *Issues in vowel harmony*. Amsterdam: John Benjamins.1-48.

Archangeli, Diana & Douglas Pulleyblank. (1986). The content and structure of phonological representations. Ms., University of Arizona and University of British Columbia.

Archangeli, Diana & Douglas Pulleyblank. (1987). Minimal and maximal rules: Effects of tier scansion. *NELS 17.* 16-35.

Archangeli, Diana & Douglas Pulleyblank. (1994). *Grounded phonology*. Cambridge, Massachusetts: MIT Press.

Archangeli, Diana & Keiichiro Suzuki. (1996). Yokuts templates: correspondence to neither input nor output. In B. Agbayani & N. Harada (eds.) *UCI Working Papers in Linguistics*. Irvine CA: Irvine Linguistics Students Association, University of California, Irvine.

Battistella, Ed. (1979). The revised directional application theory. Indiana University Linguistics Club.

Beckman, Jill. (to appear). Positional faithfulness, positional neutralisation, and Shona vowel harmony. *Phonology*.

Beckman, S. Urbanczyk & L. Walsh (eds.) (1995). *University of Massachusetts Occasional Papers in Linguistics UMOP 18*. Amherst, MA: GLSA.

Beddor, Patrice. (1993). The perception of nasal vowels. In M. K. Huffman & R. A. Krakow (eds.) *Nasals, nasalisation, and the velum*. New York: Academic Press.

Bell, Alan. (1978). Syllabic consonants. In J. H. Greenberg (ed.) *Universals of human language*. Stanford: Stanford University Press. 153-201.

Benthem, Johan van. (1983). *The logic of time - a model-theoretic investigation into the varieties of temporal ontology and temporal discourse*. Dordrecht: Reidel.

Benua, Laura. (1995). Identity effects in morphological truncation. In J. Beckman, S. Urbanczyk & L. Walsh (eds.) 77-136.

Bird, Steven & Ewan Klein. (1990). Phonological events. *Journal of Linguistics* 26. 33-56.

Bladon, R & F. Nolan. (1977). A video-fluorographic investigation of tip and blade alveolars in English. *Journal of Phonetics* 5.185-193.

Boyce, S. (1990). Coarticulatory organisation for lip rounding in Turkish and English. *Journal of the Acoustical Society of America* 88. 2584-2595.

Browman, Catherine & Louis Goldstein. (1986). Towards an articulatory phonology. *Phonology Yearbook* 3. 219-252.

Browman, Catherine & Louis Goldstein. (1989). Articulatory gestures as phonological units. *Phonology* 6. 201-251.
Browman, Catherine & Louis Goldstein. (1990). Tiers in articulatory phonology, with some implications for casual speech. In J. Kingston & M. Beckman (eds.) 341-397.
Buckley, Eugene. (1995). Cyclicity as correspondence. Ms, University of Pennsylvania, Philadelphia, PA.
Burzio, Luigi. (1994). *Principles of English stress*. Cambridge: University of Cambridge Press.
Burzio, Luigi. (1996). Surface constraints versus underlying representations. In J. Durand & B. Laks (eds.)
Carney, P.J. & K.L. Moll. (1971). A cinefluorographic investigation of fricative consonant-vowel coarticulation. *Phonetica* 23. 193-202.
Chomsky, Noam & Morris Halle. (1968). *The sound pattern of English*. New York: Harper & Row.
Clements, G. N. (1976a). *Vowel harmony in nonlinear generative phonology: an autosegmental model*. Distributed 1980, Indiana University Linguistics Club.
Clements, G. N. (1976b). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. *NELS 7*. 49-64.
Clements, G. N. (1985). The geometry of phonological features. *Phonology Yearbook* 2. 225-252.
Clements, G. N. (1990). The role of the sonority cycle in core syllabification. In J. Kingston & M. Beckman (eds.) 283-333.
Clements, G. N. (1991). Place of articulation in consonants and vowels: a unified theory. In *Working Papers of the Cornell Phonetics Laboratory*. Ithaca: Cornell University.
Clements, G.N. & Elizabeth V. Hume. (1995). The internal organisation of speech sounds. In J. A. Goldsmith (ed.) *The handbook of phonological theory*. Cambridge, MA: Blackwell. 245-306.
Clements, G. N. & Engin Sezer. (1982). Vowel and consonant disharmony in Turkish. In H. van der Hulst & N. Smith (eds.) 213-256.
Cohn, Abigail. (1990). Phonetic and phonological rules of nasalisation. PhD dissertation, UCLA, Los Angeles.
Cohn, Abigail. (1992). The consequences of dissimilation in Sundanese. *Phonology* 9. 199-220.
Cohn, Abigail. (1993). The status of nasalised continuants. In M. Huffman & R. Krakow (eds.) 329-367.
Cole, Jennifer & Charles W. Kisseberth. (1994). An optimal domains theory of harmony. Urbana-Champaign: Cognitive Science Technical Report UIUC-BI-CS-94-02, University of Illinois.
Cole, Jennifer & Charles W. Kisseberth. (1995a). Nasal harmony in optimal domains theory. Urbana-Champaign, IL: Cognitive science technical report UIUC-BI-CS-95-02, University of Illinois.
Cole, Jennifer S. & Charles W. Kisseberth. (1995b). Restricting multi-level constraint evaluation: opaque rule interaction in Yawelmani vowel harmony. Ms, University of Illinois at Urbana-Champaign.
Coleman, John. (1996). Declarative syllabification in Tashlhit Berber. In J. Durand & B. Laks (eds.)

Coleman, John & John Local. (1991). The "no crossing constraint" in autosegmental phonology. *Linguistics and Philosophy* 14. 295-338.

Durand, Jacques & Bernard Laks (eds.) (1996). *Current trends in phonology: models and methods*. Salford, Manchester: European Studies Research Institute.

Farnetani, Edda, Kyriaki Vagges & Emanuela Magno-Caldognetto. (1985). Coarticulation in Italian /VtV/ sequences: a palatographic study. *Phonetica* 42. 78-99.

Flemming, Edward. (1993). The role of metrical structure in segmental rules. MA thesis, UCLA.

Flemming, Edward. (1995a). Vowels undergo consonant harmony. Paper presented at the Trilateral Phonology Weekend, Stanford University.

Flemming, Edward. (1995b). Auditory representations in phonology. PhD dissertation, UCLA.

Flemming, Edward. (to appear). Phonetic detail in phonology: evidence from assimilation and coarticulation. In *Proceedings of the Arizona Conference: Workshop on Features in Optimality Theory*. Coyote Working Papers, University of Arizona, Tuscon.

Fowler, C.A. (1983). Converging sources of evidence on spoken and perceived rhythms of speech: cyclic production of vowels in sequences of monosyllabic stress feet. *Journal of Experimental Psychology: General* 112. 386-412.

Frisch, Stefan. (1996). *Similarity and frequency in phonology*. PhD dissertation, Northwestern University, Evanston, IL.

Gafos, Adamantios. (1996). The articulatory basis of locality in phonology. PhD dissertation, Johns Hopkins University, Baltimore, MD.

Gerfen, Chip. (1996). *Topics in the phonology and phonetics of Coatzospan Mixtec*. PhD dissertation, University of Arizona.

Gnanadesikan, Amalia. (1994). The geometry of coronal specifications. *NELS 24*. 125-139.

Goldsmith, John. (1976). Autosegmental phonology. PhD dissertation, MIT. Published 1979, New York: Garland Press.

Goldsmith, John A. (ed.) *The last phonological rule: reflections on constraints and derivations*. Chicago: University of Chicago Press.

Goldsmith, John A. (ed.) (1995) *The handbook of phonological theory*. Cambridge, MA: Blackwell.

Halle, Morris & Jean-Roger Vergnaud. (1978). *Metrical structures in phonology*. Cambridge, MA: MIT Press.

Hammond, Michael. (1988). On deriving the well-formedness condition. *Linguistic Inquiry* 19. 319-325.

Hayes, Bruce. (1984). The phonetics and phonology of Russian voicing assimilation. In M. Aronoff & R. T. Oehrle (eds.) *Language sound structure*. Cambridge, MA: MIT Press. 318-328.

Herzallah, R. (1990). Aspects of Palestinian Arabic phonology: A nonlinear approach. PhD dissertation, Cornell University.

Hombert, Jean-Marie, John J. Ohala & William G. Ewan. (1979). Phonetic explanations for the development of tones. *Language* 55. 37-58.

Howard, Irwin. (1973). *A directional theory of rule application in phonology*. PhD dissertation, MIT. Distributed by the Indiana University Linguistics Club.

Hualde, Jose Ignacio. (1989). Autosegmental and metrical spreading in the vowel-harmony systems of northwestern Spain. *Linguistics* 27. 773-805.
Huffman, Marie K. and R. A. Krakow (eds.) (1993). *Nasals, nasalisation, and the velum*. New York: Academic Press.
Hulst, Harry van der (1988). The geometry of vocalic features. In H. van der Hulst & N. Smith (eds.) *Features, segmental structure and harmony processes*. Dordrecht: Foris. 77-126.
Hulst, Harry van der & N. Smith (eds.) (1982). *The structure of phonological representations*. Dordrecht: Foris.
Hulst, Harry van der & Jeroen van de Weijer. (1991). Topics in Turkish phonology. In H. Boeschoten & L. Verhoeven (eds.) *Turkish linguistics today*. E.J.Brill. 11-59.
Hulst, Harry van der & Jeroen van de Weijer. (1995). Vowel harmony. In J. A. Goldsmith (ed.) 495-534.
Hyman, Larry M. (1989). Advanced tongue root in Kinande. Ms, UC Berkeley.
Hyman, Larry M. & Armindo Ngunga. (1994). On the non-universality of tonal association 'conventions': evidence from Ciyao. *Phonology* 11. 25-68.
Itô, Junko & Armin Mester. (to appear-a). Correspondence and compositionality: the ga-gyo variation in Japanese phonology. In I. Roca (ed.) *Derivations and constraints in phonology*. Oxford: Oxford University Press.
Itô, Junko & Armin Mester. (to appear-b). Structural economy and OCP interactions in local domains. In V. Samiian (ed.) *Proceedings of the Western Conference on Linguistics*. Fresno, CA: Department of Linguistics, California State University.
Itô, Junko, Armin Mester & Jaye Padgett. (1995). Licensing and underspecification in Optimality Theory. *Linguistic Inquiry* 26. 571-614.
Jensen, John T. & Margaret Stong-Jensen. (1976). Ordering and directionality of iterative rules. In A. Koutsoudas (ed.) *The application and ordering of grammatical rules*. The Hague: Mouton.
Johnson, C Douglas. (1972). *Formal aspects of phonological description*. The Hague: Mouton.
Jun, Jongho. (1995). Place assimilation as the result of conflicting perceptual and articulatory constraints: Paper presented at WCCFL 14.
Kager, René. (1996). Surface opacity of metrical structure in optimality theory. Utrecht: OTS working paper.
Kaisse, Ellen. (1992). Can [consonantal] spread? *Language* 68. 313-332.
Kaye, Jonathan & Jean Lowenstamm. (1984). De la syllabicité. In F. Dell, D. Hirst & J.-R. Vergnaud (eds.) *Forme sonore du language: structure des représentations en phonologie*. Paris: Hermann.
Kenstowicz, Michael. (1995). Cyclic vs. non-cyclic constraint evaluation. *Phonology* 12. 397-436.
Kenstowicz, Michael. (1996). Base identity and uniform exponence: alternatives to cyclicity. In J. Durand & B. Laks (eds.)
Kenstowicz, Michael & Charles Kisseberth. (1977). *Topics in phonological theory*. New York: Academic Press.
Kenstowicz, Michael & Charles Kisseberth. (1979). *Generative phonology: description and theory*. New York: Academic Press.

Kingston, John & Mary Beckman (eds.) (1990). *Papers in laboratory phonology 1: between the grammar and physics of speech*. New York: Cambridge University Press.
Kingston, John & Randy L. Diehl. (1994). Phonetic knowledge. *Language* 70. 419-454.
Kiparsky, Paul. (1981). Vowel harmony. Ms, MIT.
Kiparsky, Paul. (1985). Some consequences of Lexical Phonology. *Phonology* 2. 85-138.
Kirchner, Robert. (1993). Round and back vowel harmony and disharmony: An optimality theoretic account. Ms, UCLA.
Koskenniemi, Kimmo. (1983). *Two-level morphology: a general computational model for word-form recognition and production*. Helsinki, Finland: Publication no. 11, Department of General Linguistics, University of Helsinki.
Ladefoged, Peter & Ian Maddieson. (1996). *The sounds of the world's languages*. Cambridge, MA: Blackwell.
Lakoff, George. (1993). Cognitive phonology. In J. Goldsmith (ed.) 117-145.
Lees, Robert. (1961). *The phonology of Modern Standard Turkish, Uralic and Altaic series, vol. 6*. Bloomington: Indiana University Publications.
Levelt, Clara C. (1994). *On the acquisition of place*. Dordrecht: Holland Institute of Generative Linguistics dissertations.
Levergood, Barbara. (1984). Rule governed vowel harmony and the strict cycle. *NELS 14*.275-293.
Levin, Juliette. (1985). A metrical theory of syllabicity. PhD dissertation, MIT.
Liljencrants, Johan & Björn Lindblom. (1972). Numerical simulation of vowel quality systems: the role of perceptual contrast. *Language* 48. 839-862.
Lindblom, Björn. (1986). Phonetic universals in vowel systems. In J. Ohala & J. Jaeger (eds.) *Experimental phonology*. Orlando: Academic Press. 13-44.
Lindblom, Björn. (1990). On the notion "possible speech sound". *Journal of Phonetics* 18. 135-152.
McCarthy, John. (1979). Formal problems in Semitic phonology and morphology. PhD dissertation, MIT.
McCarthy, John. (1984). Theoretical consequences of Montañes vowel harmony. *Linguistic Inquiry* 15. 291-318.
McCarthy, John. (1986). OCP effects: gemination and antigemination. *Linguistic Inquiry* 17. 207-263.
McCarthy, John. (1988). Feature geometry and dependency: A review. *Phonetica* 43. 84-108.
McCarthy, John. (1991). Semitic gutturals and distinctive feature theory. In M. Eid & B. Comrie (eds.) *Perspectives on Arabic linguistics*. Philadelphia: John Benjamins. 63-91.
McCarthy, John J. (1994). On coronal 'transparency'. Paper presented at the Annual Trilateral Phonology Weekend, University of California, Santa Cruz.
McCarthy, John J. (1995). Extensions of faithfulness: Rotuman revisited. Ms, University of Massachusetts, Amherst.
McCarthy, John J. (to appear). Remarks on phonological opacity in Optimality Theory. In J. Lecarme, J. Lowenstamm & U. Shlonsky (eds.) *Proceedings of the Second Colloquium on Afro-Asiatic Linguistics*.

McCarthy, John J & Alan S. Prince. (1993a). Generalised alignment. In G. Booij & J. v. Marle (eds.) *Yearbook of Morphology*. 79-153.
McCarthy, John J. & Alan S. Prince. (1993b). Prosodic morphology I: constraint interaction and satisfaction. Ms, University of Massachusetts, Amherst & Rutgers University.
McCarthy, John J. & Alan S. Prince. (1995). Faithfulness and reduplicative identity. In J. Beckman, S. Urbanczyk & L. Walsh (eds.) 249-384.
Mester, Armin. (1986). Studies in tier structure. PhD dissertation, University of Massachusetts, Amherst.
Mohanan, K. P. (1993). Fields of attraction in phonology. In J. Goldsmith (ed.) 61-116.
Ní Chiosáin, Máire. (1994). Irish palatalisation and the representation of place features. *Phonology* 11. 89-106.
Ní Chiosáin, Máire & Jaye Padgett. (1993). Inherent VPlace: Report No. LRC-93-09, Linguistics Research Center, University of California, Santa Cruz.
Odden, David. (1991). Vowel geometry. *Phonology* 8. 261-289.
Odden, David. (1994). Adjacency parameters in phonology. *Language* 70. 289-330.
Ogden, Richard. (1995). An exploration of phonetic exponency in Firthian prosodic analysis: form and substance in Finnish phonology. PhD dissertation, University of York.
Ogden, Richard & John Local. (1994). Disentangling autosegments from prosodies; a note on the misrepresentation of a research tradition in phonology. *Journal of Linguistics* 30. 477-498.
Ohala, John J. (1975). Phonetic explanations for nasal sound patterns. In C. A. Ferguson, L. A. Hyman & J. J. Ohala (eds.) *Nasálfest: papers from a symposium on nasals and nasalisation*. Stanford, CA: Linguistics Department, Stanford University.
Ohala, John. (1978). The production of tone. In V. Fromkin (ed.) *Tone: a linguistic survey*. New York: Academic Press. 3-39.
Ohala, John J. (1981). The listener as a source of sound change. *CLS* :2
Ohala, John J. (1990). There is no interface between phonology and phonetics: a personal view. *Journal of Phonetics* 18. 153-171.
Ohala, John J. & Manjari Ohala. (1993). The phonetics of nasal phonology: theorems and data. In M. K. Huffman & R. A. Krakow (eds.)
Öhman, S.E.G. (1966). Coarticulation in VCV utturances: spectrographic measurements. *Journal of the Acoustical Society of America* 39. 151-168.
Orgun, Orhan. (1993). Alterable geminates and the release node. Ms, UC Berkeley. Paper presented at the LSA annual meeting, 1993.
Orgun, Cecil Orhan. (1994). Monotonic cyclicity and Optimality Theory. *NELS* 24. 461-474.
Orgun, Orhan. (1996). Sign-based morphology and phonology. PhD dissertation, University of California, Berkeley, Berkeley, CA.
Padgett, Jaye. (1991). *Stricture in feature geometry*: PhD dissertation, UMass, Amherst [Revised version published 1995, CSLI Press, Stanford, CA].
Padgett, Jaye. (1994). Stricture and nasal place assimilation. *NLLT* 12. 465-513.
Padgett, Jaye. (1995). Feature classes. In J. Beckman, S. Urbanczyk & L. Walsh (eds.) 285-420.
Padgett, Jaye. (to appear). Partial class behavior and nasal place assimilation. In *Proceedings of the Arizona Phonology Conference: Workshop on Features in Optimality Theory*. Coyote Working Papers, University of Arizona, Tucson.

Paradis, Carole, and Jean-François Prunet. (1989). On coronal transparency. *Phonology* 6. 317-348.
Penny, Ralph. (1969a). *El habla pasiega: ensayo de dialectología montañesa*. London: Tamesis Books.
Penny, Ralph. (1969b). Vowel harmony in the speech of the Montes de Pas (Santander). *Orbis* 18. 148-166.
Penny, Ralph. (1970). Mass nouns and metaphony in the dialects of north-western Spain. *Archivum Linguisticum* 1. 21-30.
Pierrehumbert, Janet. (1993). Dissimilarity in the Arabic verbal roots. In *NELS 23*.
Pierrehumbert, Janet & Mary Beckman. (1988). *Japanese tone structure*. LI Monograph Series No. 15. Cambridge, MA: MIT Press.
Piggott, Glyne. (1992). Variability in feature dependency: the case of nasality. *NLLT* 10. 33-78.
Poser, William. (1982). Phonological representations and action-at-a-distance. In H. van der Hulst & N. Smith (eds.). 121-158.
Prince, Alan. (1987). Planes and copying. *Linguistic Inquiry* 18. 491-510.
Prince, Alan, and Paul Smolensky. (1993). Optimality Theory: constraint interaction in generative grammar. Ms, Rutgers University and University of Colorado, Boulder.
Pulleyblank, Douglas. (1989). Patterns of feature cooccurrence: the case of nasality. In *Arizona Phonology Conference 2: Coyote Working Papers 9*. Tuscon, AZ: University of Arizona.
Pulleyblank, Douglas. (1994a). Neutral vowels in Optimality Theory: a comparison of Yoruba and Wolof. Ms, University of British Columbia, Vancouver.
Pulleyblank, Douglas. (1994b). Vowel harmony in Optimality Theory. Ms., University of British Columbia [for Proceedings of the Workshop on Phonology, University of Coimbra, Portugal, September 1993].
Pullum, Geoffrey K. & Rachel Walker. (1996). Possible and impossible segments. Ms, UC Santa Cruz.
Ringen, Catherine O. (1976). Vacuous application, iterative application, reapplication, and the unordered rule hypothesis. In A. Koutsoudas (ed.) *The application and ordering of grammatical rules*. The Hague: Mouton. 55-75.
Robins, R.H. (1957). Vowel nasality in Sundanese. In *Studies in linguistic analysis*. London: Basil Blackwell. 87-103.
Sagey, Elizabeth. (1986). The representation of features and relations in nonlinear phonology. PhD dissertation, MIT.
Sagey, Elizabeth. (1988). On the ill-formedness of crossing association lines. *Linguistic Inquiry* 19. 109-118.
Schein, Barry & Donca Steriade. (1986). On geminates. *Linguistic Inquiry* 17. 691-744.
Schlindwein, Deborah. (1987). P-bearing units: A study of Kinande vowel harmony. *NELS* 17. 551-568.
Schourup, Lawrence. (1972). Characteristics of vowel nasalisation. *Papers in Linguistics* 5. 530-548.
Scobbie, James. (1991). Attribute value phonology. PhD dissertation, University of Edinburgh.
Selkirk, Elisabeth. (1984). On the major class features and syllable theory. In M. Aronoff & R. T. Oehrle (eds.) *Language sound structures*. Cambridge, MA: MIT Press. 107-136.

Selkirk, E.O. (1991). Major place in the vowel space: vowel height features. Ms., University of Massachusetts.
Shaw, Patricia. (1991). Consonant harmony systems: the special status of coronal harmony. In C. Paradis & J. F. Prunet (eds.) *The special status of coronals*. New York: Academic Press. 125-157.
Smolensky, Paul. (1993). Optimality, markedness, and underspecification. Paper presented at the Rutgers University Optimality Workshop, New Brunswick, NJ.
Smolensky, Paul. (1996). On the explanatory adequacy of Optimality Theory. Handout of talk presented at Stanford University, June 3, 1996.
Steriade, Donca. (1986). A note on coronal. Ms, MIT.
Steriade, Donca. (1987). Locality conditions and feature geometry. *NELS* 17. 595-617.
Steriade, Donca. (1994). Complex onsets as single segments: the Mazateco pattern. In J. Cole & C. Kisseberth (eds.) *Perspectives in phonology*. Stanford: CSLI Publications.
Steriade, Donca. (1995a). Underspecification and markedness. In J. Goldsmith (ed.) 114-174.
Steriade, Donca. (1995b). Positional neutralisation: Ms, UCLA.
Stevens, Kenneth N., Samuel J. Keyser & Haruko Kawasaki. (1986). Toward a phonetic and phonological theory of redundant features. In J. S. Perkell & D. H. Klatt (eds.) *Invariance and variability in speech processes*. Hillsdale, NJ: Lawrence Erlbaum. 426-449.
Svantesson, Jan-Olof. (1989). Tonogenetic mechanisms in northern Mon-Khmer. *Phonetica* 46. 60-79.
Thurgood, Graham. (1996). Language contact and the direction of internal drift: the development of tones and registers in Chamic. *Language* 72. 1-31.
Trubetzkoy, N. (1939). *Grundzüge der Phonologie*. Güttingen: Vandenhoeck and Ruprecht.
Vago, Robert. (1973). Abstract vowel harmony systems in Uralic and Altaic languages. *Language* 49. 579-605.
Vago, Robert. (1976). Theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry* 7. 242-263.
Vago, Robert. (1988). Underspecification in the height harmony system of Pasiego. *Phonology* 5. 343-362.
Walker, Rachel. (1995). Hierarchical opacity effects in nasal harmony. In J. Fuller, H. Han & D. Parkinson (eds.) *Proceedings of the Eleventh Eastern States Conference on Linguistics*. Ithaca, NY: DMLL Publications. 318-329.
Walker, Rachel. (1996). Neutral segments and locality. Ms, University of California, Santa Cruz.
Walker, Rachel & Geoffrey K. Pullum. (1996). Possible and impossible segments. Santa Cruz, CA: Ms, University of California, Santa Cruz.
Waterson, Natalie. (1956). Some aspects of the phonology of the nominal forms of the Turkish word. *Bulletin of the School of Oriental and African Studies* 18. 578-591.
Whitney, W. D. (1889). *Sanskrit grammar*. Cambridge, MA: Harvard University Press.
Yip, Moira. (1989). Feature geometry and cooccurrence restrictions. *Phonology* 6. 349-374.
Zec, Draga. (1988). Sonority constraints on prosodic structure. PhD dissertation, Stanford University.
Zoll, Cheryl. (1996). *Parsing below the segment in a constraint based framework*. PhD dissertation, UC Berkeley.

Please address comments to both authors:

Máire Ní Chiosáin
Department of Linguistics
University College Dublin
Belfield, Dublin 4
Ireland
nichiosain@ucd.ie

Jaye Padgett
Department of Linguistics
Stevenson College
UC Santa Cruz
Santa Cruz, CA 95064
padgett@ling.ucsc.edu