Monovalency and opacity: Chicheŵa height harmony^{*}

JOHN HARRIS

1 Introduction

When it comes to providing insights into the nature of segmental structure, the kinks in phonological systems can often be every bit as revealing as the symmetries. Signs of crookedness can typically be put down to the uneven manner in which segmental primes congregate. As is well known, some logically possible pairs of primes display at best a certain reticence to combine, while others steadfastly refuse to have any truck with one another whatsoever. A selection of such uneasy bedfellows might include the following, ranked in ascending order of mutual repulsion: voiced and obstruent, round and palatal, ATR and open, ATR and nasal.

Such incompatibilities are perhaps at their starkest in vowel harmony. One of the intriguing quirks of harmony systems is the consistently skewed distribution of neutral vowels — that is, those lacking harmonic counterparts. Either such vowels behave transparently; that is, although they do not alternate, they fail to halt the advance of harmony across a domain. Or they behave opaquely, which case they not only resist alternation but also block the propagation of harmony. Specific questions arising in connection with these types of segments include the following. Why is it that, if a palatal harmony system exhibits transparency, it is the palatal vowels the typically, perhaps exclusively, have this property? Why in ATR systems is it typically *a* that acts as opaque? Why does the finger also consistently point at *a* when height harmony systems exhibit opacity? (For a useful summary and discussion of such patterns, see van der Hulst & Smith 1986.)

Facts such as these, it is now widely acknowledged, are intimately bound up with the observation that a typical harmony process (perhaps the only type of process) targets only one term of a phonological opposition. In frameworks incorporating two-valued features, asymmetries of this sort have for some time been captured through underspecification — that is, by having only one value of each distinctive feature specified in lexical representation (e.g. Kiparsky 1982, Archangeli 1984, Pulleyblank

^{*}My thinking on the theoretical issues raised in this paper has benefited greatly from discussions with Francis Moto and Al Mtenje, to whom I'm also grateful for the data. Thanks also to Wiebke Brockhaus and John Goldsmith for their insightful comments on an earlier draft. Naturally, none of these people bears any responsibility for the views expressed here.

1986). A more radical alternative is adopted in frameworks which assume that phonological primes are single-valued (e.g. Anderson & Jones 1974, Schane 1984, Kaye, Lowenstamm & Vergnaud 1985). Both types of approach define situations in which only one term of a phonological opposition is harmonically active. The question then becomes a matter of identifying which term is the active one. With underspecification, which value of a bivalent feature is lexically selected is something for individual grammars to decide (although one of the choices may be favoured by markedness conventions). Under this approach, harmonic asymmetries are thus no more than contingent facts; in principle, either value of a contrast could be harmonically active or inert. Under a monovalent approach, in contrast, such cases of lopsided behaviour are necessary facts; since only one term of each opposition possesses the distinguishing prime, it is perforce the only one available for harmonic access.

For a long while, bivalency and monovalency led more or less separate lives in the phonological literature. It is only relatively recently that the two approaches have found themselves appearing together in discussions of segmental structure. Even then, the meetings have often been no more than rather casual affairs which have ended with the two staking out separate representational territories. This is perhaps most obvious in the now-classic model of feature geometry, in which monovalency asserts itself in the form of organisational or class nodes, while bivalency continues to rule the roost at the level of features proper, the terminal nodes of the melodic hierarchy (Clements 1985, McCarthy 1988). But at some point, it has to be acknowledged that a full-blooded version of monovalency lays claim to the whole realm of segmental representation, even if much of the relevant literature continues to fight shy of the confrontation with bivalency that this inevitably implies. Only rarely do we find direct and detailed comparisons of the two types of approach (e.g. Pulleyblank, in press). This article plays host to just such a head-to-head.

For various reasons to be expanded on presently, the particular pattern of Bantu height harmony we focus on here is a suitable arena in which to stage the confrontation. The protagonists in the present contest are well established representatives of respectively bivalency and monovalency: feature-underspecification versus a model in which single-valued primes combine to form headed expressions. The former type of analysis, couched in terms of orthodox features, treats the height pattern as the spreading of [-low]; the latter treats it as the spreading of the element [A]. In several respects, the two accounts show themselves to be evenly matched. Nevertheless, they can be clearly separated on a number of counts, particularly with respect to one issue that is the main focus of this article: how they cope with the opacity of *a* in the relevant systems. Monovalent analyses of Bantu height harmony have already been proposed, at least in broad outline, by Goldsmith

(1985), Rennison (1987) and Harris & Moto (1989). Only in the last of these is monovalency, as here, directly pitted against bivalent alternatives. Moreover, none of these studies, it can be argued, provides a satisfactory account of opaque a.

§2 provides a brief reminder of how opacity is typically treated from a bivalent feature perspective. In §3, we outline the facts of Bantu height harmony, drawing on one illustrative system, that of Chicheŵa. In §4, we discuss earlier two-valued analyses of the phenomenon and show how these are internally flawed. In §5, we present and reject an underspecification analysis, focusing on the unsatisfactory nature of the various types of parochial condition to which the theory must resort in order to capture opacity. In §6, we show how such mechanisms can be dispensed with in a monovalent treatment which makes appeal to intra-segmental dependency. We conclude in §7 with some general remarks on how both opacity and transparency effects might be prosodically derived within this framework.

2 Antagonistic opacity

Investigating Bantu height harmony provides us with the opportunity to reflect on the validity of a number of devices which figure prominently in analyses of opaque vowels. In classic non-linear accounts, appeal is made to a ban on the crossing of association lines, the impact of which is to interrupt the spreading of some autosegment [α F] when its path is blocked by the occurrence of a segment bearing the opposite specification. The overall effect of the no-crossing constraint, itself derivable from general principles of precedence and locality (Sagey 1988, Archangeli & Pulleyblank 1992: 20 ff), is to rule out gapped configurations in which a multiple association skips over an intervening vowel (van der Hulst & Smith 1986). In current underspecification theory, however, the apparent attractiveness of the constraint is severely diminished by the requirement that only one value of a bivalent feature be represented underlyingly. In order for no-crossing to deliver opacity, it is necessary for the missing value to be assigned by the time harmony takes place. However, as Archangeli & Pulleyblank point out, there is frequently no motivation independent of opacity itself that such values should be present during derivation (1992: 311-312).

An alternative means of deriving opacity appeals to featural antagonisms of the sort mentioned at the outset of this article. In a rule-based approach, this is achieved through context-sensitivity: the harmonic behaviour of one particular feature value is made dependent on the presence of some other feature value in the relevant rule's structural description. Such a rule typically takes the form given in ?a.

(1) (a) $[\alpha F]$ (b) $[\alpha F]$

That is, the feature value $[\alpha F]$ spreads onto a neighbouring vowel, only if the latter is specified for the feature value $[\beta G]$. Antagonistic opacity (the term is Archangeli & Pulleyblank's (1992: 313)) arises where a vowel is specified as $[-\beta G]$ and is thus in a position to resist the harmonic advances of $[\alpha F]$. As shown in ?b, not only will the $[-\beta G]$ vowel V₃ fail to harmonise but, under locality, it will also block the further propagation of $[\alpha F]$ to any vowel (V₄) which otherwise satisfies the conditions of the rule.

Allowing for feature sensitivity grossly inflates the expressive power of a rulebased model. In principle, any combination of feature specifications is free to be attached to the target of a spreading rule. This predicts a host of non-occurring processes and misses the point that only a very limited set of quite specific voweltypes function opaquely in harmony systems. Moreover, the device is at variance with the recent move towards constraint-based treatments of phonological derivation. It is in the spirit of this general approach to assume that processes occur freely wherever their conditions are met. Any restrictions on this occurrence are then attributable not to parochial stipulations built into the formalisation of individual processes but to the intervention of independent and ideally quite general constraints. This implies that harmony should be simply statable as spread $[\alpha F]$. Opacity is then due to the operation of some separate constraint which blocks the association of the harmonic category to a particular class of vowel. In current feature-based approaches, a constraint of this sort is typically formulated as a filter which bans the co-association of particular feature values, in this case * $[\alpha F, -\beta G]$ (e.g. Kiparsky 1985, Calabrese 1989, Archangeli & Pulleyblank 1992), or as some form of feature licensing restriction (e.g. Itô, Mester & Padgett 1994).

Filters of this type are potentially open to the same criticism as that levelled above at feature-sensitive rules: there is in principle no restriction on what feature values can co-occur in a filter. In fact, only a very narrow subset of the extremely large set of possible feature combinations that are expressible by means of filters actually corresponds to observed restrictions. In other words, the recurrent blocking behaviour of particular vowels in different types of harmony system is accidental under this kind of account.

The damage of this criticism can be mitigated by demonstrating that a given filter has some independent motivation beyond the opacity effect it delivers — for example, by showing that the device also helps define the set of contrasts in the system. As we will see below, this expedient is simply not available to an analysis of

Bantu height harmony that is couched in terms of the orthodox features $[\pm high]$ and $[\pm low]$. In fact, if we seek support in this direction, we find exactly the opposite of the desired result. The particular problem in this case is that the two feature values which the relevant filter defines as antagonistic combine quite happily in the full specification of vocalic contrasts.

We will contrast this feature-underspecification analysis with one formulated within a framework in which all segmental primes (elements) are monovalent and enter into dependency relations within melodic expressions. Under this approach, the combinability of elements is determined directly by the organisation of autosegmental structure rather than indirectly by anything resembling filters. Applied to the particular Bantu system under discussion, we will argue, this account successfully achieves the twin goals outlined in the previous paragraph: not only does it inhibit spreading past *a*, but it also determines the shape of the relevant vowel systems. Far from being some ad hoc ploy solely designed to cope with the specific case of opacity under discussion, the notion of intrasegmental dependency is an intrinsic design property of the model and can be shown to be independently involved in a range of apparently unrelated melodic phenomena.

3 Chicheŵa height harmony

Of the various height-related harmony systems found in Bantu languages, the one we focus on here is perhaps the oldest and most widespread (Greenberg 1951). In it, high vowels alternate with mid under specific conditions, with *a* acting as opaque. The domain of this type of harmony is usually confined to the verbal base, composed of a root and a variable number of extensional suffixes. This domain restriction can be achieved by assigning harmony to an 'inner' layer of Bantu morphology, where extensional suffixes are attached (Level 1 of the lexicon, in Lexical-Phonological terms). Other affixes (including the final vowel suffix *-a* and all prefixes) form an 'outer' layer (Level 2), which lies outwith the scope of harmony (see Goldsmith 1985, Mtenje 1986 and Harris 1987). In this section, we illustrate this pattern of height harmony with data from Chicheŵa.

Chicheŵa, like many of its sister languages, has a canonical five-vowel system, although height harmony is also to be found in languages which preserve the original Bantu seven-term system, including Kikuyu (Clements 1991a). Historically, each Bantu verbal root contained but one vowel quality (Meeussen 1967). Some vestige of this restriction survives in Chicheŵa, where over 80 percent of roots consisting of two or more vowels conform to this pattern (figures from Scullen 1992). Most of the remainder are composed of original root-suffix complexes which are no longer

morphologically transparent. One important restriction remains fully active, however: within a span of non-low vowels, all have to be either high or mid. This constraint is evident in both static distributional and dynamic alternation patterns. The distributional evidence demonstrates that harmony functions as a condition on the vocalic composition of polysyllabic roots, as in (2). (As just indicated, the non-alternating final vowel suffix *-a* lies outside the harmonic domain of the base.)

(2)

pitiliz-a	'continue'	pelekez-a	'escort'
futuk-a	'give way'	fotokoz-a	'explain'
uzir-a	'blow cool'	kolez-a	'blow on fire'

Within roots, *a* co-occurs freely with *i* or u (3a) or with *e* or o (3b):

(3)

(a)	chiŋgamir-a luŋgam-a	'welcome' 'be straightforward'
(b)	pendam-a polam-a	'slant' 'bend face-down'

Harmony is also evidenced in alternations between high and mid vowels in extensional suffixes such as *-its-/-ets-* (causative), *-il-/-el-* (applied), *-idw-/-edw-* (passive) and *-ik-/-ek-* (descriptive passive). High suffix vowels occur after high-vowel roots, as in ?a, mid suffix vowels after mid-vowel roots, as in ?b:

(4)

		Causative	Applied	
(a)	pind-a	pind-its-a	pind-il-a	'bend'
	put-a	put-its-a	put-il-a	'provoke'
(b)	lemb-a	lemb-ets-a	lemb-el-a	'write'
	konz-a	konz-ets-a	konz-el-a	'correct'

Roots containing *a* select high-vowel suffixes:

(5)

Causative Applied

bal-a	bal-its-a	bal-il-a	'give birth'
kaŋgaz-a	kaŋgaz-its-a	kaŋgaz-il-a	'hurry up'

Suffixes containing a do not alternate and select high vowels in suffixes occurring to their right. Thus a low-vowel suffix marks the boundary between two harmonic spans: vowels to its left do not necessarily harmonise with vowels to its right. The opaque behaviour of a is illustrated by the forms in ?, in which the reciprocal suffix *-an*-intervenes between a root and at least one other suffix.

(6)

konz-an-its-a	pelekez-an-il-a	
lemb-an-its-a	kwez-ets-an-il-a	

4 Analyses with two harmonic feature values

Before proceeding to a comparison of underspecification and monovalency, it is first necessary for us to deal with a challenge facing both approaches. A potential problem for any theory predicated on the asymmetry of phonological oppositions is posed by analyses which assume that both values of a bivalent feature are harmonically active. Just this type of treatment has sometimes been proposed for Bantu height harmony. Our initial purpose is to show how these systems can be reanalysed as involving only one harmonic category.

Earlier feature-based treatments of the type of Bantu height harmony focused on here include those of Katamba (1984) for Luganda and Mtenje (1985, 1986) for Chicheŵa. Their analyses share the following assumptions: (a) all vowels are harmonically active, and (b) the harmonic sets are *e*, *o* versus peripheral *i*, *u*, *a*. Both Katamba (1984) and Mtenje (1986) characterise the peripheral set as [α high, - α low]. Harmony is achieved by spreading a [high, low] feature complex from the root onto high suffix vowels. Although the minus-alpha notation captures the high-low disjunction, it gives the wrong results in the spreading operation. While it spreads [+high, -low] from high-vowelled roots with the correct results, it also incorrectly spreads the [-high, +low] of a low root value, yielding * *lab-al-a* 'see for' in place of attested *lab-il-a*.

In his (1985) treatment of Chicheŵa, Mtenje characterises the harmonic categories as [+tense] *i*, *u*, *a* versus [-tense] *e*, *o*. Harmony consists in the spreading of both values of [tense]. A front non-low suffix vowel thus surfaces as *e* after a [-tense] (i.e. mid) root vowel, as in *lemb-el-a*. The same suffix vowel surfaces as [+tense] (i.e. high) after a [+tense] root vowel, whether this be high (as in *pind-il-a*)

or low (as in *bal-il-a*). The opacity of *a* is accounted for by assuming that it is (redundantly) specified as [+tense] at the time harmonic spreading takes place. In a form such as *lemb-an-itsa*, the rightward propagation of [-tense] from the first nucleus is blocked by the [+tense] value associated to the nucleus occupied by *a*.

While this analysis avoids the problems inherent in the minus-alpha notation, it is not without its own difficulties, particularly with regard to the use of [tense]. We believe that Mtenje's basic insight is a sound one, namely that it is desirable to be able to capture a peripherality distinction between e, o on the one hand and i, u, a on the other. The problem is that orthodox features are notoriously bad at allowing us to express this dichotomy in a direct manner. Pressing [tense] into service for this purpose can only be done at the expense of doing violence to its phonetic interpretation. True, mid vowels in Chicheŵa are typically realised as lax E/O. However, according to the usual classification and contrary to Mtenje's, *a* is also lax. In languages which have a tense-lax (or ATR) contrast in low vowels, the tense vowel typically has some form of low centralised quality (usually transcribed as &, {, %, or the like). Moreover, even setting aside the phonetic facts, invoking [tense] in order to get the harmony facts out is ad hoc to the extent that, in five-vowel systems such as Chicheŵa, this is the only process the feature is needed for. At least with orthodox features, a restrictive theory of the structure of vowel systems requires the height dimension of the canonical five-vowel type to be universally represented in terms of [high] and [low].

5 Underspecification analyses

5.0 In this section, we outline and reject an underspecification analysis of Bantu height harmony. We show how the problems of the two analyses discussed in the previous section can be overcome by abandoning the assumption that both values of the harmonic feature are active in the spreading process. Nevertheless, we will also see that this type of analysis is itself flawed in several important respects. §4.1 shows how height harmony can be characterised as the spreading of [-high], supplemented by the filling-in of redundant [+high]; §4.2 discusses the problematic treatment of opacity within this framework.

5.1 [-high] spread

According to standard underspecification assumptions, underlying representations contain only one value of each feature. Redundant values — the complement values of distinctive features and both values of any non-distinctive features — are filled in by rule. The simplest analysis of harmony under this approach is one in which the lexical value of a feature is also the one that is exclusively active in spreading. This sort of analysis is indeed available in the Chicheŵa case. One advantage of this account, it turns out, is that we avoid getting entangled in the specification of disjoint classes, the problem which dogged the analyses outlined in §4.

First, we make the assumption that *i*, *u*, *e*, *o* but not *a* are full participants in Bantu height harmony. The opacity of *a* is undeniable; but, unlike the analyses outlined in the previous section, we need not assume that it initiates its own harmonic span. In classic autosegmental terms (e.g. Clements & Sezer 1982), it is thus a non-undergoer and a blocker but not a spreader. This rules out [low] as a candidate for spreading, since it would be unable to differentiate the high and mid harmonic sets. The harmonic feature can therefore only be [high].

Deciding on which value of [high] is harmonically active is relatively straightforward. Incorrect results are obtained if we take [+high] to be the lexically specified spreading value. The low vowel of a root such as *bal-a* will correctly fail to initiate spreading, since it contains no incidence of [+high]. However, a nonlow vowel appearing in a suffix attached to such a root will erroneously be assigned the redundant [-high] value, producing for example * *bal-el-a* as opposed to grammatical *bal-il-a*.

Under one approach to underspecification, only the marked value of a given feature may be lexically specified, although the particular value may vary from one environment to another (Kiparsky 1982, 1985). Under another, the lexical value of a feature is fixed across all environments, with the particular value being a matter of language-specific selection (Archangeli 1984, 1988). According to one version of the latter approach, Combinatorial Specification, the selection of the lexical value is largely determined on the basis of its active participation in phonological processes (Archangeli & Pulleyblank 1992). At least as far as the treatment of Chicheŵa harmony is concerned, these two approaches happen to converge, in that both identify minus as the lexical value of [high] in roots. This is both the unmarked value (Archangeli 1984: 62) and the one that is harmonically active. As we will see presently, the accounts also agree with respect to their treatment of *a*: both require some kind of filter or feature-sensitive rule to capture the vowel's opaque behaviour. Having to resort to either of these devices, we will argue, crucially undermines the validity of a feature-underspecification type of analysis.

The main substantive differences between monovalency and underspecified bivalency are by definition representational. To keep these distinctions firmly in focus, it makes sense to hold any surrounding theoretical variables as constant as possible, particularly those relating to the issue of derivation. Earlier treatments of Bantu height harmony, no less than contemporary analyses of other phenomena, reflect the earlier prevalent conception of phonological processes as being characterised in terms of batteries of ordered rewrite rules located in the grammars of individual languages. As we will see presently, at least one type of feature-underpsecification analysis appears to be inextricably bound up with rule-based assumptions about phonological derivation. As such it is incompatible with more recently emerging approaches in which the mapping between lexical representation and phonological output is viewed as occurring freely in a manner that is sensitive to quite general constraints. We need not concern ourselves here with arguments about how the latter are appropriately formalised, whether as conflicting violable constraints (as in Prince & Smolensky 1993) or as parameterised conditions (as in Archangeli & Pulleyblank 1992: ch 4). In any event, at least as far as the particular phenomena involved in Bantu height harmony are concerned, the derivational questions which various constraint-based approaches address are in many respects orthogonal to the main representational issues that are at stake. To help us isolate the main differences between competing analyses of Bantu harmony, we will present them uniformly in terms of parameterised conditions which individually regulate such factors as the identity of the harmonic category and the directionality of spreading.

The underspecification account of Bantu height harmony to be sketched in this section is essentially the one discussed by Harris & Moto (1989) and proposed by Scullen (1992) for Chicheŵa and by Archangeli & Pulleyblank (1992) for Haya. Roots come in two classes: those bearing lexical [-high] and those lacking a value on the [high] tier. Extensional suffixes are unspecified for [high]. Harmony spreads an available [-high] rightwards over roots and suffixes:

(7)

Harmonic category: [-high] Direction of spreading: rightwards

The default and complement values of [high] are filled in after harmony:

(8)

Redundant [high]:

- (a) $[+low] \rightarrow [-high]$ (b)
- $[] \rightarrow [+high]$

As illustrated in ?a, a mid-vowelled root contains a lexical [-high] specification which spreads onto any nonlow suffix vowel. (In what follows, we provide truncated representations, shorn of details relating to prosodic and geometric structure that are not directly relevant to the harmonic facts. V abbreviates a nuclear position.)

(9)

(a) lembela (b) pindila [-hi] Harmony | V Nb - V | - a ? | | $\begin{bmatrix} -bk \\ -lo \end{bmatrix} \begin{bmatrix} -bk \\ -lo \end{bmatrix}$ [+hi] [+hi] | | p V Nd - V l - a | | Redundant [high] ?b $\begin{bmatrix} -bk \\ -lo \end{bmatrix} \begin{bmatrix} -bk \\ -lo \end{bmatrix}$

In a high-vowelled root such as shown in ?b, there is no harmonic spreading; here [+high] is supplied by default.

5.2 Opaque a

5.2.0 A treatment of opacity requires two components: (a) some means of ensuring that the harmonic category fails to associate with the opaque segment, and (b) a condition whereby such a failure prevents any further progress of the category across the potential harmonic domain. The latter component, as we will see in 5.2.1, is interpretable as a locality constraint. In feature-underspecification approaches, the association-failure effect in Bantu harmony is achievable through recourse to either of two devices: a feature-sensitive rule (5.2.2) or a filter (5.2.3).

5.2.1 Locality. *a* in Bantu displays that type of opacity in which a non-alternating vowel bears the same feature value as the harmonic category (at least in output) and yet fails to pass it on to a neighbouring vowel. We might think of tackling this apparent paradox by ensuring that *a* remains unspecified for [-high] until after harmony has applied. This would allow us to derive a failure to spread, as shown in ?.

balila

b V l - V l - a | |

[,,][,,]

(10)

Harmony ?	n/a
	[-hi] [+hi]
Redundant [high] ?	b V l - V l - a
	$\begin{bmatrix} +bk \\ +lo \end{bmatrix} \begin{bmatrix} -bk \\ -lo \end{bmatrix}$

In fact, this option is not available in some versions of underspecification theory, specifically those that invoke the Redundancy Rule Ordering Constraint (Archangeli 1984). The effect of the constraint is to ensure that a feature's blank values are automatically filled in before it makes an appearance in a phonological rule. The original motivation behind this proposal was to pre-empt the inadvertent use of purportedly bivalent features to conjure up ternary contrasts, specified as plus versus minus versus zero (the kind of ploy Stanley blew the whistle on in 1967). As Scullen (1992) observes, the problem in this case is that the redundant [-high] value of *a* would have to be filled in prior to the application of [-high]-spread. But this would erroneously predict that *a* should be no less harmonically active than the other nonhigh vowels, *e* and *o*.

Whatever the merits and demerits of the Redundancy Rule Ordering Constraint (and there are some who have argued that ternary power is both a necessary and a desirable consequence

of combining bivalency with underspecification (e.g. Goldsmith 1990)), there is in any case a more compelling reason for concluding that the analysis illustrated in ? falls well short of providing a full analysis of a's opacity. Having a unspecified for [-high] at the point where harmony applies fails to account for that vowel's blocking behaviour. The problem this time is to prevent a from behaving transparently, as in the following incorrect derivation, which is entirely consistent with the analysis as developed to this point.

(11)

* lembanetsa (/ lembanitsa)

Harmony ?	[-hi] V Nb - 	- V n 	- V ts	- a
	-bk -lo	+bk +lo	-bk -lo	

What an orthodox feature analysis has somehow to be able to capture here is the fact that [+low] is antagonistic to the spreading of [-high], while [-low] is not. Two basic mechanisms have been proposed for dealing with situations such as this: feature-conditioned rules and filters. Both invoke a locality requirement (formulated as the Locality Condition in Archangeli & Pulleyblank 1992), according to which a rule can apply only if an identified target is adjacent to its trigger. The effect of this condition is to place a 'no-skip' constraint on the operation of local iterative processes of the type that are active in the long-distance propagation of harmony (see van der Hulst & Smith 1986). Spreading is interrupted whenever a melody unit, for whatever reason, fails to associate with a melody-bearing unit. Gapped configurations of the following sort are thus ruled out:

(12)



The issue now becomes one of determining what it is about f_2 that resists association to $[\alpha F]$.

5.2.2 Feature-conditioned rules. Given the locality assumption, one way of achieving the blocking effect of *a* in Chicheŵa is to supplement the conditions in ? with one which requires the target and trigger of [-high]-spread to share the specification [-low]. This is illustrated in ?a.

(13)

This is an instance of what has been termed **Linked-Structure Analysis** (Cole 1987: ch 2). The device certainly achieves the desired results. A [-high] autosegment will only spread from mid to high vowels. a, being [+low], will fail to attract or spread [-high]; under the locality requirement, this interrupts the rightward propagation of harmony.

In fact, it is characteristic of this type of analysis that it does not much matter whether the relevant rule is formulated in terms of multiple linking, as in ?a, or in terms of either of the single linkings given in ?b and ?c. In the case in hand, ?b will correctly prevent *a* from spreading the harmonic feature value, while ?c will prevent *a* from receiving it. In conjunction with locality, both alternatives result in opacity.

In all essentials, the type of analysis embodied in ? amounts to an SPE context-sensitive rewrite rule, albeit dressed up in non-linear garb. As such, it suffers from one of the fundamental flaws of that format: a failure to provide a formal explanation for why a process takes place where it does. Within a theory which permits rules of this type, the prediction is made that, in principle, any harmony process is potentially conditioned by any feature. As a result, it is a matter of purest accident that the rules in ? happen to be conditional on the presence of [-low]. It would be equally possible to imagine a language in which [-high]-spread is contingent on additional linking to [+low], or for that matter to any other feature specification you might care to mention — [-round], [+continuant], [-sonorant], for example. It goes without saying that the vast majority, if not all, of these alternatives simply fail to correspond to attested patterns. In any event, the predictions that intrinsically flow from a feature-sensitive rule model fail to tally with the observation that the class of opaque segments in harmony systems is extremely restricted.¹

Before turning to an alternative underspecification account of opacity, let us briefly consider a more radical bivalent approach to the representational problems posed by height harmony. In an important contribution to the debate, Clements (1991a) discusses a range of Bantu harmony systems, including some in which traditionally distinguished categories relating to height and ATR appear to be simultaneously and perhaps inextricably entwined. One example is the southeastern Bantu pattern in which mid E/O raise to e/o under the influence of a following vowel drawn from the set i/l, u/U, e, o. This sort of pattern is a challenge to any theory, bivalent or monovalent, which treats height and ATR contrasts in terms of completely independent primes, such as [high]/[low] versus [ATR]. To do justice to the complexity of the data presented by such systems would take us well beyond our present remit. However, it is worth briefly considering how Clements' response to the challenge deals with the Chicheŵa-type pattern, since, like the analysis outlined above, it relies crucially on some form of feature-conditioned rule.

Clements' proposal is that all height and ATR contrasts be subsumed under a single bivalent feature [open]. Contrasts exceeding two degrees of openness are characterised by having multiple occurrences of this feature deployed in hierarchically organised registers, geometrically grouped under an Aperture node. (The arrangement is similar in spirit to tonal registers; see, for

¹The same brickbat can be lobbed in the direction of Harris's (1987) treatment of ATR/height harmony in southeastern Bantu, in which a rule of [-low]-spread is restricted to targets specified as [α back, α round]. Sauce for the goose, sauce for the gander.

522 John Harris

example, Clements 1983 and Hyman 1986.) The three-height inventory of a canonical five-vowel system is represented in terms of two such registers (labelled $[open_1]$ and $[open_2]$):





Under this mode of representation, height harmony of the Chicheŵa type is formalised as follows (Clements 1991a: 43):

(15)



The rightward spreading of [+open] on register 2 has the effect of lowering any following high vowel to mid. The explicit reference to $[-open_1]$ in the structural description of the rule is necessary in order to restrict the class of spreaders to mid vowels. Under this account, the failure of *a* to spread its $[+open_2]$ value is due to the fact that this vowel is specified as $[+open_1]$ on register 1, not $[-open_1]$ as required by the rule. The appearance of the latter specification in the trigger marks (15) out as a feature-conditioned rule. As such, it suffers from exactly the same flaw of arbitrariness as that which defaces the context-sensitive rules in ?.

5.3.3 Filters. The device of the context-sensitive rule has no place in a theory in which the mapping between lexical representation and phonological output is deemed to proceed freely in response to general principles and constraints. Such a theory, at least when coupled to orthodox features, would require Bantu harmony to be stated simply as **spread [-high]**. The challenge then is to identify some independent constraint which blocks the participation of *a* in this process. In feature frameworks, a by-now familiar way of expressing such constraints is in terms of filters of the sort advocated by Kiparsky (1985).

According to Kiparsky's proposal, the independence of a filter should be reflected in the fact that it not only blocks harmony in a given grammar but also helps shape the inventory of contrasts. In Guaraní, for example, the filter responsible for the opacity of voiceless segments to nasal harmony also records the fact that nasality is not distinctive in such segments (Kiparsky

1985: 130). As we will see, this independence criterion is not met in a filter-based treatment of *a* in Bantu. The filter that is required to block [-high]-spread does not correspond to any restriction on the combinability of feature values in the vowel system. Worse than that, the very combination of values it proscribes is one that is necessary for a full specification of the system.

The filter in ? could be assumed to remain in force during the operation of Bantu height harmony. (Save for certain geometric niceties which are irrelevant to the point at hand, this is in essence the device appealed to by Scullen (1992: 234).)

(16)

This would block the association of a rightward-spreading [-high] autosegment to *a*. Locality would then guarantee that [-high] could not propagate past *a*, as illustrated below:

lembanitsa

(17)

	[-hi]
Harmony ? blocked	X V Nb - V n - V ts - a
	$\begin{bmatrix} -bk \\ -lo \end{bmatrix} \begin{bmatrix} +bk \\ +lo \end{bmatrix} \begin{bmatrix} -bk \\ -lo \end{bmatrix}$

The constraint embodied in ? is essentially a negative recoding of the feature-conditioned harmony rule in ?. The same criticism of arbitrariness can thus be levelled at it. The positing of filters of this type introduces an undesirable degree of arbitrariness into phonological theory. In principle, any combination of feature values could be represented in a filter. But, as already emphasised, the class of opaque segments which are naturally attested in harmony systems is extremely restricted.

In any event, the filter in ? is suspect for at least one other reason. It will have to be switched off after harmony has applied, since the very linking it bars is later supplied by the rule in ?a, which assigns redundant [-high] to a [+low] vowel. This is evident when we complete the derivation started in ?:

(18)

lembanitsa

[-hi] [-hi] [+hi] | | |

Redundant
$$l \ V \ Nb \ - \ V \ n \ - \ V \ ts \ - a$$
[high] ? $|$ $\left[-bk \\ -lo \right]$ $\left[+bk \\ +lo \right]$ $\left[-bk \\ -lo \end{bmatrix}$

It is at best bizarre to have within the same grammar a filter and a universal default rule with precisely opposite effects.

6 A monovalent analysis

6.0 In this section, we present an analysis of Bantu height harmony which, like the underspecification account just outlined, countenances only one harmonically active category. Here, however, this effect will be seen to follow necessarily from the monovalency of the relevant phonological prime, rather than contingently from the suppression of one value of a bivalent prime. Moreover, unlike the feature-underspecification account, our analysis of opacity dispenses with contradictory filters and arbitrarily conditioned spreading rules.

6.1 briefly reviews a tried and tested version of monovalency in which vowel contrasts are represented in terms of the elements [A], [I] and [U]. 6.2 outlines an [A]-spread account of height harmony. 6.3 demonstrates how the notion of intra-segmental dependency can be invoked to derive the non-spreading behaviour of *a*. 6.4 discusses how a prosodic construal of adjacency allows us to account for the blocking behaviour of *a*.

6.1 Elements

The set of boundary conditions for a successful model of vocalic quality, it may reasonably be argued, should include some provision for capturing the unmarked status of the corner vowels a, i and u. The privileged status of these vowels is recurrently manifested in the pivotal role they play in language acquisition and in the organisation of vowel systems. This patterning should follow directly from the fundamental design properties of any viable model of vocalic quality. As has long been recognised, the orthodox feature set composed of [high], [low], [back] and [round] does not measure up particularly well to this requirement. The range of vowel-system types generated by the free combination of these features fails to provide any direct expression of the fact that certain types are more naturally preferred than others. For example, rather than the triangular arrangement constituted by a, i and u, the intersection of either of the height features with [back] predicts a basic quadrangular system (e.g. [+low, +back], [-low, +back], [+low, -back], [-low, -back]). The only way in which this framework can capture the organisational preferences of observed systems is through recourse to auxiliary markedness statements which are in and of themselves arbitrary (e.g. 'for a [+low] vowel, the unmarked value for [back] is plus)'. Any direct derivative of this model is likely to suffer from the same fundamental flaw.

The bivalency of the traditional feature system has been identified as one reason for its unsatisfactory performance on this score. One response has been to limit the contrastive potential of individual features by redefining them as monovalent, such as in the set [low], [front] and [round] or some equivalent (e.g. Donegan 1978). The potential for greater restrictiveness that this move promises, however, has sometimes been compromised by specifying certain features as bivalent and others as monovalent or by claiming that features are privative at one level of derivation but equipollent at another (e.g. Goldsmith 1990: 298 ff, 1993).

The most direct way of capturing the basic tricorn structure of vowel quality is by taking *a*, *i* and *u* themselves to be the embodiment of monovalent primes of vocalic representation, which we will symbolise here as [A], [I] and [U]. Mid vowels constitute amalgams of these primes; for example, [A] and [I] combine to produce *e*, while [A] and [U] yield *o*. In modern times, the longest established formalisation of this proposal is to be found in Dependency Phonology (Anderson & Jones 1974, Anderson & Ewen 1987). It has subsequently been taken up in Particle Phonology (Schane 1984), Government Phonology (Kaye, Lowenstamm & Vergnaud 1985) and in the work of various other researchers, including Rennison (1984), Goldsmith (1985), van der Hulst & Smith (1985) and van der Hulst (1989). Within this general framework, the relative markedness of a vowel is directly coded in the relative complexity of its representation; unmarked *a*, *i* and *u* are structurally simpler than more marked *e* and *o*.²

In a full-blooded interpretation of the tricorn arrangement, each prime enjoys stand-alone phonetic interpretability; that is, it is phonetically expressible without needing to be combined with any other primes. This immediately sets elements apart from orthodox features, each of which is only physically interpretable once harnessed to a set of other features. A [+high] segment, for example, is only manifested as such when supported by a slate of other feature values; in conjunction with, say, [-back], [+sonorant] and [-consonantal], the specification [+high] contributes to the definition of i/y. (For further discussion of this point, see Harris & Lindsey, in press.) The idea that each phonological prime has an independent embodiment but can appear in compounds with other primes suggests an analogy with physical matter. It thus seems appropriate to follow Kaye, Lowenstamm & Vergnaud (1985) in dubbing the primes **elements**.

One consequence of assuming that all phonological primes are individually interpretable is that phonological representations are phonetically mappable at all levels of derivation. That is, there is no call for redundancy rules which, as in underspecification approaches, serve the function of filling in details which are lexically absent but which are necessary to ensure that representations receive phonetic interpretation.

The combinability of [A], [I] and [U], let us assume, is constrained by the manner in which they are deployed in autosegmental structure. The basic a-i-u system is derived by taking the three elements to be arrayed on a single quality tier (cf. Rennison 1987):

(19)

²Monovalent feature systems have been proposed which fail to provide a structural representation of the unmarked status of the corner vowels. One example is Goad's (1993) characterisation of the height dimension in terms of the features [open] and [low]: *a* bears both features; *e* and *o* bear only [open]; *i* and *u* are unspecified. Under this arrangement, unmarked *a* displays more structure than marked mid vowels.

Residing on a single tier in this way, the elements are unable to combine with one another. The simple three-vowel system that this necessarily defines constitutes the initial state in language acquisition and is the structure that is retained in languages such as Djingili, Nyangumarda, Warlpiri (van der Hulst & Smith 1985) and Classical Arabic.

The development of more complex systems proceeds via what can be termed **tier division** — the autosegemental unpacking of melody. Once the representation of two elements is split into distinct tiers, the elements are in a position to fuse with one another. The next least marked type of system, as established on the basis of familiar acquisitional and distributional evidence, is the canonical five-vowel type encountered in languages such as Chicheŵa and Spanish. This suggests that the first stage of tier division results in [A] being hived off onto a separate 'aperture' line, while [I] and [U] remain as cohabitees of what is now a 'tonality' or 'colour' line. Under this arrangement, which we provisionally sketch in ?, [A] is in a position to fuse with either [I] or [U]. (The co-registration of elements under a single skeletal slot indicates fusion.)

(20)



Some additional provision must be made for ATR contrasts. Whether this should be achieved by positing an independent [ATR] element (as in Kaye, Lowenstamm & Vergnaud 1985) or by some structural means (as in van der Hulst 1990) is not a matter that need detain us here. (See Harris & Lindsey (in press) for a comparison of various element-based approaches to this issue.) In any case, in Chicheŵa, the incidence of this property (high and mid vowels are typically ATR and non-ATR respectively) is non-distinctive and plays no active role in harmony.

Still more complex systems emerge if the colour tier is allowed to divide, with the result that every element resides on its own tier and is thus free to combine with any other element. This potential is exploited in languages with front rounded vowels, which are defined by permitting the fusion of [I] with [U]: [I, U] = \ddot{u} ; [A, I, U] = \emptyset .

Phonological oppositions defined by elements are privative in the Trubetzkoyan sense. In the context of the present discussion, the main significance of this point is that only one term of each opposition is phonologically accessible. This follows from the quite reasonable assumption that phonological processes can only manipulate what is present in a representation; they cannot refer to the absence of an element. It is this design property of monovalency that gives formal expression to the observed asymmetries of harmonic and other processes (den Dikken & van der Hulst 1988). With the elements discussed so far, only three basic types of spreading system are definable. One is labial harmony, in which [U] is active. The lack of a complement value ('[\sim U]') captures the striking absence of robust cases of 'nonround' harmony (predicted to be equally natural by a bivalent framework containing [-round]). A second type is palatal harmony, in which [I] is active. The third is the one we are concerned with here: height harmony is only expressible in terms of an active [A] element.

6.2 [A]-spread

Element-based analyses of Bantu height harmony have been proposed by Rennison (1987) and Harris & Moto (1989) for Chicheŵa and by Goldsmith (1985) for the related pattern in Yaka. The basic system is derived by establishing the following conditions:

(21)

(a) Harmonic category: [A](b) Direction of spreading: rightwards

Mid-vowelled roots possess [A], while high-vowelled roots lack it. The alternating suffix vowels are also lexically [A]-less; they appear in their mid guise as a result of [A] spreading from a mid-vowelled root:

(22)

The fusion of [I] with incoming [A] defines *e*.

No harmonic activity takes place in bases containing high-vowelled roots, since these lack [A], the only active element in this system:

These forms illustrate the point that high vowels in both roots and suffixes are simply the independent embodiment of [I] or [U]. There is no need for anything resembling default fill-in rules of the type required in a feature-underspecification analysis.

6.3 Headed melodic expressions

As developed to this point, the element-based analysis erroneously predicts that lowvowelled roots should select mid-vowelled variants of the alternating suffixes. This follows from the fact that *a*, like *e* and *o*, possesses [A]:

The challenge is thus to account for the fact that [A] is harmonically active in mid vowels but not in low.

Goldsmith's (1985) solution to this problem is to formulate a harmony rule incorporating the stipulation that [A] only spreads if the nucleus to which it is attached also contains an association to some other element. The relevant part of the rule is given in ?, where X indicates the presence of some element on the [I]/[U] tier.³

(25)



This correctly distinguishes the harmonically active mid vowels, which display an association to [I] or [U], from inert a, which does not. However, the rule is subject to the same sort of criticism as that levelled at the feature-based rules given in ?: it

³The rule given here is adapted to the conventions of the present paper, without prejudice to the spirit of Goldsmith's (1985) proposal.

resorts to the expedient of allowing the spreading of some prime to be arbitrarily conditional on the presence of some other prime.⁴

Element theory provides an alternative solution which dispenses with anything resembling a context-sensitive rule and which exploits a resource that is motivated by a range of melodic facts not specifically related to harmony. Most researchers within the [A]-[I]-[U] tradition subscribe to the idea that relations between elements within a melodic compound are asymmetric. That is, one element can be assumed to predominate over others occurring within the same compound, an arrangement that is expressible in terms of head-dependent relations (Anderson & Jones 1974). This opens the way towards the recognition of **isomeric** compounds (the use of the term is Kaye, Lowenstamm & Vergnaud's (1985)) — melodic expressions containing the same elements but in different head-dependent configurations. According to one well established interpretation of this notion, the compound [A, I] defines either e or x, according to which of the two elements is head of the expression (Kaye, Lowenstamm & Vergnaud 1985, van der Hulst 1989). As shown in ?a, e is [I]-headed, while x is [A]-headed (heads underlined).⁵

(26)

(a)
$$\begin{bmatrix} A, \underline{I} \end{bmatrix}$$
 e (b) $\begin{bmatrix} A, \underline{U} \end{bmatrix}$ o
 $\begin{bmatrix} \underline{A}, I \end{bmatrix}$ \boldsymbol{x} $\begin{bmatrix} \underline{A}, U \end{bmatrix}$ $\boldsymbol{\varepsilon}$

As shown in ?b, the same type of isomeric relation distinguishes o (like e, colourheaded) from \mathcal{E} ([A]-headed). In a simplex expression, the lone element can be assumed to enjoy head status. This means that a, i and u are respectively [A]-, [I]- and [U]-headed.⁶

⁴Goad's analysis of Chicheŵa invokes a similar stipulative condition, except that it is couched in terms of the **absence** of a feature (1993: 174 ff). Under her account, the feature [open] spreads only if it is not linked to [low]. This isolates mid vowels as harmonically active, since they are exhaustively specifed as [open]. By contrast, *a* is inert, since it bears both [open] and [low].

⁵The interpretation proposed by Kaye, Lowenstamm & Vergnaud (1985), van der Hulst (1989) and others differs somewhat from that of Dependency Phonology (Anderson & Ewen 1987), in which three types of intrasegmental dependency are recognised. Applied to the front vowel space, this yields the following classification which subsumes ATR contrasts: [A] dependent on [I] = e; [A] and [I] mutually dependent = E; [I] dependent on [A] = α .

⁶The notion of dependency outlined here is quite distinct from the sense in which the term has more recently been employed in feature geometry. In the latter theory, a given node is sometimes referred to as being a dependent of some other node; for example, the Labial, Coronal and Dorsal nodes are dependents of the Place node (Sagey 1986). The main function of non-terminal geometric nodes is to

530 John Harris

Intra-segmental dependency is invoked by Harris & Moto (1989) as a means of deriving the opacity of *a* in height harmony. Specifically, the condition in ? is supplemented by a restriction which narrows the identity of the harmonic category to **dependent** [A]. This has the effect of allowing [A] to spread from mid vowels, where it has dependent status, but not from *a*, where it occurs as a head. While this proposal exploits an inherent resource of the dependency model, it is nevertheless stipulative to the extent that it provides no good reason why it is a dependent element that spreads rather than a head. In what follows, we will try to remedy this defect by building further on the notion of tier division.

By exploiting the full range of head-dependent configurations, a system with a split between an [A]-tier and an undivided [I]/[U]-tier has access to three heights, *viz. i/u, e/o, \frac{x}{a}_{\mathcal{E}}* (as in English). In order to capture the fact that this type of system is more marked than the canonical five-vowel type, we refine the mechanism of tier division by proposing that the extrication of the aperture tier from the colour tier actually takes place in two stages. The initial disengagement is only partial to the extent that aperture remains dependent on colour. This suggests a geometric configuration along the following lines:



This arrangement is somewhat reminiscent of Rennison's (1987) proposal that [A] in e and o may be indirectly linked to the skeletal tier via [I] or [U]. In keeping with the notion of tier division, however, the dependency relation in ? should be construed as holding between lines rather than between elements themselves. Contrary to the conventions adopted in feature geometry, this means that an element residing on a dependent tier will not spread parasitically on the back of some element spreading along the dominant tier.⁷

group other nodes — a relation that might be termed **dependency of occurrence** (Ewen 1993). In this case, 'there is no claim that the **content** of the features involved is in any way affected by the dependency relation: the dependent is in no sense less prominent than the "dominating" feature' (ibid: 20).

⁷Under Rennison's (1987) account, the harmonic anchors for [A]spreading are specified as [I]/[U] rather than nuclear positions. This means that indirectly linked [A] spreads in, for example, *lembela* (a),

а

Thus it is necessary to distinguish **tier dependency** from **element dependency**, the latter defining a relation of headedness between elements that allows us to identify isomeric compounds of the sort illustrated in ?. That is not to say that the two notions are completely independent of one another. A hierarchical relation between the tiers on which two distinct elements roost, let us assume, constrains the relations of headedness the elements may contract with one another. Specifically, we make the following proposal, already implicit in ?:

(28)

An element on a dominant tier is always the head of a melodic expression.

To put it somewhat differently, an element on a dependent tier can only head a melodic expression in the absence of an element on the dominant tier. The impact of this principle on a system with the partial colour-aperture split portrayed in ? is to bar a head [A] from supporting a dependent [I] or [U]. In this way, we ensure that the basic five-vowel system is one in which the compound vowels are mid *e* and *o* rather than low æ and æ. This proposal, as we will see presently, also derives the opaque behaviour of *a* under [A] spreading.

The next stage in the disengagement of the [A]- and [I]/[U]-tiers is total schism. In the absence of inter-tier dependency, the principle in ? is inapplicable, and pairs of elements on separate lines are free to combine in any relation of headedness. This results in a seven-vowel system of the following type:

(29)



while directly linked [A] remains inert in a form such as *balila* (b):

(a)	[A] 	(b) [A]
	[ʊ] [I] 	[I]
	l v Nb - v l - a	b v l - v l -

This approach appears to make the erroneous prediction that [A] will parasitically spread in systems with [I] and/or [U] harmony. The inevitable effect on an [I]-harmony system, for example, would be to cause a non-high front vowel not only to palatalise but also to lower an alternating high vowel. Turkish is just one example of a palatal-harmony language where no such implication holds; hence attested *dere-sl* \rightarrow *dere-si* 'river (poss.)', as opposed to * *dere-se*.

532 John Harris

-----[A]----[<u>A</u>]----[<u>A</u>]-----[A]------[A]------

In short, what distinguishes the canonical five-vowel inventory in ? from the more marked system in ? is the inability of head [A] to fuse with a dependent on the [I]/[U] tier.

The relevance of intra-segmental dependency to the analysis of Bantu height harmony is that it provides an independent means of distinguishing the harmonic behaviour of e and o from the non-harmonic behaviour of a. The key to the distinction lies in the observation noted above that [A] only spreads when it is a dependent, its status in mid vowels. This is illustrated in ?, repeated from ? but now amended to show the tier and element dependencies appropriate to the canonical five-vowel system:

(30)	
k V Nz - V l - a	l V Nb - V l - a
- [<u>U]</u> [<u>I</u>] -	- [<u>I</u>] [<u>I</u>] -
\	\
[A]	[A]
konzela	lembela

In contrast, when [A] occurs as a head, the status it enjoys in *a*, it fails to spread:

```
balila
```

In order to capture this distinct behaviour, it is not necessary to complicate the statement of harmony by adding a stipulation to the effect that [A] is active only as a dependent, as proposed in Harris & Moto (1989). The result we seek is already supplied by the manner in which the five-vowel system is constructed, provided we make the reasonable assumption that, in accordance with Structure Preservation, lexically established dependency relations remain stable under spreading. Thus, a spreading element, through multiple association, cannot change its spots from head to dependent or vice versa. This prevents a head [A] from associating as a dependent to the following vowel in forms such as ?. Nor can the lexical head of a melodic

expression be usurped from its position of pre-eminence by the arrival of a spreading element. This would be the undesired result, were a head [A] allowed to spread and impose its headship on a following [I] or [U]. The relegation to dependent status that this would mean for elements on the dominant colour tier in any event violates the principle in ?:

(32)

According to this account then, the non-spreading behaviour of *a* in Bantu height harmony follows from an independent property of headed monovalent representations. However, lest a suspicion might linger that intra-segmental dependency is some *deus ex machina*, improvised solely to bolster this particular analysis, let us briefly summarise a number of additional arguments which can be mounted in its support.

One argument makes straightforward appeal to economy. The range of possible vocalic contrasts attested in languages of course exceeds that generated by the simple amalgamations of [A], [I] and [U] shown in ?. Exploiting the structural concept of headedness, which is in any event independently necessary for the representation of prosodic relations, allows us to approach the full definition of this range without adding to the set of basic primes.

Another argument relates to the phonetic interpretation of elements in compounds. In terms of their spectral mappings in the speech signal, there is a clear sense in which e is more *i*-like and less *a*-like than x. This is consistent with the claim that [I] predominates in e, whereas [A] predominates in x. The same point can be made in relation to o versus a. (On this matter, see Harris & Lindsey 1991.)

Furthermore, there is evidence that the headedness of melodic expressions is actively exploited by certain phonological processes other than those treated in terms of spreading. One class of process involves a type of harmonic phenomenon which, it has been proposed, is appropriately analysed in terms of head agreement (see, for example, Lowenstamm & Prunet's (1988) treatment of Tigre). This consists in a requirement that the head elements of all vowels within a particular domain must occupy the same tier.

6.4 Locality and opacity

Intra-segmental dependency successfully captures the non-spreading property of a in Bantu height harmony. The question now is whether this analysis also derives that vowel's blocking behaviour. The immediate answer would seem to be yes, provided we make two assumptions: (a) some kind of locality requirement is in force, and (b) head and dependent occurrences of an element are distinct representational objects. These notions are illustrated in the following representation of *kwez-ets-an-il-a*, where a dependent [A], spreading from the first nucleus V₁, encounters a head [A] in V₃:

kwezetsanila

In §6.1, we invoked the notion that different elements are unable to fuse if they are deployed on the same tier. If we construe head and dependent instances of the same element in exactly the same terms, we may conclude that, in ?, dependent [A] is unable to OCP-merge with head [A]. Locality then blocks the dependent [A] from spreading past V_3 to V_4 .

The validity of this account of a's blocking behaviour depends to a great extent on how we understand adjacency between melody units to be defined in phonological representations. One extremely widespread view on this matter relies on a notion which might be termed **melodic locality**:

(34)

Melodic locality

Two autosegments are adjacent if

- (a) they reside on the same tier, and
- (b) no other autosegment intervenes.

Melodic locality is what is appealed to in, for example, Archangeli & Pulleyblank's (1986) notion that certain types of rule opt for **minimal scansion**; that is, their scope is confined to the tier on which a particular feature or geometric node is located.

A melodic construal of locality lies behind much recent work in feature geometry. It is, for example, implicit in the proposal that the same place features may

appear twice in a representation under separate C(onsonant)-place and V(owel)-place nodes (or some notational equivalent) (Clements 1991b, Odden 1991, Clements & Hume, forthcoming). The main aim of this segregation is to capture the rarity of two types of phenomena: (a) feature-spreading between consonants across an intervening vowel, and (b) the blocking of feature-spreading between vowels by an intervening consonant. The first result is achieved by assuming that the relevant features appear under the C-node in both consonants and vowels. In this way, so the argument goes, the absence of a harmonic interaction between consonants flanking a vowel reflects the fact that their C-nodes are not adjacent, separated as they are by the C-node of the intervening vowel. The second result is derived by stipulating that vowels, but not consonants, also bear specifications for the relevant features under the V-node; the non-participation of consonants in vowel harmony is then accounted for by assuming that spreading in the relevant cases exclusively picks out V-nodes.

There are various reasons for judging this to be an unconvincing representation of adjacency. For one thing, it entails a wholesale duplication of feature specifications under C- and V-nodes. One logical conclusion of the approach is that vowels must be C-specified for any feature which fails to spread between flanking consonants. In practice, this means virtually all features relevant to consonants, not just those required for place. The fact is that manner and source features such as [continuant], [sonorant], [consonantal], [slack vocal cords] and [stiff vocal cords] are rarely if ever observed to spread between non-contiguous consonants. Are we then to attribute this to the appearance of specifications for these features in intervening vowels?

As is now well known, the lay terms consonant and vowel, which the geometric distinction between C- and V-nodes calls to mind, are ambiguous. On the one hand, vowel-consonant can refer to the phonetically definable distinction between (constricted) contoids and (non-constricted) vocoids. On the other, it refers to different positions in syllabic constituent structure — roughly speaking, nuclear (vowel) versus non-nuclear (consonant). The lack of isomorphism between these two dimensions is confirmed by the existence of non-nuclear vocoids (glides) and nuclear contoids (such as syllabic nasals). In recognition of this fact, each of the dimensions merited its own feature in SPE — [consonantal] for vocoid versus contoid, [syllabic] for nuclear versus non-nuclear. It is the latter dimension that the distinction between C- and V-nodes is evidently designed to capture. In this respect, the CV-node arrangement harks back to pre-autosegmental days, since it is a melodic coding of what is now generally acknowledged to be a prosodic matter. It was formal recognition of this point that initially prompted the decision to replace the earlier CVslots of the timing tier by skeletal (and later moraic) positions devoid of melodic content (Kaye & Lowenstamm 1984, Levin 1985). The representation of syllabic

organisation in terms of prosodic constituent structure renders the feature [syllabic] wholly superfluous.

The reversion to a melodic coding of syllabic affiliation that CV-node geometry entails arises out of an acknowledgment that certain positions can enter into relations of locality even though they may not be string-adjacent. But such relations must in any event be independently represented in the prosodic hierarchy. This is perhaps most familiarly true of metrical structure: stress-bearing syllable heads can be adjacent at higher levels of prosodic projection, including the foot and the word, where any non-heads (onsets, codas) that might intervene at some lower level are simply invisible. This insight carries over to phenomena other than stress, such as tone. Interactions between nuclear projections have also been shown to be at work in the conditioning of vowel syncope/epenthesis (Kaye 1990, Charette 1991).

Significantly from the standpoint of the present discussion, locality at higher levels of nuclear projection must also be invoked in the treatment of harmony (Halle & Vergnaud 1981, Kaye, Lowenstamm & Vergnaud 1985). Short-distance assimilations, such as those responsible for the homorganicity of certain consonant clusters, invoke string adjacency, defined at the level of the skeletal tier, containing the terminal nodes of the prosodic hierarchy. Long-distance assimilations of the type involved in vowel harmony respond to relations holding between nuclear projections at higher levels of the hierarchy, such as the word.

In light of these considerations, we are entitled to ask what value there is to be had in melodically duplicating prosodic information in the form of CV geometric nodes. CV geometry, it seems to us, leaves us in a position of not being able to see the prosodic wood for the melodic trees.

Under a prosodic construal of locality, occupation of the same tier is a necessary but not a sufficient condition for establishing a relationship between two autosegments. In addition, as expressed in ?b, they must be prosodically adjacent.

(35)

Prosodic locality

Two autosegments are adjacent if

- (a) they reside on the same tier, and
- (b) the positions to which they are attached are adjacent on some projection.

This interpretation subsumes Archangeli & Pulleyblank's (1986) notion of **maximal** scansion, the operation whereby a process targeting a particular feature or node α scans the highest level of syllabic structure which provides access to α .

Archangeli & Pulleyblank's view of locality in rule application allows for two types of rule, those that scan melodically and those that scan prosodically. We cannot

embrace this all-encompassing approach without first rejecting the more restrictive alternative, under which relations of locality are exclusively defined with reference to prosodic structure, as in ?. In the absence of proof that the latter model is empirically underpowered, it is the one we continue to assume here.

A consistently prosodic interpretation of locality is only possible within a theory which subscribes to the Strict Layer Hypothesis, according to which each constituent on any given level of the prosodic hierarchy must be properly included in the next higher level (Selkirk 1984). This requirement is not met by the version of moraic theory in which onset melody units are deemed to link directly to syllable nodes, bypassing the moraic level (as in Hayes 1989). Adherence to this view of representation entails the recognition of rules which simultaneously scan melody, in the form of root nodes, and prosody, in the form of morae (Archangeli & Pulleyblank 1992: 314-316).⁸ A strictly layered representation, on the other hand, is one in which every melody unit is directly anchored to the skeletal tier (whether this take the form of syllabic positions à la Kaye & Lowenstamm (1984) or morae to which onset segments are adjoined, as in Zec 1989). With this mode of representation, it is possible to maintain the position that the lowest level of scansion in phonological processing is no less consistently defined in prosodic terms (involving string adjacency on the skeletal tier) than higher levels.

On higher levels of projection, where harmony takes place, spreading may be thought of as instantiating a relation in which a head nuclear position licenses all other nuclear positions within its domain (Halle & Vergnaud 1981, Kaye, Lowenstamm & Vergnaud 1985, Demirdache 1988). Specifically, some element that is lexically lodged in the head position spreads to other positions within the same domain; or, to put it somewhat differently, a licensed position is bound to its licensor by the spreading element. In the Chicheŵa height case, the rightward directionality of harmony, spelt out in the specific conditions given in ?b and ?b, reflects the fact that, in this language, licensing relations between nuclei within the verbal base are headinitial.

The validity of the blocking analysis illustrated in ? rests on the assumption that the protagonists in harmony, be they spreaders (V_1) , undergoers (V_2) or blockers (V_3) , are adjacent in the sense of ?. Inherent in this assumption is the prediction that, at a given level of projection, the appearance of adjacent elements on the same tier will give rise to opacity if they are of unlike head-dependent status. (This includes the sort

⁸This hybrid scansion potential is just as much a feature of constraint-based theories which assume the same mode of representation. This is perhaps most obviously true in the case of the ALIGN family of constraints, particularly those which require the edge of some prosodic domain to coincide with the edge of some morphological domain (Prince & Smolensky 1993, McCarthy & Prince 1994). In many instances, violations of such constraints are calculated simultaneously in terms of root nodes and morae.

of case set out in ?, where a dependent is unable to spread past a head occurrence of the same element; but it presumably extends to the converse situation, in which a head is predicted to be unable to spread past a dependent.)

On the face of it, this prediction is wrong. According to Demirdache (1988), in at least two harmony systems, Finnish and Hungarian, spreading by a dependent element, in this case [I], is **not** blocked by a head appearing within the same harmonic span. This is demonstrated by the transparent behaviour of high and mid palatal vowels, both [I]-headed, in the two languages. Thus in the hoary Finnish example *værtinæ* 'spinning wheel', the dependent [I] emanating from the first æ spreads to the second æ without being impeded by the head [I] in the intervening *i*. Demirdache seeks to elevate this transparency effect to a universal by attributing it to a general property of dependent elements mentioned in §6.3, namely the recessive contribution they make to the phonetic interpretation of compound expressions. On her account, a dependent [I] does fuse with and can thus spread past *i/e* in systems such as Finnish and Hungarian. However, adding dependent [I] to *i* or *e* makes no difference to their phonetic interpretation, she claims, since both already contain [I] as head. Under her proposal, the complementary type of system, in which the harmonically active element spreads as a head, is predicted to lack transparency. In its capacity as head, an element is able to stamp its authority on the phonetic complexion of an expression and will thus force harmonic agreement on any vowel that so much as crosses its path. This is precisely the situation encountered in the [I]-harmony system of Turkish. In this case, all vowels within a palatal harmonic span are [I]-headed, resulting for example in the selection of e-vowelled alternants of the suffixes -lar/-ler (plural) and -dan/-den (genitive) (e.g. dil-ler-in-den 'language (pl. poss. gen.)').

According to Demirdache, the purported link between dependent status and transparency extends to all types of harmony. She points to the Montañes dialect of Pasiego Spanish as exemplifying systems in which this principle holds of [A]-spreading, although she does not offer a detailed analysis. The Pasiego example is not particularly compelling, since there are good grounds for surmising that it is not a classic case of harmonic spreading at all.⁹ In any event, her claim is flatly contradicted

⁹The general outline of height harmony in Pasiego, the reader might recall, is as follows: if the tonic vowel of a word is high, then any non-low vowel to its left must also be high (Penny 1969). This gives rise to alternations such as *koxeré* 'take (fut. 1 sg.)' versus *kuxirí:s* (fut. 2 pl.). Low vowels, both *a* and its non-peripheral counterpart &, do indeed behave transparently, as Demirdache (1988) suggests. This is confirmed by forms such as *legatérna* 'lizard' and *lsk&l&mbrÚxU* 'dog-rose', in which the first vowel of each word harmonises with the tonic across one or more occurrences of a/&. As with the Bantu system reviewed here, earlier feature-based analysis of the Pasiego pattern incorporated the symmetrical spreading of two feature-values (in this case [+high] and [-high] — see McCarthy 1984). Vago (1988) has subsequently shown that only one active feature-value need be invoked, [+high], with [-high] being filled in by default. This is consistent with the informal obervation that Pasiego height harmony, unlike

by the Bantu pattern reviewed here: in systems of this type, there is no escaping the conclusion that a head [A] placed in the path of a spreading dependent [A] blocks harmony.

7 Concluding remark on opacity versus transparency

We leave it to a wider survey of harmony types to determine whether a solid connection can be established between intra-segmental dependency and harmonically transparent vowels. However, on the strength of the contradictory findings just outlined, the initial results do not look particularly promising. This, considered in conjunction with the unsatisfactory nature of CV-geometric attempts to capture restrictions on spreading, suggests that we will search in vain for an exclusively melodic account of transparency versus opacity.

On the other hand, the prosodic hierarchy already makes structural provision for just such a distinction. We conclude by briefly indicating how this resource might be exploited. This will involve us summarising the conditions we have been assuming

Bantu, is of the raising type.

From an element-based perspective, the Pasiego height system is a variant of a widespread pattern of vowel reduction, in which mid vowels are banished from unstressed positions — the sort of phenomenon found, for example, in Bulgarian, Catalan and some types of Portuguese. The phenomenon is straightforwardly treated in element terms as a diminution in the melodic complexity of vocalic expressions in prosodically weak positions; mid vowels, compounds of [A] and [I] or [U], simplify to high through the suppression of [A]. (The gist of this analysis, more fully developed in Harris 1990a, is provided in Harris & Lindsey, in press.) The twist to the Pasiego pattern is that the reduction is harmonically adjusted; that is, a lexically mid vowel in a weak nucleus reduces to high, unless its [A] is sanctioned by an [A] in the dominant nucleus. The latter situation is illustrated in (a) (where the arrows indicate the directionality of inter-nuclear licensing).



In (b), in contrast, there is no [A] in the licensing nucleus to sustain the lexically present [A]s in the licensed positions, and delinking (indicated by \approx) ensues. Each of the residual elements independently defines a high vowel. This part of the process manifests a general type of complexity agreement, whereby the elemental complexity of a licensed position cannot exceed that of its licensor (the Complexity Condition discussed in Harris 1990b).

to be at play in opacity. In the case of transparency, our remarks can be no more than suggestive, since this phenomenon does not figure on the agenda we have set ourselves in this article.

Compare the configurations in ?, which portray nuclear positions at different levels of projection. In each of ?a and ?b, position x_1 , as the ultimate head of the representation, licenses positions x_2 and x_3 (indicated by $\rightarrow \rightarrow$). [*E*] and [ε] stand for elements on the same tier which differ in terms of whether or not they occur as the head of a melodic expression.

(36)



Consider a situation in which the [*E*] lexically specified in x_1 is harmonically active and potentially targets x_3 . The issue is whether or not position x_2 , by virtue of its attachment to [*e*], interferes with the spreading of [*E*].

The configuration in ?a depicts opacity in the manner we have been assuming for Chicheŵa. Here harmony is defined with reference to some level of the prosodic hierarchy (Nⁿ) at which all nuclei are projected. Since x_1 and x_2 are adjacent at this level, and since [*E*] is unable to merge with [ε], the locality requirement prevents [*E*] from spreading to x_3 .

?b indicates how transparency might be prosodically characterised. In this configuration, two levels of projection are relevant to the consideration of harmony. Within one domain, Nⁿ, x_1 licenses x_2 . At a higher level, Nⁿ⁺¹, x_2 is not projected. Since x_1 and x_3 are adjacent within this domain, they can be harmonically bound with respect to [*E*], irrespective of the melodic content of x_2 , which is invisible at this level.

This arrangement simply recapitulates mechanisms that are familiar from metrical theorising. We may conceive of a transparent nucleus as failing to be projected at some level in the same way that, say, the weak nucleus of a stress foot fails to be projected at the word level. If this parallel is valid, then it is not surprising that a transparent nucleus should display the same degree of reduced contrastive potential as is characteristically found in unstressed nuclei.¹⁰ Both cases can be related to a general property of prosodically recessive positions — a diminished ability to support melodic content.

References

- Anderson, John M. & Colin J. Ewen (1987). *Principles of Dependency Phonology*. Cambridge: Cambridge University Press.
- Anderson, John M. & Charles Jones (1974). Three theses concerning phonological representations. *Journal of Linguistics* 10. 1-26.
- Archangeli, Diana (1984). Underspecification in Yawelmani phonology and morphology. PhD dissertation, MIT. Published 1988, New York: Garland.
- Archangeli, Diana (1988). Aspects of Underspecification Theory. *Phonology* 5. 183-205.
- Archangeli, Diana & Douglas Pulleyblank (1986). The content and structure of phonological representations. Ms. University of Arizona/University of Southern California.
- Archangeli, Diana & Douglas Pulleyblank (1992). Grounded phonology. Ms. Tuscon, AZ: University of Arizona/Vancouver, BC: University of British Columbia.
- Bogers, K., H. van der Hulst & M. Mous (eds) (1986). *The representation of suprasegmentals in African languages*. Dordrecht: Foris.
- Calabrese, Andrea (1989). Phonological variations. In Paola Benincà (ed.), *Dialect variation and the theory of grammar*, 9-39. Dordrecht: Foris.
- Clements, George N. (1983). The hierarchical representation of tone features. In I.R. Dihoff (ed.), *Current approaches to African linguistics*, vol. 1, 145-176. Dordrecht: Foris.
- Clements, George N. (1985). The geometry of phonological features. *Phonology* 2. 223-50.
- Clements, George N. (1991a). Vowel height assimilation in Bantu languages. Working Papers of the Cornell Phonetics Laboratory 5. 37-76.
- Clements, George N. (1991b). Place of articulation in consonants and vowels: a unified theory. *Working Papers of the Cornell Phonetics Laboratory* 5.77-123.

¹⁰It would be interesting to discover whether stress-feet necessarily coincide with 'harmonic feet' in languages possessing both phenomena (such as Finnish). Although we suspect this not to be so, this is a matter for further research.

- Clements, George N. & Elizabeth V. Hume (forthcoming). Segment structure. To appear in John Goldsmith (ed.), *A handbook of phonology*. Oxford: Blackwell.
- Clements, George N. & Engin Sezer (1982). Vowel and consonant disharmony in Turkish. In van der Hulst & Norval Smith (1982) Part II, 213-256.
- Cole, Jennifer (1987). Planar phonology and morphology. PhD dissertation, MIT.
- Demirdache, Hamida (1988). Transparent vowels. In van der Hulst & Smith 1988 Part II, 39-76.
- den Dikken, Marcel & Harry van der Hulst. (1988). Segmental hierarchitecture. In van der Hulst & Smith 1988 Part I, 1-78.
- Donegan (Miller), P.J. (1978). On the natural phonology of vowels. PhD dissertation, Ohio State University.
- Durand, Jacques & Francis Katamba (eds) (in press). *Frontiers in phonology: atoms, structures and derivations*. Harlow, Essex: Longman.
- Ewen, Colin J. (1993). Dependency relations in phonology. Ms. Rijksuniversiteit te Leiden. To appear in John Goldsmith (ed.), *A handbook of phonology*. Oxford: Blackwell.
- Goad, Heather (1993). On the configuration of height features. PhD dissertation, University of Southern California.
- Goldsmith, John A. (1985). Vowel harmony in Khalkha Mongolian, Yaka, Finnish and Hungarian. *Phonology* 2. 251-274.
- Goldsmith, John A. (1990). Autosegmental and metrical phonology. Oxford: Blackwell.
- Goldsmith, John A. (1993). Harmonic Phonology. In J.A. Goldsmith (ed.), *The last phonological rule*, 21-60. Chicago, IL: University of Chicago Press.
- Greenberg, Joseph (1951). Vowel and nasal harmony in the Bantu languages. *Zaïre* 5.8.
- Halle, Morris & Jean-Roger Vergnaud (1981). Harmony processes. In W. Klein & W. Levelt (eds), Crossing the boundaries in linguistics, 1-23. Dordrecht: Reidel.
- Hayes, Bruce (1989). Compensatory lengthening in moraic phonology. *Linguistic Inquiry* 20. 253-306.
- Harris, John (1987). Non-structure-preserving rules in lexical phonology: southeastern Bantu harmony. *Lingua* 73. 255-292.
- Harris, John (1990a). Reduction harmony. Paper read at the GLOW Phonology Workshop, University of London.
- Harris, John (1990b). Segmental complexity and phonological government. *Phonology* 7. 255-300.
- Harris, John & Geoff Lindsey (1991). Segmental decomposition and the signal. Paper read at the London Phonology Seminar, University of London. To appear in

Phonologica 1992: proceedings of the 7th International Phonology Meeting, Krems, 1992. Turin: Rosenberg & Sellier.

- Harris, John & Geoff Lindsey (in press). The elements of phonological representation. In Durand & Katamba, in press.
- Harris, John & Francis Moto (1989). Bantu height harmony. Paper read at the 20th Annual Conference on African Linguistics, University of Illinois, Urbana.
- Hulst, Harry van der (1989). Atoms of segmental structure: components, gestures and dependency. *Phonology* 6. 253-84.
- Hulst, Harry van der (1990). The segmental spine and the non-existence of [±ATR]. In Joan Mascaró & Marina Nespor (eds), *Grammar in progress: GLOW essays* for Henk van Riemsdijk, 247-257. Dordrecht: Foris
- Hulst, Harry van der & Norval Smith (eds) (1982). *The structure of phonological representations, Parts I & II*, 131-175. Dordrecht: Foris.
- Hulst, Harry van der & Norval Smith. (1985). Vowel features and umlaut in Djingili, Nyangumarda and Warlpiri. *Phonology* 2. 275-302.
- Hulst, Harry van der & Norval Smith (1986). On neutral vowels. In Bogers, van der Hulst & Mous 1986, 233-279.
- Hulst, Harry van der & Norval Smith (eds) (1988). *Features, segmental structure and harmony processes, Parts I & II.* Dordrecht: Foris.
- Hyman, Larry M. (1986). The representation of multiple tone heights. In Bogers, van der Hulst & Mous 1986, 109-152.
- Itô, Junko, Armin Mester & Jaye Padgett (1994). Licensing and redundancy: underspecification in Optimality Theory. Ms. University of California at Santa Cruz.
- Katamba, Francis. (1984). A nonlinear analysis of vowel harmony in Luganda. Journal of Linguistics 20. 257-275
- Kaye, Jonathan (1990). Government in phonology: the case of Moroccan Arabic. *The Linguistic Review* 7. 131-159.
- Kaye, Jonathan & Jean Lowenstamm (1984). De la syllabicité. In François Dell, Daniel Hirst & Jean-Roger Vergnaud (eds.), *Forme sonore du language*, 123-159. Paris: Hermann.
- Kaye, Jonathan, Jean Lowenstamm & Jean-Roger Vergnaud. (1985). The internal structure of phonological elements: a theory of charm and government. *Phonology Yearbook* 2. 305-328.
- Kaye, Jonathan, Jean Lowenstamm & Jean-Roger Vergnaud (1990). Constituent structure and government in phonology. *Phonology* 7. 193-232.
- Kiparsky, Paul (1982). From Cyclic Phonology to Lexical Phonology. In van der Hulst & Smith 1982, 131-175.

- Kiparsky, Paul (1985). Some consequences of Lexical Phonology. *Phonology* 2. 85-138.
- Levin, Juliette (1985). A metrical theory of syllabicity. PhD dissertation, MIT.
- Lindsey, Geoff & John Harris (1990). Phonetic interpretation in generative grammar. *UCL Working Papers in Linguistics* 2. 355-369.
- Lowenstamm, Jean & Jean-François Prunet (1988). Tigre vowel harmonies. Rapport Annuel du Groupe de Recherche sur la Linguistique Africaniste au CRSH 1987-88. Université du Québec à Montréal.
- McCarthy, John J. (1984). Theoretical consequences of Montañes vowel harmony. *Linguistic Inquiry* 15. 291-318.
- McCarthy, John J. (1988). Feature geometry and dependency: a review. *Phonetica* 45. 84-108.
- McCarthy, John J. & Alan Prince (1994). Generalized alignment. *Morphology Yearbook* 1.
- Meeussen, A.E. (1967). Bantu grammatical reconstructions. *Africana Linguistica* 3. 79-121.
- Mtenje, Al D. (1985). Arguments for an autosegmental analysis of Chicheŵa vowel harmony. *Lingua* 66. 21-52.
- Mtenje, Al D. (1986). Remarks on extra-harmonicality in Bantu harmony systems. Sheffield Working Papers in Language and Linguistics 3. 62-76.
- Odden, David (1991). Vowel geometry. Phonology 8. 261-290.
- Penny, Ralph J. (1969). Vowel-harmony in the speech of the Montes de Pas (Santander). *Orbis* 18. 148-166.
- Prince, Alan & Paul Smolensky (1993). Optimality Theory: constraint interaction in generative grammar. *Technical Report # 2 of the Rutgers Center for Cognitive Science*.
- Pulleyblank, Douglas (1986). Tone in Lexical Phonology. Dordrecht: Reidel.
- Pulleyblank, Douglas (in press). Feature geometry and underspecification. To appear in Durand & Katamba, in press.
- Rennison, John (1984). On tridirectional feature systems for vowels. *Wiener linguistische Gazette*. 33-34, 69-93. Reprinted in Jacques Durand (ed.) (1986), *Dependency and non-linear phonology*, 281-303. London: Croom Helm.
- Rennison, John (1987). Vowel harmony and tridirectional vowel features. *Folia Linguistica* 21. 337-354.
- Sagey, Elizabeth (1986). The representation of features and relations in non-linear phonology. PhD dissertation, MIT.
- Sagey, Elizabeth (1988). On the ill-formedness of crossing association lines. *Linguistic Inquiry* 19. 109-118.

- Schane, Sanford S. (1984). The fundamentals of Particle Phonology. *Phonology* 1. 129-156.
- Scullen, Mary Ellen (1992). Chicheŵa vowel harmony and Underspecification Theory. *Linguistic Analysis* 22. 218-245.
- Selkirk, Elisabeth O. (1984). *Phonology and syntax: the relation between sound and structure*. Cambridge, MA: MIT Press.
- Stanley, Richard (1967). Redundancy rules in phonology. Language 43. 393-436.
- Vago, Robert M. (1988). Underspecification in the height harmony system of Pasiego. *Phonology* 5. 343-362.
- Zec, Draga (1989). Sonority constraints on prosodic structure. PhD dissertation, Stanford University.