This file is the scanned faksimile of the following reference:
Harris, John & Geoff Lindsey 1995. The elements of phonological representation. Frontiers of Phonology, edited by J.Durand & F.Katamba, 34-79. Harlow, Essex: Longman.

Chapter 2

The elements of phonological representation

John Harris and Geoff Lindsey

1 Introduction[1]

What size are the primes of which phonological segments are composed? From the standpoint of orthodox feature theory, the answer is that each prime is small enough to fit inside a segment, and not big enough to be phonetically realized without support from other primes. Thus, [+high], for example, is only realizable when combined with values of various other features, including for instance [−back, −round, −consonantal, +sonorant] (in which case it contributes to the definition of a palatal approximant). This view retains from earlier phoneme theory the assumption that the segment is the smallest representational unit capable of independent phonetic interpretation.

In this chapter, we discuss a fundamentally different conception of segmental content, one that views primes as small enough to fit inside segments, yet still big enough to remain independently interpretable. The idea is thus that the subsegmental status of a phonological prime does not necessarily preclude it from enjoying stand-alone phonetic interpretability. It is perfectly possible to conceive of primes as having autonomous phonetic identities which they can display without requiring support from other primes. This approach implies recognition of 'primitive' segments, each of which contains but one prime and thus reveals that prime's autonomous phonetic signature. Segments which are non-primitive in this sense then represent compounds of such primes.

This notion – call it the *autonomous interpretation* hypothesis – lies at the heart of the traditional notion that mid vowels may be considered amalgamations of high and low vowels. Adaptations of this idea are to be found in the work of, among others, Anderson

and Jones (1974, 1977) and Donegan (1978). Anderson and Jones' specific proposal is that the canonical five-vowel system should be treated in terms of various combinations of three primes which we label here [A], [I] and [U]. Individually, these manifest themselves as the primitive vowels *a*, *i* and *u* respectively (see (1a)). As shown in (1b), mid vowels are derived by compounding [A] with [I] or [U].

(1)

(a)			(b)		
	[A]	a		[A, I]	e
	[I]	i		[A, U]	o
	[U]	u			

Not all work which incorporates this type of analysis has maintained the notion of autonomous interpretation. In current Dependency Phonology, the direct descendant of Anderson and Jones' (1974) proposal, the view is in fact abandoned in favour of one in which primes such as those in (1) ('components') define only the resonance characteristics of a segment and must be supplemented by primes of a different sort ('gestures') which specify manner and major-class properties (see Anderson and Durand 1987, Anderson and Ewen 1987 and the references therein). A similar line has been followed in van der Hulst's 'extended' dependency approach (1989, this volume).

Nevertheless, the principle of autonomous interpretation continues to figure with varying degrees of explicitness in other approaches which employ the primes in (1). This is true, for example, of Particle Phonology (Schane 1984a) and Government Phonology (Kaye et al. 1985, 1990), as well as of the work of, among others, Rennison (1984, 1990), Goldsmith (1985) and van der Hulst and Smith (1985). Although the use of the term *element* to describe such primes is perhaps most usually associated with Government Phonology, we will take the liberty of applying it generically to any conception of subsegmental content which incorporates the notion of autonomous interpretation. The analogy with physical matter seems apt in view of the idea that phonological elements may occur singly or in compounds with other elements. Moreover, to the best of our knowledge, Government Phonology is the only framework to have explicitly pushed the autonomous interpretation hypothesis to its logical conclusion – namely that each and every subsegmental prime, not just those involved in the representation of vocalic contrasts, is independently interpretable.

In this chapter, we explore some of the consequences of adopting

a full-blooded version of element theory. This entails a view of phonological derivation which is radically different from that associated with other current frameworks, particularly those employing feature underspecification. One significant implication of the autonomous interpretation hypothesis is that phonological representations are characterized by full phonetic interpretability at all levels of derivation. This in turn implies that there is no level of systematic phonetic representation. Viewed in these terms, the phonological component is not a device for converting abstract lexical representations into ever more physical phonetic representations, as assumed say in underspecification theory. Instead, it has a purely generative function, in the technical sense that it defines the grammaticality of phonological structures. In the following pages, we will argue that these consequences of the autonomous interpretation hypothesis inform a view of phonology and phonetic interpretability that is fully congruent with current thinking on the modular structure of generative grammar.

The various element-based and related approaches share a further theoretical trait which distinguishes them from traditional feature theory: they all subscribe to the assumption that phonological oppositions are, in Trubetzkoyan terms, *privative*. That is, elements are single-valued (monovalent) objects which are either present in a segment or absent from it. Traditional feature theory, in contrast, is based on the notion of *equipollence*; that is, the terms of an opposition are assigned opposite and in principle equally weighted values of a bivalent feature.

It is probably fair to say that, of the various facets of element-based theory, monovalency is the one that has attracted the closest critical attention. The empirical differences between privative and equipollent formats continue to be disputed, and the debate is hardly likely to be resolved in a short chapter such as this. Since the arguments have been well aired elsewhere (see especially van der Hulst 1989, den Dikken and van der Hulst 1988), we will do no more summarize the main issues (Section 2). Our intention here is to shift the focus of the comparison between features and elements to the fundamental conceptual differences that centre on the issue of autonomous interpretation, a matter that up to now has received only scant attention in the literature. In Section 3, we discuss the place of the autonomy hypothesis in a view of phonology which at first blush may seem paradoxical but which is in fact fully in tune with a modular theory of language. On the one hand, the view is 'concrete' to the extent that phonological representations are assumed to be phonetically interpretable at all levels of derivation. On the other, the representations are uniformly cognitive; the

primes they contain, unlike orthodox features, do not recapitulate modality-specific details relating to vocal or auditory anatomy. We then go on to show how this view is implemented in the specification of elements involved in the representation of vowels (Section 4) and consonants (Section 5).

2 Monovalency

Phonological oppositions, it is widely agreed, are inherently binary, rather than multivalued or scalar. That is, each phonological prime defines a bifurcate partition of segments into classes. One question that arises in connection with this observation is whether the two terms of an opposition are equal, in the sense that both have a role to play in phonological processing. If they are, the opposition is said to be equipollent; that is, the prime in question is deemed to have two values, usually expressed as a plus vs. minus co-efficient, both of which are potentially addressable by the phonology. In the case of a private opposition, on the other hand, only one term is considered phonologically significant; the relevant prime is then monovalent, being present in one class of segments and absent from the complement set.

The issue of how the valency of phonological primes is best expressed impinges only indirectly on the issue of whether or not they should be credited with autonomous interpretation. The assumption that features are not independently interpretable is compatible with both monovalency and bivalency. In SPE, all oppositions were treated as equipollent. Direct descendants of this tradition, however, incorporate varying degrees of privativeness. In feature geometry, for example, non-terminal class nodes are inherently monovalent (e.g. Clements 1985a, Sagey 1986). Within the latter model, there is some disagreement about whether the terminal nodes of the hierarchy, the features themselves, are uniformly bivalent, uniformly monovalent, or mixed. (For competing views on this matter, see for instance Avery and Rice (1989), Mester and Itô (1989) and Archangeli and Pulleyblank (1994).)[2]

It is not immediately clear whether adherence to the notion of interpretational autonomy necessarily entails a commitment to privativeness. In practice, however, all approaches based on the autonomy principle subscribe to uniform monovalency.

Drawing direct comparisons between available privative and equipollent models is not at all straightforward, for the reason that other theoretical variables are typically implicated, some of which interact with how the valency of primes is treated. A completely level playing field would require concurrence on such matters as

the basic set of phonological operations (spreading and delinking, for example), the universal set of primes, and the question of whether or not primes enjoy interpretational autonomy. The nearest we get to such controlled conditions is within feature theory itself, where it is possible to isolate privative versus equipollent versions of the framework which differ minimally in most other respects. (Compare, say, Avery and Rice's 1989 monovalent feature approach with the bivalent approach of Archangeli and Pulleyblank 1994.)

Nevertheless, one thing we can be certain of is that, all other things being equal, a full-blown privative model of phonological oppositions has considerably less expressive power than one based wholly on equipollence. This weaker generative capacity naturally favours the privative approach, unless it can be shown to be significantly empirically underpowered. Assuming for the sake of argument that the controlled conditions mentioned in the last paragraph are in place, we can compare how the two models treat the possibilities of spreading involving an opposition defined in terms of the prime [round], such as might be observed in vowel harmony. The equipollent approach predicts three types of harmony system: one in which [+round] spreads, one in which [−round] spreads, and one in which both values are active. In a privative framework, the only expressible spreading pattern is one involving [round]. There is no possible way of defining a complementary system of nonround spreading, since there is no object equivalent to [− round] to which the spreading operation could have access.

A strictly privative model makes the strong prediction that all phonological oppositions will rigidly display this sort of asymmetry. That is, only one term of each distinction, the one possessing the relevant element, has the potential to participate in phonological activity. The element is available for, say, spreading or for blocking the propagation of some other element. By contrast, the complement set of segments, those lacking the element in question, is predicted to be phonologically inert; they will fail to trigger spreading and will be transparent to it.

There is little doubt that asymmetries of this type exist. Markedness theory and its implementation in underspecification frameworks represent attempts to graft this insight onto the equipollent approach (Kiparsky 1982a, Archangeli 1984 and Archangeli and Pulleyblank 1994). According to this notion, only one value of an opposition is specified underlyingly and is thus accessible to phonological processing at the outset of derivation. In the normal case, it is the marked term that is lexically present, while the

unspecified value is filled in by universal default rule. Again in the normal case, the latter operation does not take place until the final stage of derivation, with the result that the unspecified value remains invisible to phonological processes. Thus far, underspecification theory resembles a privative model. However, the power of the former is greatly increased by allowing supposedly universal markedness conventions to be overturned in individual grammars, something that can be achieved in two ways. Universal specification preferences can be reversed such that it is the marked value of a feature that is represented underlyingly. Moreover, an underlyingly unspecified value is permitted to be filled in at any stage of derivation, at which point it is free to become no less phonologically active than the underlyingly specified value.

When coupled to a bivalent feature theory, the set of markedness conventions constitutes an independent look-up table against which the marked/unmarked status of individual feature values is gauged. Within a uniformly privative framework, in contrast, a more radical alternative is adopted whereby markedness relations are built directly into phonological representations. Thus, while the privative and equipollent-underspecification approaches are united in acknowledging the skewed nature of phonological oppositions, they part company on the issue of how the universality of this asymmetry is implemented. According to the equipollent account, the asymmetries are relative; the favouring of one distinctive term over its complement is a matter of preference, potentially reversible on a language-specific basis. In a genuinely privative approach, by contrast, the universality of distinctive imbalances is absolute. Viewed from a bivalent perspective, the monovalency of the privative position implies that one value of each prime is *inherently* underspecified (Archangeli 1988).

Potential counter-evidence to the more restrictive privative model comes in the form of any equipollent analysis which refers to a feature value for which there is no direct privative equivalent. And it has to be acknowledged that such accounts are myriad – hardly surprising, considering the almost unchallenged ascendancy that the equipollent view enjoyed during a period which included the publication of SPE. In some cases, such examples can be straightforwardly reanalysed in terms of the opposite feature value. In other words, they turn out not to constitute counter-evidence at all but rather reflect one of the recurring maladies that is symptomatic of over-generation – the possibility of analysing a single phenomenon in more than one way. For example, McCarthy's (1984a) account of the height harmony system of Pasiego Spanish involves spreading both values of the feature [±high]. Vago (1988) shows how the

spreading process can be reanalysed so as to access only [+high]. An element-based treatment of the same facts, to be briefly outlined below, need refer only to a single monovalent element ([A]).

Nevertheless, there no doubt exists a corpus of more robust equipollent accounts that need to be reassessed on a case-by-case basis. It is clearly beyond the scope of this chapter to tackle such cases one by one. What we can do, however, is provide a brief illustration of the more general issues that arise when particular examples such as these are re-investigated from an element-based perspective.

Let us return to two of the factors cited above as potentially disruptive of any attempt to undertake a direct comparison between the privative and equipollent approaches to phonological analysis. First, disagreement on the issue of autonomous interpretation probably removes the possibility of complete agreement on what constitutes the universal set of phonological primes in the first place. It often seems to be taken for granted that there exists an ultimate level of systematic phonetic representation defined in terms of traditional features (see, for example, Kaye et al. 1985, Pulleyblank, this volume). However, any such presumption immediately queers the pitch in favour of an equipollent account. Segments defined in terms of SPE-type features cannot be considered pre-theoretical entities. For example, in the version of element theory to be presented below, strict adherence to the principle of interpretational autonomy means there is no equivalent of (traditionally bivalent) features such as [consonantal], [strident] or [sonorant].

Second, there is the challenge of pinning down the particular phonological operations that are deemed to manipulate primes in individual phenomena. This issue may appear reasonably straightforward in view of the emerging consensus that phonological processes should as far as possible be reduced to a small set of formal operations. These days, most phonologists, whether of equipollent or privative persuasion, would probably subscribe to the view that all processing should, ideally at least, be reducible to two fundamental operations – linking and delinking, or perhaps more generally composition and decomposition (Kaye et al. 1985, Mascaró 1987). Composition involves the fusion of primes, achieved through spreading or OCP-triggered coalescence. Decomposition describes either the active rupture of associations between primes, achieved through delinking, or the failure of compositional potential to be realized. The complicating factor in any privative–equipollent comparison, as we will see presently, is that we cannot assume prior agreement on which of these basic operations is at work in a particular phenomenon.

With these provisos in mind, we will try to illustrate some of the theoretical variables that need to be controlled for in a comparison between bivalency and monovalency by contrasting the feature [±high] with the element [A]. Although the segment classes defined by these two primes are clearly not isomorphic, the sets of phenomena to which they are applicable overlap sufficiently for at least some degree of close comparison to be possible. The phenomena to be focused on here are two kinds of height harmony, one a 'lowering' type, the other a 'raising' type. Elsewhere these have been analysed as the spreading of [−high] and [+high] respectively. This particular comparison is especially instructive, since it has been alleged that an element-based approach is unable to express the raising pattern (Clements 1991a).

Lowering harmony is widely represented in the central Bantu languages. In the typical case, mid vowels are lexically absent from certain types of suffix and only appear in such positions during derivation as a result of high vowels lowering under the harmonic influence of a mid vowel appearing in the root. This is the pattern encountered in, for example, Luganda (Katamba 1984), Yaka (van den Eynde 1968) and Chichewa (Mtenje 1985). Compare, for instance, high (2a) and mid (2b) suffix vowels in Chichewa (Mtenje 1985):

(2)

		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'

In earlier feature accounts, this type of system was treated in terms of the rightward spreading of both values of [±high] (Katamba 1984, Mtenje 1985). Within an underspecification framework, it can be reanalysed in such a way that only [−high] spreads (from mid root vowels), with [+high] being filled in by default (Harris and Moto 1989).

In element-based analyses, there is nothing akin to fill-in treatments of harmony (or of any other type of phenomenon for that matter). The nearest equivalent is a situation in which one set of harmonic spans within a system remains unaffected by spreading and simply manifests the identity of lexically present elements. Exactly this state of affairs can be shown to hold in the central Bantu case. Alternations between mid and high vowels are treated in terms of the presence versus absence of the element [A] in other

than low vowels. In the central Bantu pattern, mid vowels in harmonically recessive suffixes arise through the rightward spreading of [A] from a mid vowel in the root nucleus (see Goldsmith 1985, Rennison 1987).[3] As shown in (3a), fusion of [A] with the [I] of a suffix such as *-il-* produces the mid-vowel alternant *-el-* in Chichewa.

(3)

(a) *lemb-el-a*

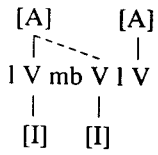

(b) *pind-il-a*

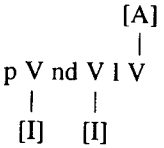

In the complement set of harmonic spans, illustrated by the form in (3b), a high-vowel nucleus in the root lacks [A], and no spreading occurs. The *i* reflex of *-il-* in such cases is then simply the independent realization of the lexically present element [I].

Thus far, [A] and [− high] are roughly equivalent in their coverage of height harmony. The challenge to the element-based approach would now seem to come from feature analyses of harmony which invoke the complement value [+ high]. Just such a case is presented by the treatments of Pasiego Spanish mentioned above.

Very briefly, harmony in Pasiego is controlled by the vowel in the stressed syllable of the word; if it is high, then all non-low vowels to its left are also high. Hence alternations such as the following:[4]

(4)

Infinitive	Future first person singular	Future second person plural	
bebér	*beberé*	*bibirí:s*	'drink'
komér	*komeré*	*kumirí:s*	'eat'
koxér	*koxeré*	*kuxirí:s*	'take'

McCarthy's (1984a) account of this phenomenon actually incorporates two operations: delinking and spreading. According to the first part of his analysis, a lexically present [± high] value in a harmonically recessive nucleus is delinked under the influence of a harmonically dominant nucleus bearing a value for the same feature. Any vowel affected by this process then picks up the dominant value through spreading. Under Vago's (1988) reanalysis, only [+ high] is lexically represented; this spreads from a dominant nucleus to any harmonizing vowels in its domain. In the complement set of harmonic domains, [− high] is filled in by a later default rule.

The element-based analysis of Pasiego to be sketched here retains McCarthy's (1984a) insight that lowering harmony involves delinking (cf. van der Hulst 1988a). On the other hand, it parts company with both types of feature analysis in dispensing with spreading altogether (Harris 1990a). The analysis makes appeal to the licensing relations that obtain between vowels within harmonic spans. A nucleus which determines the harmonic category of a span can be said to license the other nuclei within that domain. The specific proposal is as follows: [A] is sustainable in a licensed nucleus only if it is sanctioned by an [A] in the licensing nucleus. Any [A] that fails to receive such a sanction is delinked. In (5a), the occurrence of [A] in the two recessive nuclei is supported by an [A] in the stressed nucleus. (The arrows indicate the directionality of the licensing relation.)

(5)

(a) *komeré*

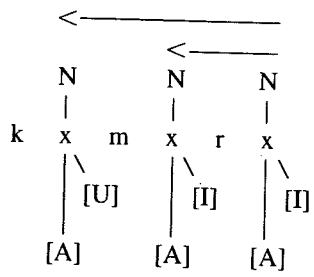

(b) *kumirí:s*

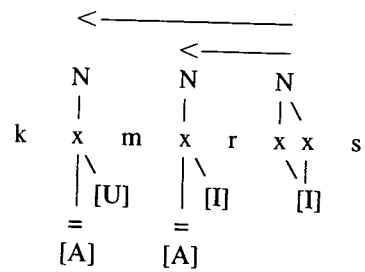

In (5b), in contrast, there is no [A] in the licensing nucleus to sustain the lexically present [A]s in the licensed positions. Delinking (indicated by =) ensues. Each of the residual elements independently defines a high vowel.

Under this analysis, Pasiego harmony is expressed in terms of a single generalization governing the appearance of [A] in adjacent nuclei. Moreover, the mechanism it invokes, inter-nuclear licensing, is motivated by a wide range of facts, most of which are quite independent of height harmony (Kaye 1990a, Charette 1991). These include vowel syncope and metrical phenomena, as well as harmony involving other elements. In a still wider perspective, this mechan-

ism is itself subsumed under the general notion of phonological licensing, the fundamental principle by which all segmental material and constituents, not just nuclear positions, are integrated into the phonological hierarchy (Itô 1986, Kaye et al. 1990).

The process illustrated in (5b) manifests a type of complexity agreement whereby complex vowels in licensed positions undergo reduction if a simplex vowel occupies the licensing position. This is but one instance of a much more general principle according to which the elemental complexity of a licensed position cannot exceed that of its licensor (the Complexity Condition discussed in Harris 1990b). The generality of this principle is reflected in its applicability to an apparently disparate range of contexts that goes beyond the inter-nuclear relation at issue here. Other domains displaying similar complexity effects include those formed by branching onsets and coda-onset clusters.

The delinking analysis of harmony is further supported by the observation that it parallels the treatment of vowel reduction under weak stress. That is, the single formal operation of delinking is invoked as a means of deriving the curtailed distributional potential of prosodically recessive nuclei, irrespective of whether this manifests itself harmonically or metrically. In the Pasiego example, the two aspects of prosodic weakness happen to cooccur: harmonizing nuclei appear in unstressed positions.

Our comparison of height harmony analyses illustrates the more general point that *harmony*, like its hypernym *assimilation*, is a descriptive term with no formal status (see Mascaró 1987). The notion, it is sometimes assumed, can be directly equated with the formal operation of spreading, and this was certainly the state of affairs envisaged in early autosegmental research. In classic (i.e. feature-based) autosegmental treatments of vowel harmony, for example, all harmonic spans in a given system are derived via spreading. That is, both values of the harmonic feature are deemed to spread, as illustrated in several of the feature-based analyses referred to above. (The original example is Clements' 1981 treatment of Akan ATR harmony.) However, in the light of more recent proposals, the claim that assimilation uniformly reduces to spreading is no longer tenable. As shown in our discussion of Pasiego Spanish, two other operations have been invoked: delinking and, in underspecification approaches at least, blank-filling redundancy rules.

The immediate relevance of this point to the treatment of height harmony is that the deployment of [− high] in some feature-based accounts and [+ high] in others does not imply that the monovalency of [A] must be abandoned or that some additional height

element must be posited. The nearest privative equivalent of spread-[− high] (lowering) is spread-[A]; the effect of spread-[+ high] (raising), on the other hand, is achieved by delink-[A].

The absence of a unitary formal treatment of assimilation has a significant impact on how we assess whether each term of an opposition is phonologically active or inert. And this in turn has a direct bearing on how we weigh up the merits and demerits of the equipollent and privative approaches. At the beginning of this section, we noted that bivalency predicts three types of process for each harmonic feature, involving the spreading of either or both values. In fact, taking account now of two possible operations in harmony, spreading and delinking, we find the number of expressible systems increases to six, illustrated in (6) for some given prime [P]. This is in contrast to the two types of system generated by the equivalent monovalent model.

(6)

Equipollent		*Privative*	
Spread	[+ P]; [− P]; [± P]	Spread	[P]
Delink	[+ P]; [− P]; [± P]	Delink	[P]

Including the blank-filling operations of underspecification theory, together with the different orderings these permit, actually multiplies the possibilities defined by bivalency still further. Admittedly, the trend in more recent equipollent work has been to eschew harmonic analyses which manipulate both values of a feature simultaneously. Nevertheless, within this framework the intersection of the two terms of a bivalent feature with the formal operations of spreading, delinking and ordered blank-filling increases the likelihood of more than one analysis being available for a particular harmony system.

To do full justice to the whole issue of monovalency would take up more space than is available here. Nevertheless, we believe that the discussion of the analyses outlined in this section, however brief, illustrates the care that needs to be exercised in comparing the performance of the equipollent and privative models. Given the inherently more constrained nature of the privative approach, the onus is on proponents of equipollence to prove the need to compromise on the universality of markedness asymmetries by showing that both terms of phonological oppositions are potentially phonologically active. The equipollent game-plan should thus be to force advocates of privativeness to admit the necessity of recognizing two monovalent elements for any given bivalent feature. (In the height harmony case, the extra privative element would be

['counter-A'].) If completely successful, the strategy would lead to a situation in which the privative model contained exactly twice as many primes as the equipollent model. At this point, all notion of the absoluteness of universal markedness asymmetries would be relinquished, and the two models would be to all intents and purposes indistinguishable. However, as the foregoing comparison of harmony analyses suggests, any such concession to the equipollent viewpoint should not be made too hastily.

3 Phonetic interpretation in generative grammar

3.1 There is no level of 'systematic phonetic' representation

Recognizing the interpretational autonomy of phonological primes gives rise to a view of phonological representations and derivations which is fundamentally different from that associated with orthodox features. Within the latter tradition, it has long been usual to suppose that only a subset of the feature specifications appearing in final representation is present at the outset of derivation (Halle 1959). Underlyingly absent values, those that are non-distinctive or predictable, are filled in during the course of derivation by redundancy rules or phonological rules proper. This view achieves apotheosis in Radical Underspecification Theory, in which all predictable feature values are stripped from underlying representation (Archangeli 1984, 1988, Pulleyblank 1986a). There has been a slight retreat from this extreme position in recent descendants of the theory. In the Combinatorial Specification approach of Archangeli and Pulleyblank, for example, a criterion of representational simplicity, which favours maximal despecification of lexical representations, may be overridden in certain cases where otherwise non-distinctive values can be shown to play an active role in underlying association patterns (1994: 88–9).

Coupled to the assumption that features lack interpretational autonomy, the classic underspecification arrangement implies that any non-final representation containing blank feature values is phonetically uninterpretable. The realization of lexically represented values is contingent on the support of non-distinctive or predictable values; and the full complement of mutually supporting feature values is not mustered until the final stage of derivation, the level of systematic phonetic representation. This view too has been modified somewhat in more recent underspecification approaches. As a result of work by Keating (1988), Pierrehumbert and Beckman (1988) and others, it is now assumed that fragments of underspecification may persist into phonetic implementation (Archangeli and Pulleyblank 1994: 43, Pulleyblank, this volume). This

situation supposedly arises in certain instances of what used to be termed *sequence redundancy*, where the phonetic interpretation of a segment with respect to a particular feature is entirely predictable on the basis of the value this feature has in neighbouring segments. In a sub-class of such cases, it has been argued, the intervening sound fails to be assigned a value for the feature in question and is thus submitted to motor planning without an inherent articulatory target. The relevant articulatory dimension then allegedly manifests itself through simple linear interpolation between the motor targets associated with the specified values of the flanking segments. (Reservations about the allegedly targetless nature of such segments have been expressed by, among others, Boyce et al. 1991). Even when adapted to allow for sporadic instances of persistent underspecification, it nevertheless remains true of this overall approach that a significant proportion of feature values, particularly those which are predictable from other values within the same segment (*segment redundancy*), must be specified in phonetic representation before phonetic interpretation is possible.

According to this line of thinking, one of the main jobs of the phonological rule component is to transform abstract representational objects into ever more physical ones. This mismatch between underlying and surface representations is sometimes justified on the grounds that the two levels allegedly perform quite different functions: underlying representations, it is sometimes claimed, serve the function of memory and lexical storage, while surface representations serve as input to articulation and perception (Bromberger and Halle 1989). The validity or otherwise of this view hinges on a number of considerations, only some of which we have space to consider here.

One issue concerns the degree of circularity that is inherent in the strategy of stripping lexical representations down to the distinctive bone, a problem fully discussed by Mohanan (1991). In many instances, the presence of a given pair of phonological properties in a representation may involve balanced mutual dependence. That is, in such cases there is no non-arbitrary way of determining the directionality of the relation whereby one property is deemed lexically distinctive and the other predictable and hence derived. A well-known example concerns the relation between syllable-structure and segment-structure information. As observed by Levin (1985), Borowsky (1986) and others, it is often the case that the former can be extrapolated from the latter with a facility equal to that with which the latter can be extrapolated from the former.

Another issue concerns the desirability of having underspecified

representations as addresses for lexical storage and retrieval. The despecification of lexical feature values may be viewed as a type of archiving programme which compresses information for compact storage.[5] As noted by Lass (1984: 205), Mohanan (1991) and others, one of the motivations for proposing redundancy-free lexical representations in generative phonology seems to have been the assumption that long-term memory constraints prompt speakers to limit storage to idiosyncratic information and to maximize the computing of predictable information. This view has never been seriously defended in the psycholinguistic literature.[6] But if underspecification is not justified by storage considerations, nor does it constitute a particularly plausible model of efficient lexical access. The goal of maximal economy of lexical representation can only be achieved at the expense of greatly increasing the amount of computation to be performed at retrieval (Braine 1974). Just as an archived computer text file has to be de-archived before it can be accessed, so would a speaker-hearer have first to 'unpack' the condensed, underspecified form of a lexical entry before submitting it to articulation or recognition.

This last point leads on to perhaps the most fundamental objection to the notion that the function of phonological derivation is to prepare cognitively represented objects for phonetic implementation. Bromberger and Halle (1989: 53) express this notion quite succinctly when they make the claim that systematic phonetic representations 'are only generated when a word figures in an actual utterance'. It is as well to be quite clear about how this view relates to the Chomskyan dichotomy between I-language and E-language. In equating the notion of generation with speech production, it immediately places the phonological component outside the domain of grammar proper. It is not just that the systematic phonetic level constitutes a buffer between the representation of internalized phonological knowledge and its articulatory or perceptual externalization; the phonological component as a whole is geared towards fulfilling a goal which is essentially extragrammatical, that of turning out utterance-bound phonetic forms.

The validity of this view cannot be taken for granted. In what follows, we will consider the consequences of adopting an alternative position, in which phonology remains on the competence side of the competence-performance divide. According to this view, phonological processes, in line with general principles, do no more than capture generalizations holding over alternations and distributional regularities. This function can be served quite independently of any provision that needs to be made for articulation and perception. That is, processes can be construed as purely generative

in the technical sense of specifying the membership of the set of grammatical phonological structures in a language. Under this view, initial and final representations in phonological derivation are isotypic: processes map phonological objects onto other phonological objects rather than onto phonetic ones. Such a view places phonology firmly in the grammatical camp.

If phonological processes map like onto like, it follows that initial representations should be no less phonetically interpretable than final representations. In fact, initial representations can in principle be envisaged as being wholly indistinguishable in kind from final representations. Such an arrangement is of course impossible if non-final representations are subject to underspecification. But it is perfectly consistent with an approach in which phonological primes enjoy autonomous interpretability. In a full-blown element approach, there is no sense in which a final representation is any more physical or concrete than an initial one. There is thus nothing corresponding to a systematic phonetic level, since any representation at any stage of derivation is directly mappable onto physical phonetics.

It is necessary to bear this point in mind when comparing the empirical content of element theory with that of feature underspecification. As noted in the last section, it is sometimes assumed that the two approaches somehow share the same realizational outcome, namely a systematic phonetic representation consisting of fully specified matrices of bivalent features. In the light of the foregoing discussion, it is clear that this opinion is mistaken. The assumption gives the incorrect impression that one theory is directly translatable into the other, an impression that has hardly been discouraged by the occasional practice of explicitly spelling out the phonetic exponence of elements in traditional feature terms (Kaye et al. 1985, Rennison 1990). These researchers are quick to point out that features employed in this way serve no more than a phonetic implementation function and as such are not accessible to the phonology. However, it seems to us better to return to Anderson and Jones' (1974) basic proposal that features have no place in element theory whatsoever. Not only does invoking features muddy the water as far as presentation of elemental phonology is concerned; in several respects, it actually hamstrings attempts to provide an accurate account of how elements are mapped onto physical phonetics, a point we will return to below.

3.2 Jakobson *redivivus*

This last point introduces yet another motive for dispensing with references to orthodox features in specifying the phonetic expo-

nence of elements. The use of SPE-type features entails acceptance of the notion that phonological primes are mapped in the first instance onto the articulatory dimension, in spite of what is usually claimed about a generative grammar being neutral between speaker and hearer. This orientation is perhaps most vividly illustrated in the close congruence that exists between current conceptions of feature geometry and Browman and Goldstein's (1989) gestural model of speech production (Clements 1992). To the slight extent that researchers working within this tradition have concerned themselves with the hearer side of the equation, it has been assumed that the primarily articulatory features can be mapped, albeit indirectly, onto acoustic and perceptual dimensions. (For a recent proposal along these lines, formulated within feature geometry, see Clements and Hertz 1991).

It would perhaps be unfair to speculate that the articulatory slant of much feature theory is not entirely unconnected to the clear articulatory bias of most introductory courses in phonetics. Nevertheless, it is worth pointing out that a belief in the centrality of speech production, however tacit, implies varying degress of commitment to the motor theory of speech perception, the notion that the listener decodes speech by some species of internal articulatory synthesis. The theory, in various guises, has met with rather less than general agreement in the phonetic literature (see Klatt 1987 for discussion). Many phoneticians and phonologists remain to be convinced of the wisdom of abandoning the Jakobsonian insight that the phonetic exponence of subsegmental primes should in the first place be defined in acoustic terms.[7] The speech signal, as Jakobson was wont to point out, is after all the communicative experience that is shared by both speaker and hearer (Jakobson et al. 1952: 13). Its primacy in phonetic interpretation should hardly be in question, at least if we are to pay more than lip service to the idea that generative grammar is neutral between production and perception.

It is for this reason that the specifications of elements provided in the following sections are couched in primarily acoustic terms (an orientation long associated with Dependency Phonology). That is not to say that elements should be construed as acoustic (or articulatory) events. They are properly understood as cognitive objects which perform the grammatical function of coding lexical contrasts.[8] Nevertheless, continuing the essentially Jakobsonian line of thinking, we consider their phonetic implementation as involving in the first instance a mapping onto sound patterns in the acoustic signal. Viewed in these terms, articulation and perception are parasitic on this mapping relation. That is, elements are internally represented pattern templates by reference to which listeners

decode auditory input and speakers orchestrate and monitor their articulations.

4 Elements for vowels

This section covers four main topics: the specification of [A], [I] and [U] (Section 4.1), the headedness of compound segmental expressions (Section 4.2), evidence supporting the existence of a 'neutral' element which defines the resonance baseline on which the other elements are superimposed (Section 4.3), and the representation of ATR (Section 4.4).

4.1 Elemental patterns: [A], [I], [U]

The phonological evidence supporting recognition of the elements [A], [I] and [U] is reasonably well established. For example, we may refer to the pivotal role played by the independent manifestations of these elements in the organization of phonological systems. The 'corner' vowels *a*, *i* and *u* figure with much greater frequency than any other segments in the vocalic systems of the world's languages (Maddieson 1984). Moreover, there is a substantial body of distributional and alternation evidence supporting the conclusion that primes of this nature are individually accessible to phonological processing. Thus we find harmony processes, for instance, which address natural classes identifiable with dimensions such as 'palatality' (characterized by presence of [I]), 'labiality' ([U]), or, as illustrated in the last section, 'height' ([A]). (For examples of all three types of harmony process, see van der Hulst and Smith 1985, van der Hulst 1988b.)

As suggested by representations such as those in (3) and (5), we are following current thinking in assuming that each prime resides on its own autosegmental tier. (The question of whether these tiers are hierarchically organized is something we take up in Section 5.3.) A compound segment thus involves the co-registration of elements on separate tiers. This is illustrated in (7), which shows

(7)

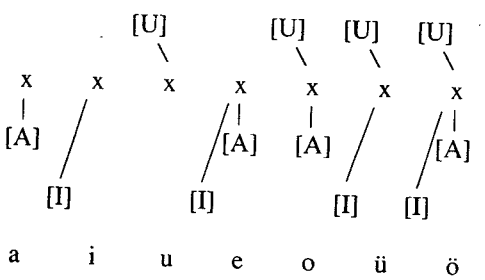

not only the five canonical vowels already given in (1) but also *ü* and *ö*, the two other segments derivable by possible [A]-[I]-[U] combinations of the type we have been assuming up to now.

Of course, not every language opts to exploit the full range of combinatorial possibilities presented by this fully autosegmentalized model. According to Kaye et al. (1985), such systems can be derived by means of parameterized tier conflation. The effect of collapsing two tiers is to prevent the relevant elements from fusing with one another. For example, as depicted in (8), a system in which [I] and [U] take up residence on the same tier is one lacking front rounded vowels.

(8)

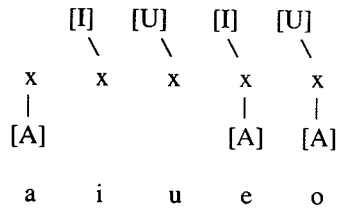

In what follows, we will only indicate autosegmental structure where the discussion demands it.

How are [A], [I] and [U] to be defined, and how do we derive the results of their combination? For reasons explained in the previous section, we begin by proposing for them definitions which may be mapped rather directly onto the acoustic signal. Before embarking on such an exercise, it is as well to ensure that we disabuse ourselves of a common misconception – the idea that quantitative values specified with reference to waveforms or spectrographic displays are to acoustic phonetics what qualitative categories specified in terms of vocal tract diagrams are to articulatory phonetics. In fact, a more apt proportion is one that relates a spectrogram to a three-dimensional X-ray movie of vocal-tract gymnastics. The closest acoustic analogues of the relatively abstract categories associated with vocal-tract diagrams are idealised spectrographic patterns. (The nearest precedents are the 'pattern-playback' diagrams pioneered by Haskins Laboratories; see, for example, Cooper et al. 1952.)

It usually goes without saying that the articulatory specification of phonological representations is appropriately characterized in terms of qualitative categories rather than in terms of the continuously varying quantitative values encountered in speech production.

By the same token, the specification of the acoustic signatures of phonological categories should be couched qualitatively in terms of overall quasi-spectral shapes, and not quantitatively. It would therefore be misguided to express the resonance characteristics of elements such as [A], [I] and [U] as quantitative values relating, say, to formant frequencies. Rather it is necessary to determine the gross quasi-acoustic shapes – what we may call elemental *patterns* – which constitute these categories, and for which symbols such as [A], [I] and [U] are no more than shorthand notations (Lindsey and Harris 1990). The grossness of these patterns should presumably be related to their interpretational autonomy, which they exhibit regardless of whether they appear alone or in compounds.

In Figure 2.1 we diagram the elemental patterns which, we have proposed elsewhere, characterize [A], [I] and [U] (Harris and Lindsey 1991). We display each pattern in a frame mimicking a spectral slice in which the vertical axis corresponds to intensity and the horizontal axis to frequency. The latter coincides with what we may term the *sonorant frequency zone*, the frequency band containing the most significant information relating to vocalic contrasts (roughly speaking between 0 and 3 kHz).

The elemental pattern of [A], shown in Figure 2.1a, is appropriately labelled *mAss*. The signal specification of the vowel *a*, the element's independent manifestation, is a spectral energy mass in the middle of this zone, interpretable as the convergence of Formants 1 and 2;

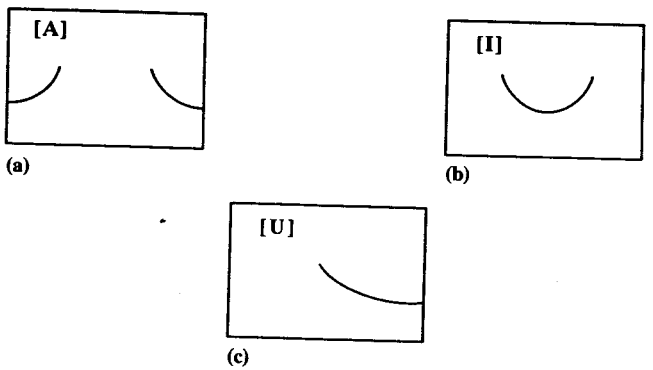

FIGURE 2.1 Elemental patterns: [A], [I], [U] (from Harris and Lindsey 1991). The contours are schematic spectral envelopes, plotted here in frames which map onto acoustic spectral slices (vertical axis: amplitude, horizontal axis: frequency (ca. 0-3kHz)). The solid lines specify regions of low energy. The energy contained within the blank regions is at a higher level than the specified minima, but its precise envelope is not criterial to the definition of the patterns.

that is, crucially there are energy minima at top and bottom of the zone. In the diagram for [A], the precise structuring of the massed energy in the middle of the frequency range is not a criterial part of the pattern's definition and is thus left blank.

The spectrum of *i* contains a low first formant coupled with a spectral peak at the top of the sonorant frequency zone, the latter peak being relatable to the convergence of Formants 2 and 3. This configuration, with energy lower in the middle of the zone than at either side, may be taken to correspond to a *dIp* elemental pattern characterizing [I] (see Figure 2.16).

The signal evidence relating to *u* indicates what we may term a *rUmp* elemental pattern for [U]. The vowel displays a spectral peak at the lower end of the sonorant frequency zone (produced by a convergence of the first and second formants); that is, there is no significant energy above the middle of the zone. The corresponding elemental pattern is displayed in Figure 2.1c; here the lower portion of the contour is left blank, since the precise structuring of the higher-amplitude energy in this part of the frame is not criterial to the pattern's definition.

In summary, the elemental patterns associated with the three resonance elements under discussion are: [A] mAss (energy minima at top and bottom), [I] dIp (energy minimum in middle), and [U] rUmp (energy minimum above middle).

The effects of compounding elements are derived by overlaying elemental patterns on one another. The two complex profiles depicted in Figure 2.2 result from fusing pairs of patterns shown in Figure 2.1. Figure 2.2a shows *e*, the outcome of fusion between [A]

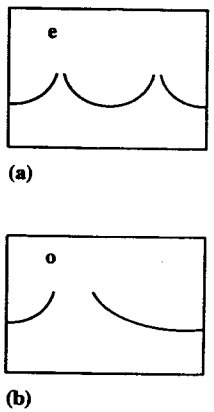

FIGURE 2.2 Compounded elemental patterns: (a) [A, I] *e*, (b) [A, U] *o* (from Harris and Lindsey 1991).

and [I]. The profile here, which might be described as 'dIp within a mAss', can be viewed as an amalgam of two patterns: (i) energy minima at top and bottom, indicating the presence of [A]; and (ii) energy minimum in middle, the pattern associated with [I]. In *o*, a compound of [A] and [U] (Figure 2.2b), we see a pattern that might be dubbed 'mAss at the rUmp'; that is, we have both (i) energy minima at top and bottom, i.e. [A], and (ii) energy minimum above middle, i.e. [U].[9]

Elemental patterns are templates which hearers endeavour to detect in speech input and speakers endeavour to match in the production and self-monitoring of speech output. When a given element is input to speech production mechanisms, the speaker will marshal whatever articulatory resources are necessary or available for the spectral realization of the target elemental pattern. For example, the desired acoustic effect associated with the mAss pattern [A] can be achieved through maximal expansion of the oral tube and constriction of the pharyngeal tube. Note that the imaginary point of maximal height of the tongue body, long cited in impressionistic phonetics (at least since Sweet 1877) as the crucial reference point in the definition of articulatory categories such as 'high' and 'low', has no particular importance in the specification of how [A], or indeed any element, is produced. As is well known, a single vowel sound can be produced with widely varying tongue body contours by different individuals and even by the same individual on different occasions. (See for example, Ladefoged et al. 1972 and the results of bite-block experiments of the type reported by Folkins and Zimmerman 1981.) It is by manipulating the overall shape of the vocal airway that the speaker targets particular acoustic effects. Nor is any elemental specification required to determine that the vocal cords should vibrate. Indeed, if anything, such specification would be undesirable, since the speaker may, depending on circumstance, choose whisper rather than voice as the acoustic source.

The articulatory incarnation of dIp [I] calls for maximal expansion of the pharyngeal cavity and maximal constriction of the oral cavity. The articulatory implementation of rUmp [U] involves a trade-off between maximal expansion of both the oral and pharyngeal tubes. Labial activity (rounding and/or protrusion, for example) is but one of the factors that contribute to the overall size of the oral cavity. One implication of this is that lip rounding is not an articulatory prerequisite in the production of vocalic expressions containing [U].[10]

Consideration of the articulatory implementation of elemental patterns helps underline the distinction between element-based and

current feature-based approaches to the specification of phonetic detail. Feature *under*specification refers to the lexical suppression of properties that are phonologically represented at later stages of derivation. As noted earlier, this notion is incompatible with the autonomous interpretation hypothesis, since specified properties are in most instances uninterpretable without support from temporarily suppressed properties. Implicit in element theory, on the other hand, is the conclusion that certain properties assumed by feature theory are *non*-specified. The non-specification of a property implies that it has no representational status whatsoever, either lexically or during derivation. In fact, some such properties, it can be argued, do not even exist as independent entities in physical phonetics. Note that non-specification, as understood in this sense, in no way impairs the interpretational autonomy of elements. The element [A], for example, has a direct and independent physical interpretation that can be defined without so much as even a passing reference to features such as [± back], [± low], [± sonorant], [± consonantal], or whatever. Recognition of this point in recent element-based research is reflected in the abandonment of the earlier notion (outlined in Kaye et al. 1985, for example) that phonetic implementation involves the translation of each element into a traditional distinctive feature matrix. Such a featural halfway house, it is now acknowledged, is both logically and empirically redundant.

A further consequence of this view is that there is no place in element theory for anything resembling the featural interpretation of the traditional phonemic notion of contrastivity. A familiar feature matrix includes values whose primary function is to distinguish the sound it specifies from other sounds with which it is in opposition. The matrix for *a*, for example, might include [− round] (whether lexically present or filled in by redundancy rule), which helps differentiate the vowel from a round vowel such as *u*. In element theory, on the other hand, *a* is identified solely on the basis of the only element of which it is composed, namely [A]. To be sure, there is a descriptive sense in which a sound may be viewed as entering into Saussurian relations of contrast with other sounds in a given system; but this does not necessarily imply that each such distinction is directly coded in representation as a particular unit of segmental content. In element terms, the definition of *a* does not include reference to properties that identify other vowels with which it happens to be in contrast, such as the [U] present in *u*. That is, *a* is **non**-specified with respect to the pattern characteristics that are relevant to the definition of *u* (acoustically, concentration of energy in the low frequencies, typically realised in articulation by lip rounding). No specification of [U]-related characteristics

('non-rUmp', 'non-round', or whatever) ever enters into the identification of this vowel. Not being *u* is no more a criterial property of *a* than not being an orange is a criterial property of a banana.

4.2 Elemental weightings in compound expressions

Most element-based approaches incorporate some means of representing the notion that the phonetic manifestation of a compound segment reflects the weighting of one element over others occurring in the same expression. The precise implementation of this idea varies from one theory to another.

In Particle Phonology (Schane 1984a), preponderance is formalized by allowing multiple occurrences of the same element to be stacked within a single expression. The relative openness of a vowel, for example, can be reflected in the number of [A]s it contains. The sum of [A]s contained in the expression of a particular vowel is a language-specific matter, varying according to the number of distinctive heights exploited in the system. The same object, *a* say, then receives different representations in different grammars; it might consist of two [A]s in one system and only one in another. The relativism that is implicit in this arrangement clearly represents a retreat from the view that elements are universally defined and uniquely interpretable.

In other versions of element theory, preponderance is represented through relations of headedness between elements occurring within the same expression. In Dependency Phonology, a pair of primes α and β can enter into one of three relations: (a) α dependent on β, (b) β dependent on α, and (c) mutual dependency (Anderson and Jones 1974, Anderson and Ewen 1987). In Kaye et al. (1985) and van der Hulst (1989), on the other hand, only relations (a) and (b) are recognized. Assuming the latter mode of representation, let us compare two [A, I] compounds, one headed by [A], the other by [I]. Informally, we can think of the [A]-headed expression as a palatalized version of an essentially open vowel; the [I]-headed expression meanwhile can be considered an open version of an essentially palatal vowel. These asymmetric fusions are assumed to define the vowels *e* and æ respectively, as shown in (9a) (head element underlined).

(9)

(a)	[A, <u>I</u>]	e	(b)	[A, <u>U</u>]	o
	[<u>A</u>, I]	æ		[<u>A</u>, U]	ɒ

As shown in (9b), the contrast between *o* and *ɒ* is treated in parallel fashion. (The contrast between ATR (advanced tongue

root or tense) *e*/*o* and non-ATR (lax) *ɛ*/*ɔ* has also to be characterized somehow, a matter we take up in Section 4.4.)

In terms of its effect on the signal, intrasegmental dependency is reflected in the predominance of one elemental pattern over another. The compounded profile of [A, I], for example, can be interpreted as a relatively less salient mAss pattern located in the middle of a relatively more salient dIp. The expression [A, I], on the other hand, is realized as a preponderantly mAss pattern with a less salient dIp at its centre. These two elemental profiles simulate the spectral characteristics of *e* and *æ* respectively.

The fusion asymmetries illustrated in (9) provide a straightforward means of representing widely attested processes of raising and lowering involving low and mid vowels.[11] Take for example the raising of *æ* to *e* and of *ɒ* to *o*. (As attested in one portion of the English Great Vowel Shift, for example; see Schane 1984b and Jones 1989 for element-based analyses.) In terms of their elementary make-up, the inputs and outputs of each of these raisings are isomers (Kaye et al. 1985). That is, the elements of which they are composed are identical but are arranged in different ways. Viewed in this manner, the raising of low vowels involves neither the loss nor the addition of elements; it consists rather in a switch in the headedness of the relevant segmental expression. Lowering of mid vowels presents the inverse operation, in which [A] in a compound switches from dependent to head status.

The introduction of intrasegmental relations, it has to be acknowledged, adds to the expressive power of the theory. One way of preventing this power from getting out of control is to take compensating steps towards reducing the number of primes. In fact, much work in this area has been directed towards striking a balance between structural and elementary modes of representation. The enrichment of intrasegmental structure has enabled Dependency Phonologists to reduce the manner dimension of segmental contrasts to two fundamental primes or 'gestures'. This avenue has been most thoroughly explored by van der Hulst (this volume), who reduces all segmental content to two such gestures. This is balanced by an elaborated model of dependency relations, in which segmental contrasts are represented in terms of X-bar structure. This contraction in the role played by segmental primes comes at the expense of sacrificing the notion of autonomous interpretation. Under this view, stand-alone units such as [A], [I] and [U] are derived from syntactic configurations of two primes which do not enjoy independent interpretability. There is apparently a minimum set of primes below which it is impossible to maintain the interpretational autonomy principle. That is, below this point primes become

'too small' to be independently interpretable. On current assumptions, the threshold is somewhere around ten elements (a total not much smaller than the number of primes proposed by Jakobson et al. 1952).

When we turn our attention to compounds containing three or more elements, two issues arise which have a significant impact on the generative capacity of element theory. One has to do with whether or not the fusion of elements occurs in a pairwise fashion. If it does, as assumed by Kaye et al. (1985) for example, then multiple element compounds can be compiled with different patterns of embedding. For example, a three-element compound composed of [X], [Y] and [Z] could be constructed as [X, [Y, Z]], [Y, [X, Z]], or [Z, [X, Y]]. Second, there is the issue of whether any expression may form the head of a compound, irrespective of whether it is itself simplex or complex. In the three-element example [X, [Y, Z]], either [X] or [Y, Z] could in principle act as head of the expression. Together, the possibilities of pairwise fusion and free directionality of headedness greatly inflate the generative power of the theory to an extent that is empirically unjustified. And the expressive potential of course increases exponentially as more elements enter the equation.

These issues are fully discussed in a series of exchanges between Coleman (1990a, 1990b) and Kaye (1990c), in the light of which it seems prudent to abandon the notion of pairwise fusion. In the simpler model that results, one which has no recourse to nested compound structures, an unbounded relation exists between the head element of an expression and any dependent elements that may also be present. This is the model we will assume in the remainder of this chapter, although we will only indicate the headedness of a compound expression where the context demands.

Before leaving this topic, we should note that the free combinability of elements is further constrained by quite general and independent principles. One of these is the Complexity Condition referred to in Section 3, which has the effect of severely restricting the number of elements that can appear in adjacent positions.

4.3 The neutral element

Occurring singly or in combination with one another, [A], [I] and [U] help define the basic set of vocalic contrasts given in (7). Additional elements are necessary for the definition of other dimensions of contrast, including that of peripherality, to which we now turn.

The spectral peaks associated with *a*, *i* and *u* are inherently large and distinct. From an articulatory point of view, this is a reflection

of the fact that these vowels, the universally limiting articulations of the vowel triangle, represent extreme departures from a neutral position of the vocal tract. The supralaryngeal vocal-tract configuration associated with the neutral position approximates that of a uniform tube and produces a schwa-like auditory effect. The resonating characteristics of this configuration are such that the first three formants are fairly evenly spaced, with the result that it lacks the distinct spectral peaks found in *a*, *i* and *u*. Most researchers within the element-based tradition accord this neutral quality some special status, either by treating it as a segment devoid of any active elementary content or by taking it to be the manifestation of an independent element, which we will symbolize here as [@]. Broadly speaking, this corresponds to the centrality component in Dependency Phonology (Lass 1984, Anderson and Ewen 1987), to the 'cold' vowel of Government Phonology (Kaye et al. 1985), and to an 'empty' segment lacking any vocalic content in Particle Phonology (Schane 1984a) and in the work of van der Hulst (1989).

The element [@] may be thought of as a blank canvas to which the colours represented by [A], [I] and [U] can be applied. From a production point of view, this metaphor reflects the point just made that [A], [I] and [U] are realized by means of articulatory manoeuvres that perturb the vocal tract from its neutral state. From a signal point of view, this implies that the dispersed formant structure of [@] constitutes a base line on which the elemental patterns associated with [A], [I] and [U] are superimposed.[12]

The idea that [@] is signally blank may come as a surprise to those who assume that the phonologically relevant *dramatis personae* of vocalic acoustics are formants. To be sure, schwa-like vowels exhibit formants just as much as other vowels. However, in the light of what was said in Section 4.1, it is an assumption we should have no truck with. It is elemental patterns that are the phonologically relevant actors to be detected in vowel signals. Their absence from the dispersed acoustic spectrum associated with [@] indicates a stage on which the phonetic sound and fury of [@]'s formants phonologically signify nothing.

In phonemic and quasi-phonemic transcription 'ə' is frequently employed as a cover symbol to designate a vowel-reduction reflex. The range of phonetic qualities indicated by the symbol in the literature is quite impressive; in terms of the traditional vowel diagram, it covers varying degrees of openness, backness and roundness. (In transcriptions of Catalan, for example, it symbolizes a relatively open value, in Moroccan Arabic relatively close, and in French front rounded.) Some aspects of this variability are phonologically insignificant, but others evidently involve distinct phonological categories. In the former case, variability across languages can be taken to reflect indeterminacies in the fixing of the base line on which other resonance components are superimposed. From a speech production viewpoint, this variability is sometimes characterized in terms of different articulatory or vocal settings. (For a review and discussion of the relevant literature, see Laver 1980.)

In element theory, the independent realization of [@] may be understood as covering that area of the traditional vowel diagram which is non-palatal, non-open and non-labial. Non-peripheral categories that are potentially distinct from this base line can then be thought of as displaced versions of the neutral quality. In terms of their segmental make-up, these different reflexes can be characterized as compounds in which [@] is fused with some other element(s). The relatively open ə of Catalan, for example, can be represented as a combination of [@] and [A].

The idea that [@] defines the base line on which other resonances are superimposed can be implemented by assuming that it does not reside on an independent autosegmental tier. Rather it is omnipresent in segmental expressions but fails to manifest itself wherever it is overridden by any other element(s) that may be present. Viewed in these terms, reduction to a centralized vocalic reflex does not involve the random substitution of one set of elements by [@]. Rather it consists in the stripping away of elementary content to reveal a latently present [@].

One of the advantages of viewing centralization in this way is that it unifies the representation of the process with that of certain processes of raising and lowering which, although not involving reduction to non-peripheral reflexes, nevertheless occur under the same prosodically weak conditions. A widespread phenomenon in the world's languages is a tendency for mid vowels to be banished from prosodically recessive nuclear positions. In metrical systems, recessiveness refers to positions of weak stress; in harmony systems, it refers to nuclei whose harmonic identity is determined by an adjacent dominant nucleus. Under such conditions, it is common to find neutralization of vocalic contrasts in favour of either non-peripheral reflexes or the 'corner' vowels *a*, *i*, *u* or some mixture of both. The height harmony systems reviewed in Section 3 are uniformly of the peripheral type. Non-harmonic cases exhibiting similar patterns of neutralization include Bulgarian and Catalan, in which raising and centralization cooccur. In Bulgarian, the stressed five-vowel system (*i*, *e*, *a*, *o*, *u*) contracts to three vowels under weak stress, with the mid vowels raising to high and *a* undergoing centralization (Petterson and Wood 1987). The stressed seven-term system of Catalan (*i*, *e*, *ɛ*, *a*, *ɔ*, *o*, *u*) gives way to three terms under

weak stress: the *a-ɛ-e* contrast is neutralized centrally, the *ɔ-o-u* contrast is neutralized under *u*, and *i* remains as it is (Palmada 1991: ch. 2).

The fact that the processes just mentioned, raising or lowering of mid vowels and centralization, all potentially occur in the same general recessive context indicates that we are dealing with a single phenomenon. Although this commonality has long been recognized, it has not always been clear how it should be captured formally. In terms of element structure, however, the processes in question are uniformly expressible as decomposition; all involve the total or partial suppression of segmental material.

Summarizing the foregoing discussion of [@], we may identify two main respects in which it differs from other elements. First, it lacks the distinct peak-valley patterns of [A], [I] and [U]. Second, it is latently present in all segmental expressions. One question that remains is how the latter notion is to be accommodated in the mechanism of element fusion. According to Kaye et al. (1985), the desired result can be achieved by assuming that the only circumstances under which [@] contributes anything to the phonetic interpretation of a compound segment are when the element acts as the head of the expression. It will thus fail to make its presence felt in any expression in which it occurs as a dependent. The only circumstances under which latent [@] can become audible are when it is promoted to headship as a result of other elements in the expression undergoing suppression or relegation to dependent status.

4.4 ATR

Let us now turn to ATR, another dimension implicated in peripherality contrasts, most famously those associated with some kind of harmonic alternation. The treatment of ATR harmony remains a hotly disputed topic (one that is potentially bound up with height harmony, since, for some researchers at least, the two dimensions are of a piece). Issues on which there continues to be rather less than general agreement include the following. What are the relevant primes? Is there evidence to support the recognition of retracted tongue root harmony systems? How are transparency and opacity effects involving ATR to be represented? It is not part of our brief to examine these points in detail. All we can do is make some very general observations about the status of ATR in element theory.

Broadly speaking, there have been two element-based approaches to the representation of ATR. One is to posit an independent [ATR] element, the solution preferred by Kaye et al. (1985). Another is to derive the distinction between ATR and non-ATR vowels structurally, by means of different combinations of [A], [I], [U] and [@]. One version of the latter approach resorts to the element-stacking device mentioned in Section 4.2. According to Smith (1988) and van der Hulst (1989), ATR *i*, for example, contains two instances of [I] (with differing dependency status), while non-ATR *ɪ* contains only one. (In Dependency Phonology, ATR is represented both structurally and in terms of an independent prime; see Anderson and Ewen 1987.)

There are at least two factors weighing against the positing of a privative [ATR] element. For one thing, it detaches the theory from the original insight that the bounds of vowel space are defined by the extremes represented by *i*, *a* and *u*. It implies that non-ATR vowels are less complex than their ATR congeners. The independent manifestations of [I] and [U] are thus taken to be non-ATR *ɪ* and *ʊ* respectively – not *i* and *u*, which are now represented as [I, ATR] and [U, ATR]. (The realization of [A] remains as *a*, which is non-ATR in any event.) The prediction inherent in this arrangement – that the unmarked three-vowel system is *a-ɪ-ʊ* – fails to tally with the empirical record.

There is a further, this time theory-internal, reason for being suspicious of an independent [ATR] element. As Kaye et al. (1985) acknowledge, it is anomalous in being the only element whose contribution to the make-up of compound expressions must be considered constant, irrespective of whether it is a head or dependent. According to these authors, the ability of elements to combine within a compound is controlled by their 'charm' values. (In brief, elements of opposite charm attract, while those of like charm repel one another.) The overall charm value of an expression is determined by the head element. This is assumed to hold of all cases – except those involving [ATR], which anomalously imposes its charm value on an expression even as a dependent.

There are reasons for supposing that ATR can and should be derived by exploiting otherwise well-established properties of the theory rather than by adding to the pool of elements. The question is whether a purely structural definition of ATR is achievable without sacrificing the principle of interpretational autonomy through recourse to element stacking. One proposal is that non-ATR high and mid vowels, unlike their ATR counterparts, contain an active [@] (Lass 1984: 277 ff). Given the recessive behaviour of this element in compounds, this implies that such vowels are [@]-headed. Thus, as before, the basic set of non-low vowels is assumed to be ATR *i*, *u*, *e*, *o*, as shown in (10a). The non-ATR set is then represented as in (10b).[13]

(10)

(a) ATR		(b) non-ATR	
[I, @]	i	[I, @]	ɪ
[U, @]	u	[U, @]	ʊ
[A, I, @]	e	[A, I, @]	ɛ
[A, U, @]	o	[A, U, @]	ɔ

Besides representing the unmarkedness of *i*/*u vis-à-vis* *ɪ*/*ʊ* (Maddieson 1984, Lindsey 1990), this arrangement makes sense from a signal point of view. In non-low vocalic space, each non-ATR vowel is characterized by an attenuation of the well-defined peak-valley pattern that is associated with its ATR counterpart. This is consistent with the mode of representation of *ɪ*, *ʊ*, *ɛ*, *ɔ* given in (10b). This allows us to view the elemental patterns contributed by the active resonance elements [A], [I] and [U] as being muffled as a result of their being subordinated to the neutral pattern defined by head [@]. It should be emphasized that it is not formants but the patterns of other elements that are attenuated by [@] under such circumstances. The autonomous realization of [@] as some schwa-like object constitutes 'pure' attenuation, the absence of other elements' patterns.

The question now is whether this treatment is capable of representing processes involving ATR, particularly those harmonic cases which have been analysed in feature terms as the autosegmental spreading of [+ATR] or [−ATR].[14] In the absence of an independent privative [ATR] element, we have to assume that all such cases conform to the non-spreading pattern illustrated in the analysis of Pasiego discussed in Section 3. That is, they involve changes in the internal representation of harmonizing vowels, triggered by particular conditions obtaining in the dominant vowel within the harmonic span. ATR alternations take the form of switches in the headship of vocalic expressions, with [@]-headedness in non-low vowels representing non-ATR. Note that this does not involve the insertion or spreading of [@]. Given the latent presence of [@] in all segments, such alternations are entirely isomeric. That is, the manifestation of this element within a particular vocalic expression is simply a reflection of its promotion to head status. ATR harmony is thus a matter of what might be termed *head agreement*; that is, the head elements of all vowels within a given harmonic span are aligned (Lowenstamm and Prunet 1988, Charette and Kaye 1993). In Akan (Clements 1981), for example, the head elements of non-low vowels are aligned as in (11).

(11)

(a) ɛ-bʊ-ɔ 'stone'

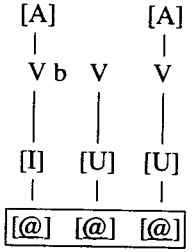

(b) e-bu-o 'nest'

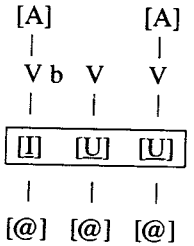

Either all vowels within a span are [@]-headed, in which case a non-ATR domain is defined (as in (11a)), or they are headed by an element on the [U/I] tier, in which case an ATR domain is defined (as in (11b)).[15]

5 Elements for consonants

In this section, we focus almost exclusively on the resonance characteristics of consonants (Section 5.1) and on the dimension traditionally referred to as 'manner' (Section 5.2). Lack of space precludes us from discussing laryngeal contrasts in any detail, and we will abstract away from this dimension in much of what follows. (On the nature of independent laryngeal elements, see Kaye et al. 1990 and Brockhaus 1992.) Nor will we have anything to say about nasality; given the current state of our knowledge, it is not clear whether this should be represented by an autonomous nasal element or is more appropriately subsumed under one of the laryngeal elements. (On the phonological evidence supporting a laryngeal–nasal connection, see, for example, Piggott 1992.) We conclude the section with a consideration of the applicability of segmental geometry to element theory.

5.1 Resonance elements

It is now widely acknowledged that the resonance/cavity characteristics of consonants and vowels are represented in terms of the same set of primes, rather than in terms of separate sets as assumed in SPE (e.g. Smith 1988, Clements 1991b). This commonality is most obviously illustrated in the case of glides, which are segmentally identical to vowels and differ only in terms of their syllabic affiliation. The transcriptional distinction between *u* and *w*, for example, records the difference between occupation of a nuclear head and occupation of any other type of position (such as an onset). In terms

of segmental content, we are dealing with the same object – an expression containing a lone [U]. A similar point can be made in respect of *i* versus *y* (= IPA *j*), except in this case the relevant element is [I]. There is plenty of phonological evidence to support this glide-vowel identity. In many languages, alternations between high vowels and glides are straightforwardly analysable as the reassignment of [I] or [U] from a nuclear to an adjacent onset position, as in French *avu – avwe* <avoue – avouer> 'confess', *si – sye* <scie – scier> '(to) saw' (Kaye and Lowenstamm 1984).

There are good grounds for assuming that this commonality extends to the resonance characteristics of non-vocalic segments. According to one element-based proposal, the resonance properties contributed to consonants by [A], [I] and [U] vary according to their status as heads or dependents (Smith 1988, van der Hulst 1989). Our more optimistic assumption, following Jakobson, is that the elements in question maintain relatively stable, albeit gross patterns. So, for example, the rUmp elemental pattern [U], besides inhering in labial vowels, also shows up as a spectral associate of labial consonants (what Blumstein and Stevens 1981 identify as a 'diffuse-falling' spectral pattern).

[I] is present in palatal and palatalized consonants, while [A] is present in uvulars and pharyngeals. According to one view, coronality requires an additional element ([R]), on which more presently. The exponence of [@] can be informally described in articulatory terms as non-coronal, non-palatal, non-labial and non-low, which suggests that it should be considered the resonance element in velar consonants.

The assumption that the resonance elements are shared by all types of segments, vocalic as well as non-vocalic, is supported by a range of assimilatory processes showing consonant-vowel interactions. These are straightforwardly treated in terms of element spreading, as in the palatalization of consonants before front vowels, where the active element is [I]. In similar fashion, [U] is implicated in labialization ([U]), while [A] is active in the lowering of vowels in the context of a uvular or pharyngeal consonant.

Non-assimilatory processes, particularly those involving lenition, provide further support for the notion that the resonance elements inherent in consonantal segments are identical to those in vowels. One class of lenition process takes the form of vocalization, the reduction of a consonantal segment to its homorganic vocalic counterpart. (In the case of a plosive target, the process historically passes through a fricative stage, as in $b > \beta > w$.) Typical vocalizing outcomes include the following:

(12)

	Residual element
(a) p → w, m → w	[U]
(b) č → y, λ → y	[I]
(c) k → (γ →) θ	[@]
(d) t → r	[R]

Each of these is exemplified in the following representative alternations:[16]

(13)

(a)	Korean:	*ki:p-t'a – kiw-ə* 'sew'
	Irish:	*mo:r – wo:r* (fem.) 'big'
(b)	Arbore:	*gerrač – gerray-me* 'thief'
	Spanish:	*vaλe – vaye* 'valley'
(c)	Turkish:	*inek – ine*[ø] *– i* 'cow'
	Welsh:	*garð – (ei* 'his') [ø]*arð* 'garden'
(d)	English:	< a[t]omic > – < a[D]om >
	Korean:	*si:t-t'a – sir-ə* 'to load'

Processes of this type are straightforwardly represented as the suppression of all elementary content save that relating to resonance (Harris 1990b). In each example in (13), the residual reflex reflects the primary resonance property of the leniting segment. Each of these properties can be equated with an element, as shown in (12).

Vocalization of velars (13c) typically results in reduction to zero, sometimes via *y*. This development is not unexpected, given the assumption that velar resonance is associated with the element [@]. Independently, [@] manifests itself as approximant *y* (non-syllabic i), but the lack of an active resonance component in this element is predicted to make it particularly likely to be eclipsed when not supported by other elementary material.

It would be consistent with the line of reasoning pursued thus far to consider the tapped *r* reflex in (13d) to be the independent realization of a coronal element, [R]. This is indeed the position taken, for example, in Kaye et al. (1989) and Harris and Kaye (1990), and it is the one we will simply assume here without further argument. However, certain considerations suggest that this view is probably in need of reappraisal. For one thing, a single pattern signature for [R] has proved somewhat elusive. (Some attempt at a definition is made in Lindsey and Harris 1990.) For another, there is a growing body of evidence indicating that specific representational provision needs to be made for the special status of coronals among the resonance categories of consonants. (The issues are conveniently summarized in Paradis and Prunet 1991a.) Among the well-known

peculiarities are the following: coronals are more prone to assimilation than other classes; consonant harmony exclusively affects coronals (at least in adult language); and coronals, unlike other resonance classes, behave transparently with respect to many processes.

Facts such as these have prompted a variety of analyses in which coronals are represented as 'placeless' consonants. In Dependency Phonology, for example, it has been suggested that they lack a place component (Anderson & Ewen 1981). According to various feature-geometric analyses, coronal is deemed the default place category, underlyingly unspecified for the PLACE node. (For a selection of examples, see Paradis and Prunet 1991b.) In this way, it is possible to account for the propensity for coronals to assimilate: they are underspecified at the point where spreading of the PLACE node takes place. For the same reason, they can be transparent to vowel-harmony processes involving the spreading of PLACE.

Treating coronality as an independent [R] element, on a par with any other, fails to capture the special properties of coronals. Following the lead of the analyses just mentioned, we might suggest that coronality has no elementary representation. However, the implications of such a move for element theory would be much more radical than anything countenanced in current underspecification approaches. It would banish coronality from phonology altogether. That is, *coronal* would be similar to *glottal* in being an exclusively articulatory detail. (This line of enquiry has recently been pursued by Backley 1993.) The question is whether this move can be made without jeopardizing the insight that glottal segments are the classic 'placeless' reduction consonants.

5.2 'Manner'

From the perspective of element theory, vocalization is on a representational par with the lowering and raising of mid vowels discussed in Section 2. All of these processes take the form of decomposition – the dissolution of compound segmental expressions resulting from the suppression of elementary material. The significance of such processes is that their outcomes allow us to identify the independent manifestations of individual elements. The same principle can be applied to the task of determining further elements implicated in consonantal contrasts, for example those that are suppressed in the vocalization processes illustrated in (13). An obvious place to start looking in this case is debuccalization, the process by which the resonance properties of consonants are stripped away. Lenitions of this type can reasonably be expected to lay bare those elements that are associated with the 'manner' dimensions of consonants.

As far as obstruents are concerned, there are two types of

weakening process that can be considered to have this disclosing effect. These occur on two lenition trajectories, which may be schematized as follows (cf. Lass and Anderson 1975: Ch. 5):

(14)

(a) Spirantization > 'aspiration' > deletion
plosive > fricative > h > ∅

(b) Loss of release > glottalling > deletion
plosive > unreleased stop > ? > ∅

Lenition trajectories such as these are established on the basis of cross-linguistic observations of the directionality of diachronic change. There is no implication that every lenition process inexorably culminates in elision. Historical progression through the various stages on a particular path may be arrested at some point, with the result that two or more stages on a particular trajectory may be retained within the same phonological grammar as stable alternants or distributional variants. Typical alternations involving the 'opening' types of lenition schematized in (14) include the following:[17]

(15)

(a) Central American Spanish: *mes* – *meh* – *me* 'month'
(b) Tiberian Hebrew: *malki* 'my king' – *melex* 'king'
(c) Malay (Johore): *masak-an* – *masaʔ* 'to cook'
(d) English: <ge[t]no> – <ge[t⁻]no> – <ge[?]no>

The weakest sounds on a lenition path are those occupying the penultimate stage; they represent the last vestige of a segment before it disappears altogether. Combining the autonomous interpretation hypothesis with a view of lenition as segmental decomposition leads us to conclude that the penultimate stages in (14), namely *h* and *ʔ*, are primitive segments; that is, each is the independent embodiment of a single element. The recognition of these two elements, [h] and [ʔ] (Kaye et al. 1989), is strongly reminiscent of Lass and Anderson's (1975) insight that the segments in question are the 'reduction' consonants *par excellence*. Elsewhere we have explored the consequences of taking [h] and [?] to be the main 'manner' elements in consonantal contrasts (Harris 1990b, Lindsey and Harris 1990).

The elemental pattern associated with [?] may be described as *edge* or *stop*. In signal terms, it manifests itself as an abrupt and sustained drop in overall amplitude. This effect is achieved by a non-continuant articulatory gesture of the type that characterizes oral and nasal stops and laterals. The independent manifestation of [?] as a glottal stop is due to the fact that the element lacks any

inherent resonance property. In element theory, glottal place is thus non-specified in the sense described in Section 3.1. The use of glottal location to produce the independent manifestation of [?] is a purely articulatory affair: it is the only articulatory means of orchestrating the amplitude drop of the relevant elemental pattern without introducing resonance characteristics into the signal.

This point underlines the care that needs to be exercised in making sense of the notion that elements are, as is sometimes claimed, individually *pronounceable* (Kaye et al. 1985: 306). To say that each element is independently interpretable is not to say that it can be targeted by executing a unique articulatory gesture. The performance of a particular elemental pattern typically involves the arrangement of one or more of an ensemble of gestures.

In compound segments, the constriction necessary to produce the edge pattern of [?] will be located at whatever place produces the acoustic effects associated with a resonance element occurring in the same expression. For example, fusion of [?] with [U] implies labial constriction. Stops with other places of articulation are formed by fusion of [?] with [A] (uvular), [I] (palatal), [@] (velar), or, if such an element is recognized, [R] (coronal).

The elemental pattern of [h] may be identified as 'noise', manifested in the speech signal as aperiodic energy. The articulatory execution of this effect involves a narrowed stricture which produces turbulent airflow. Noise defined in these terms is present in released obstruents (plosives, affricates, fricatives) but is absent from sonorants and unreleased oral stops. Just as with [?], the absence of any supralaryngeal gesture in the independent articulation of [h] is entirely a function of the fact that it lacks its own resonance property. In compounds with other elements, however, the location of the noise-producing gesture will be determined by whatever resonance element may also be present.

The contrast between strident and non-strident fricatives can be expressed in terms of different relations of headedness involving [h] and those elements which contribute supralaryngeal resonance. Strident fricatives, in view of their relative noisiness (in the sense of displaying higher-intensity aperiodic energy), may be assumed to be [h]-headed, unlike their non-strident counterparts:

(16)

[h, U]	f	[h, U]	ʍ
[h, R]	s	[h, R]	ɵ
[h, I]	ʃ	[h, I]	ç
		[h, @]	x
[h, A]	χ	[h, A]	ħ

Treating all types of lenition as segmental decomposition implies that movement along any of the trajectories in (12) or (14) takes the form of decomplexification – a progressive depletion of the stock of elements contained in a segment. Let us pursue the consequences of this view for the various stages on the opening trajectory in (14a). If *h* is the least complex segment, the plosive input must be the most complex and oral fricatives of intermediate complexity. An oral fricative differs from *h* by one degree of complexity: as indicated in the last paragraph, the former contains a resonance element that is absent from the latter. By the same token, the internal structure of a plosive includes whatever elementary material is present in a homorganic fricative but is more complex than the latter by virtue of the presence of an additional element, the stop property represented by [?].[18] This line of reasoning leads us to conclude that [h] inheres in all released obstruents, both plosives and fricatives. Thus weakening along trajectory (14a) may be expressed as the progressive suppression of elementary material, here illustrated by labials:

(17)

Spirantization/aspiration	p	[h, U, ?]
	f	[h, U]
	h	[h]

Spirantization is frequently preceded by affrication, as illustrated in the High German Consonant Shift: $p > pf > f$; $t > ts > s$; and $k > kx > x$ (the *kx* reflex is only attested in some dialects). We make the standard current assumption that affricates are contour segments consisting of stop and fricative expressions attached to a single position. In element terms, this implies that [h] and [?] are not fused in such structures. Thus the [U] that is present in, say, *pf* is separately fused with [?] and [h].

The difference between released and unreleased oral stops is represented in terms of the presence versus absence of [h]: plosive *p* is [h, U, ?], whereas unreleased p^- is [U, ?]. Suppression of [h] is thus one stage on the lenition trajectory in (14b), while debuccalization of an unreleased stop to *?* results from the suppression of its resonance element:

(18)

Loss of release/glottalling	p	[h, U, ?]
	p^-	[U, ?]
	?	[?]

Vocalization of a plosive consists in the suppression of both [h] and [?], leaving a lone resonance element:

(19)

Vocalization	p	[h, U, ?]
	w	[U]

Vocalization of a true voiceless obstruent also implies the suppression of the relevant laryngeal element. (No such change accompanies the vocalization of so-called 'partially voiced' or neutral obstruents, since they already lack a laryngeal element.) By its very nature, a vocalic segment is produced with physiological voicing. But it would be misguided to attempt to represent vocalization as the acquisition of an active laryngeal prime (such as [+voice]). The phonetic voicing of a vocalized reflex, being of the spontaneous type, is simply a secondary effect of the phonological process which strips away the segment's obstruent-defining properties.[19]

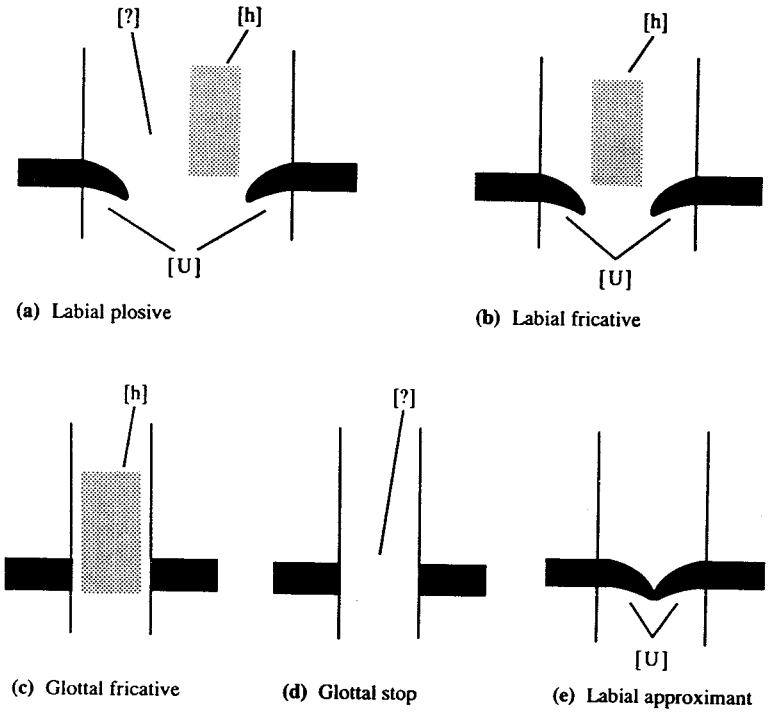

FIGURE 2.3 Stylized spectrograms showing various combinations of the elements [?] (edge), [h] (noise), and [U] (rUmp). Different types of lenition are defined by the suppression of particular elemental patterns: (b) spirantization, (c) 'aspiration' (debuccalization to *h*, (d) glottalling (debuccalization to *?*, and (e) vocalization.

Elsewhere, we have demonstrated how the various effects of lenition can be simulated by excising certain portions of the acoustic signal associated with plosives. Each of these portions can be taken to instantiate a particular elemental pattern (Lindsey and Harris 1990). The results relating to the various stages in (17), (18) and (19) are illustrated in the stylized spectrograms shown in Figure 2.3.

The term *fortition* or *strengthening* is usually applied to processes which turn approximants or fricatives into homorganic stops. That is, they represent the converse of a subset of the lenition processes schematized in (12) and (14). The classic fortition cases are assimilatory in nature. For example, a fricative may be hardened to a plosive in the context of a nasal consonant, as in Sesotho *fa* 'give'— *m-phɛ* 'give me'. The strengthening effect commonly displayed by nasal stops indicates that they, just like oral stops, contain [?]. Fortition in such cases is thus straightforwardly treated as the spreading of [?] from the nasal into the fricative. Fusing [?] with [U, h] (= *f*) yields [?, U, h], a labial plosive.

Different head-dependent configurations in expressions containing [?] enable us to capture various types of manner and resonance contrasts among non-continuant consonants. Labial-velar stops, like their approximant and fricative congeners, can be assumed to be [U]-headed. This leaves plain labial stops as [?]-headed. On theory-internal grounds, velars must be considered [@]-headed. This follows from the assumption that [@] is unable to make an active contribution to the resonance profile of a segment unless it occurs as the head of the expression.

The contrast between laterals and coronal stops, it has been argued, is also a matter of headship (Kaye et al. 1989). According to this analysis, *l* is [R, ?], while coronal stops are [R]-headed. The phonological plausibility of the relation that is predicted by this arrangement is based on frequently observed alternations between *l* and *d*/*t* (e.g. Sesotho *bal-a* – *bad-ile* 'read/count'). It also has a certain phonetic plausibility: whether or not [R] is the head of an expression might be expected to be reflected in the degree of lingual contact associated with the closure contributed by [?]. The [R]- headedness of coronal stops is consistent with full contact, both medial and lateral. The partial (medial) contact required for the production of laterals is consistent with the relegation of [R] to dependent status.

5.3 Elements and geometry

Together with nasality, the elements [?] and [h] allow us to represent all the major manner contrasts among constricted segments. The nearest corresponding feature values might be taken to be [−continuant] and [+continuant] respectively, but the equivalence

is only very rough. For one thing, [h], unlike [+continuant], inheres in plosives. Note, moreover, that the version of element theory presented here lacks anything equivalent to the features [sonorant] or [consonantal]. These categories are becoming increasingly anomalous within current feature theory in any event. It is quite possible that, like the now-abandoned feature [syllabic], they represent obsolescent throw-backs to pre-linear theory. That is, the major class distinctions they were originally designed to express are more appropriately represented not in terms of primes but in terms of prosodic constituency. (This may not seem quite so obvious in the case of [sonorant], particularly in view of the widespread process of final obstruent devoicing, the phenomenon perhaps most often cited as evidence in support of the feature. However, this argument is no longer persuasive, in view of the increasing acknowledgement that only obstruents bear a distinctive (non-spontaneous) voice feature and are thus the only segments susceptible to this type of devoicing (see, for example, Rice and Avery 1990, Brockhaus 1992).)

In recent feature theory, [sonorant], [consonantal] and other manner features, particularly [continuant], have had something of a chequered history. This is reflected in ongoing disagreements over whether an independent manner node needs to be recognized in the hierarchical arrangement of features that is assumed in most current research. This mode of representation is motivated by the insight that primes, no less than segments, pattern into natural classes. If we are to maintain the constraining principle that each phonological process can access only one unit in a representation, then we need to make the assumption that primes are hierarchically organized into classes (Clements and Hume, forthcoming). The by-now familiar geometric model in which this assumption culminates

(20)

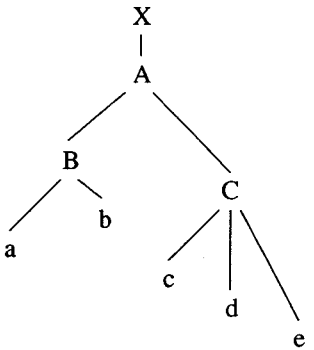

is illustrated in (20), where upper-case letters stand for class nodes, while the terminal nodes in lower-case represent primes (Clements 1985a, Sagey 1986, McCarthy 1988, Pulleybank, this volume). Under this arrangement, each process may address either an individual prime or a class node. In the latter case, the whole class of primes dominated by that node is automatically also affected.

The class nodes for which there is the firmest empirical support are those shown in the geometric fragment in (21).

(21)

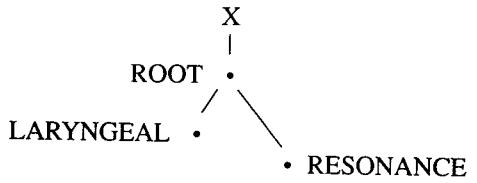

Each of these nodes can be shown to correspond to a particular recurrent grouping of primes. Under the RESONANCE (or PLACE) node are gathered those primes which can be observed to pattern together, for example, in place assimilation processes. The LARYNGEAL node is motivated by the independent and unified behaviour of primes involved in tonal and phonation-type contrasts. Both of these class nodes are grouped under ROOT, the matrix node which defines the integrity of the melodic unit. The latter concept is justified on the basis of the potential for segments to be spread or delinked in their entirety.

The nodes in (21) have a purely organizing function; that is, they are devoid of any intrinsic phonetic content. In element geometry, this can be assumed to be a defining property of class nodes; in keeping with the autonomous interpretation hypothesis, only elements themselves are independently interpretable. In this respect, feature geometry is quite different, in part a reflection of the articulatory bias of the theory. Particular organizing nodes in the feature hierarchy are deemed to have 'intrinsic phonetic' content (Sagey 1986). Of this type are the so-called articulator nodes, e.g. [coronal], under which may be grouped, depending on the version of the theory, such features as [anterior], [distributed], [lateral] and [strident].

If there is more or less general agreement on the class nodes given in (21), this can certainly not be said of nodes which have been posited for grouping manner and major-class distinctions. Even among phonologists who assume there is evidence to support such a view, there is considerable disagreement about where such nodes should be located in the geometric model and about whether these dimensions are themselves subdivided into further class nodes.

(For competing views on this matter, see, for example, Clements 1985a, Sagey 1986, McCarthy 1988, Avery and Rice 1989 and the articles in Paradis and Prunet 1991b.) Take the features [sonorant], [consonantal] and [continuant] for example. According to one proposal, the first two of these are lodged in the ROOT node, which directly dominates [continuant] (among other things) (e.g. McCarthy 1988). Another approach assumes the same relation between ROOT and [continuant] but posits an independent SUPRALARYNGEAL node dominating [sonorant] and PLACE (e.g. Avery and Rice 1989). In their gestural analogue of feature geometry, Browman and Goldstein (1989) suggest a third arrangement, under which separate specifications for [constriction degree] (roughly equivalent to [continuant]) are attached to each terminal node in the hierarchy specifying place.

Whichever of these solutions is adopted, none is able to represent the lenition processes discussed in the last section in a unified manner. Debuccalization to *h*, for example, calls for the simultaneous delinking of PLACE and a change from [−continuant] to [+continuant]. Vocalization of plosives requires simultaneous changes in [sonorant] (minus to plus), [consonantal] (plus to minus) and [continuant] (minus to plus). In both cases, arbitrary conjunctions of features and nodes have to be manipulated, in clear violation of the principle that each phonological process should only be allowed to address a single representational unit. And in the first case the two independent operations are of different types: delinking and feature-change.

These particular problems, we believe, are ungainly artefacts of the articulatory pre-occupation of orthodox feature theory. More generally, this bias can be considered ultimately responsible for the lack of agreement on the featural representation of manner contrasts.

(22)

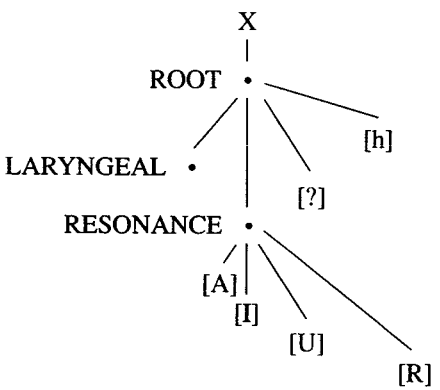

Freed of articulatory bias, element theory is able to represent lenition in a uniform and direct manner. Assuming the rather simple geometric arrangement in (22), we can express each step along a weakening trajectory as the delinking of a single element or node. Given the notion that lenition is uniformly expressed as decomposition, the number of weakening processes to which a segment is susceptible is logically limited by the number of elements of which it is composed. Debuccalization involves delinking of the RESONANCE node (and thus any element it dominates). In spirantization, it is [?] that is delinked. Subsequent vocalization implies delinking of [h], resulting in a loss of release burst. None of these operations needs to make reference to anything resembling major-class features. In keeping with the autonomous interpretation hypothesis, the outcome of any decomposition process is automatically defined by the independent manifestation of any element that remains present in the representation. Once a given element is delinked, no auxiliary operations are needed to adjust the representation to ensure that remaining segmental material can be phonetically interpreted.

6 Summary

The elements of phonological representation are monovalent entities which enjoy stand-alone phonetic interpretability throughout derivation. This conception of phonological primes informs a view of derivation which is in the strict sense generative and which dispenses with a systematic phonetic level of representation.

Elements are cognitive categories by reference to which listeners parse and speakers articulate speech sounds. They are mappable in the first instance not onto articulations but rather onto sound [*sic*] patterns. They constitute universal expectations regarding the structures to be inferred from acoustic signals and to be mimicked in the course of articulatory maturation. No phonological process, we claim, requires modality-specific reference to auditory or vocal anatomy. In particular, we deny any need to recapitulate the latter in the manner of feature-geometric biopsies.

Notes and Acknowledgements

1. Our thanks to Outi Bat-El, Wiebke Brockhaus, Phil Carr, Jacques Durand, Edmund Gussmann, Doug Pulleyblank and Neil Smith for their helpful comments on earlier drafts.

2. Some versions of Dependency Phonology permit segment classes to be identified by means of the Boolean operator ~ (e.g. ~ [A] 'not [A]'), a move which immediately reclassifies an opposition as equipollent (Anderson & Ewen 1980).

3. The low vowel *a* in Chichewa and related systems does not trigger lowering harmony, e.g. *bal-its-a* 'give birth (causative)' (* *bal-ets-a*). Moreover, it blocks the rightward propagation of harmony, e.g. *lemb-an-its-a* 'write (reciprocal-causative)' (* *lemb-an-ets-a*). In an element-based approach, this effect is derived by specifying that [A] is harmonically active only when it occurs as a **dependent** within a segmental expression (Harris and Moto 1989). (On the notion of intra-segmental dependency, see Section 4.2.) [A] is thus active in *e* and *o*, where it has dependent status, but is inert in *a*, where it is the head of the segment.

4. For the purposes of this comparison, we may set on one side the treatment of low vowels, which exhibit neutral behaviour in this system. An element-based analysis of this phenomenon (see Harris 1990a) exploits the notion of intrasegmental dependency, in a fashion similar to the Bantu height analysis just discussed.

5. We owe this analogy to Jonathan Kaye (personal communication).

6. This issue is not addressed in recent psycholinguistic applications of feature underspecification (e.g. Lahiri and Marslen-Wilson 1991, Stemberger 1991).

7. There have always been at least some dissenting voices against the Gadarene rush from acoustic to articulatory features, Andersen's being a particularly eloquent example (1972, 1974: 42–3).

8. As Phil Carr has pointed out to us, this means that phonetic events cannot be considered *tokens* of element *types*. Assuming such a relationship of instantiation would imply that the two sets of entities were of the same ontological status. This view is incompatible with the notion that elements, unlike physical articulatory and acoustic events, are uniformly cognitive.

9. Precise modelling of the computations by which the speaker-hearer detects such patterns in acoustic signals is of course no trivial matter. But we assume it is no easier, and probably more difficult, to model detectors of [high], [low], [back] and [round] in X-ray movies of natural speech articulation.

10. This means that our occasional labelling of [U] for convenience as 'labiality' must be taken with an especially large pinch of phonetic salt. Jakobson's term *flatness*, regrettably out of fashion except among Dependency Phonologists (Anderson and Ewen 1987), is much to be preferred.

11. Again, precise modelling of the cognitive computations by which headedness can be detected in acoustic signals is hardly likely to be a simple research programme. (For a preliminary element-based attempt, see Williams and Brockhaus 1992.) But there is no reason to suspect that the challenge is in principle any more difficult than the corresponding one of modelling the detection of assorted values of [high], [low] and [back] in articulatory movies.

    It should be noted that taking relative and gradient signal preponderance as the phonetic interpretation of phonological headedness does not entail a gradient conception of headedness itself. The number of ways in which multi-element compounds may be headed is constrained by purely phonological considerations.

12. By analogy with the terms mAss, dIp and rUmp, we might dub the element [@] **neutr@l**.

13. As noted above, inherently non-ATR *a* is still to be taken as simplex [A]. The neutral behaviour this vowel typically exhibits in ATR harmony systems is then related to its representational distinctiveness *vis-à-vis* non-low vowels, which are either [@]-headed (non-ATR) or headed on the [I/U]-tier (see the discussion of Akan below).

14. Not all harmony systems that have been analysed in terms of the feature [±ATR] should automatically be submitted to the treatment outlined here. A subset of such systems can be shown to involve the spreading of other elements. For reanalyses of some allegedly ATR patterns in terms of [A]-spread, see for example van der Hulst (1988b) and Anderson and Durand (1988).

15. The participation of *a* in a more restricted form of harmony in Akan requires a separate analysis. For the arguments, see Clements (1981) and Archangeli and Pulleyblank (1992: 186).

16. Sources for the data in (13): Sang Jik Rhee and Yong Heo (personal communication) (Korean), and Hayward (1984) (Arbore).

17. Sources for the data in (14): James Harris 1983 (Spanish), Leben 1980 (Hebrew), Farid 1980 (Malay).

18. The notion that the relation among plosives, fricatives and *h* involves a progressive loss of closure is made explicit in the description of opening provided by Lass and Anderson (1975).

19. For this reason, Lass and Anderson's (1975) term for this phenomenon, **sonorization**, seems more appropriate than **voicing**.