The skeleton*

Péter Szigetvári

Only theories of phonology that attach significance to representations of phonological objects and, in addition, subscribe to an autosegmental version of these representations face the question of what the phonological skeleton looks like. Therefore, this chapter presupposes an autosegmental view of phonological representations.

The motivation for the autosegmental model is the fact that the segmentation of the speech signal can never result in absolutely discrete segments. Here segmentation is taken to mean practically the conversion of the continuous speech signal into the alphabetical symbols of the IPA. Some of these symbols pertain to more than one segment, for example, the stress mark to the syllable after it, tones potentially to even longer stretches. Take the question *you live by the sea?* Its last word, carrying the most prominent stress in the sentence, the tonic, might be transcribed as [↗siː]. In this transcription, the tone mark has a scope lasting all through the word (basically its only vowel): the pitch rises on steadily until the end of the utterance. The same holds if the string after the tone mark is longer, for example, as in *you live by the ↗seaside, Martin?* It would take a very complicated mechanism to maintain that pitch was a property of individual segments and in some cases this rising pitch was realized on a single vowel, while in others it was split into low, higher, even higher, and highest pitch and added to several other vowels following. Tone is clearly not an immanent property of a vowel; it is an ephemeral phenomenon (from the point of view of a vowel) controlled by syntactic and pragmatic factors. If so, it is useful to represent it separately from the rest of the properties of the sound string. Such autonomous sound properties came to be known as autosegments.

* I acknowledge useful comments and advice from two anonymous reviewers, Marc van Oostendorp, Ádám Nádasdy, Péter Siptár, and László Varga. I thank them all. Remaining errors are mine.

If the phonological shape of an utterance is represented as a string of discrete feature bundles, the only option of representing the rising pitch in ↗*sea* includes a feature [rising tone] (here R) in the set of features corresponding to the vowel, as in (1a). In ↗*seaside, Martin,* on the other hand, a set of features [low tone], [higher tone], [even higher tone], etc. (here 0H, 1H, etc.) has to be assigned to the vowels following the tonic, as in (1b).

(1) a. s iː^{R}
b. s iː^{0H} s aɪ^{1H} d m ɑː^{2H} t ɪ^{3H} n

The greatest flaw of such representations is obvious: there appears to be nothing in common between the two rising tones, nothing to indicate their relationship. It is clear that the same tone is spread over the available vowels, but this does not show at all in (1).

Not only tone, but many other sound properties turn out to be similarly promiscuous, with the potential of simultaneously belonging to several segments, and being manipulable independently of the segment(s) they belong to. (For further discussion, see chapter 9 AUTOSEGMENTS.)

The more sound properties extracted from their feature bundles, the less there remain. There are two widespread views on how many: according to one — historically the earlier —, one feature remains in the "bundle", [syllabic] (e.g., McCarthy 1979, Halle & Vergnaud 1980, Clements & Keyser 1983), according to the other, none (e.g., Levin 1985, Lowenstamm & Kaye 1985). The string of segmental positions thus vacated is called the phonological skeleton (a name suggested by Halle & Vergnaud (1980 : 83)) — or, alternatively, timing tier or skeletal tier. The former type, in which the skeletal positions hold the feature [syllabic] is the CV skeleton (discussed in §2), the latter, absolutely empty one is the X skeleton (discussed in §3). A non-segment-based framework involving only syllables and moras is introduced in §4. I will then argue that there is a way of incorporating moras in the old CV skeleton with downright advantages over the moraic framework (§5).

To begin with, let us examine the types of relation that may exist between skeletal positions and phonetic features associable to them.

1 Melody–skeleton relations

Skeletal positions represent the presence of a segment, they serve as an anchoring site for the phonetic properties associated with that portion of the speech signal. If the relationship of feature bundles—referred to as melody, also following Halle & Vergnaud (1980) among many others—and skeletal positions were always one to one, the latter would be superfluous. But, as we have already seen in the case of tone, this is not so. Let us take the nontrivial options one by one.

1.1 One-to-many relations

The standard textbook examples for this type of skeleton–melody relationship are affricates and prenasalized plosives. To focus on the less exotic type: phonologists have long been upset by the feature [±delayed release] (and the marginal oppositions it creates), which Chomsky & Halle (1968) introduced to distinguish affricates from plosives. The altenative approach (discussed by Gimson (1989 : 172f) and Roca (1994 : 3f), among others), that affricates are bisegmental, suggested by the IPA symbols used to represent them, is undermined by many facts. In most cases, the distribution of affricates shows that they are not clusters, but single segments. It even occurs that an affricate does not contrast with a homorganic fricative in a system (e.g., Castilian Spanish has [ʧ], but no [ʃ]), rendering the cluster analysis more than unlikely.

The separation of quantity (skeleton) and quality (melody) offers an opportunity for handling the quantitatively simplex, but qualitatively complex affricates in an intuitive way, as so-called contour segments. (2) depicts the view of the affricate [ʦ] along these lines. (The skeletal slot is represented as '×', but this is not meant to indicate a standpoint in the CV vs. X skeleton debate.)

(2) *An affricate as a contour segment*

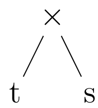

The representation in (2), however, incorporates a misconception, namely, that the melody of segments, without the slot they attach to, forms some kind of unit, two of which are here associated with a single skeletal slot. In reality, the letters 't' and 's' above have no theoretical status. What exist in an autosegmental framework (or for that matter any other phonological theory since the middle of the last century) are features, many of which are common to the two parts of the affricate (place of articulation, laryngeal properties). Another difficulty with the contour model of affricates lies in the interpretation of autosegmental representations. Any melody linked to a slot of the skeleton — also known as the timing tier — is to be interpreted simultaneously. Temporal sequencing is managed by the skeleton, that is, what is linked to an earlier slot is interpreted earlier than what is linked to a later slot. Associating the stop part of the affricate to the left leg of the contour segment and the fricative part to the right is then just a graphical trick, which in theory cannot have any realizational consequences. The standard solution of this problem, involving root nodes, is commented on in §4. As Clements (1999) and chapter 4 AFFRICATES argue affricates are best thought of as noncontour segments (strident stops), as Jakobson & al. (1952) have proposed.

It seems then that we are left without one-to-many relations between the skeleton and melodic material. In fact, such relations are the most commonplace occurrences in representations, since it is not segments but features that are associated with the slots of the skeleton. Thus, most segments embody the one-to-many relation, as the partial representations of two very common segments, [d] and [ɑ], show in (3).

(3) *Partial autosegmental representation of* [d] *and* [ɑ]

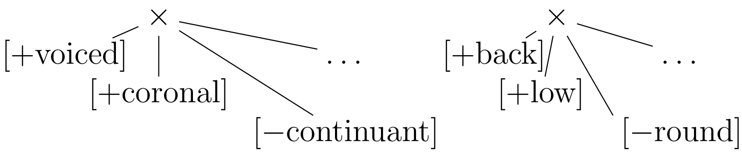

1.2 Many-to-one relations

The long–short contrast of a vowel could be encoded in a feature [long], so that the long vowel is [+long], the short one is [−long]. It is evident however that this is not an adequate way of modelling length contrasts. Vowel length (or consonant length for that matter) is not a property like vowel height (or the voicing of obstruents): it does not harmonize or trigger or suffer assimilation of any type. Furthermore, changes in segmental length are usually unlike common assimilatory changes. Take, for example, the so-called Rhythmic Law of Slovak, which shortens a suffixal long vowel after a long vowel in the stem. The agentive *-ník* (acute accent marks length) inherently contains a long vowel (e.g., *rol-ník* [rolɲiːk] 'farmer'), which shortens when added to a long-vowelled stem (e.g., *stráž-nik* [straːʒɲik] 'guard'; Kenstowicz & Rubach 1987). The rule could be categorized as a dissimilatory process. What is conspicuously missing in languages is any assimilation of this type: i.e., changes where a short vowel would lengthen in the vicinity of a long vowel, and, crucially, because of that long vowel, or a long vowel would shorten purely because of the shortness of a neighbouring vowel.

An even more telling phenomenon is compensatory lengthening (also see chapter 20 COMPENSATORY LENGTHENING).[1] A synchronic comparison of the forms of the 1sg copula in two varieties of

[1] Much of the literature limits the term COMPENSATORY LENGTHENING to cases involving the lengthening of a vowel. The lengthening of a consonant is called inverse compensatory lengthening by Hayes (1989 : 280–281). Here I will refer to both processes by the same name.

Ancient Greek, Attic [eːmi] and Aeolic [emːi] ‘I am’ suggests a simple shift in the host of the alleged feature [+long]. In light of the reconstructed Proto-Greek etymon *[esmi], however, a different analysis is called for. The loss of the [s] triggers the lengthening of one of the neighbouring segments, that of the preceding vowel in Attic, and that of the following consonant in Aeolic. If length were encoded by a feature, the change could only be described by a pair of simultaneously applying rules, one deleting the coda consonant, the other lengthening the segment next to the deletion site. It is clear that the two rules are interrelated: spontaneously neither open-syllable lengthening is attested in Attic, nor intervocalic gemination in Aeolic, these changes only occur in tandem with the loss of the coda consonant. It is difficult to understand why these two rules occur together so commonly. If the quantity of segments is stored separately from their quality, this process, and any similar one, obtains a very neat explanation: it is only the quality (melody) of the coda [s] that is lost — more precisely only its association with the skeleton is lost —, its place, that is, the time it had occupied in the string of sounds is retained (cf. Ingria 1980, Steriade 1982, Hock 1986, Hayes 1989, among others). It is this empty place that one of the neighbouring segments fill in, as shown in (4). (In this and the following diagrams, dotted lines stand for severed association lines, dashed lines stand for newly established association lines, i.e., spreading.)

(4) *Compensatory lengthening: the stability of the skeleton*

a. Attic

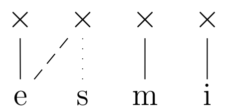

b. Aeolic

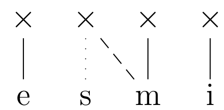

While cases like the above could also be analysed as the total assimilation of the [s] to the preceding vowel or the following consonant, there are more complicated types of compensatory lengthening, for which such an analysis is not the least viable. Cases in point include Middle English *tale* [talə] > [taːl] (Minkova 1982), Old Church Slavonic *bógŭ* > pre-Serbo-Croatian *bôg* [bóog] ‘god’, *bobŭ́* >

bóób [boób] ‘bean’ (Hock 1986 : 435), or Old Hungarian [hida] > *híd* [hiːd] ‘bridge’, [lɛvɛlɛ] > [lɛvɛːl] (Modern Hungarian *levél* [lɛveːl] ‘leaf’, with subsequent closing of the second vowel, E. Abaffy 2003 : 331). The English case is debated (Lahiri & Dresher 1999), but for the others there is evidence that they are not cases of open syllable lengthening ensued by apocope. In Slavonic the original bisyllabic stress pattern is preserved on the long vowel of the monosyllabic forms. In Hungarian no lengthening takes place before suffixes that retain the stem-final vowel: ModHu *hidam* [hidɑm] ‘my bridge’, *levelek* [lɛvɛlɛk] ‘leaves’. Lengthening due to a minimal word constraint is also excluded by the last example: the process takes place in mono- and polysyllabic words alike.

The bipositional status of long vowels is also made likely by the fact that they behave similarly to “vowel clusters”, that is, diphthongs. In English, for example, neither category occurs before noncoronal consonant clusters and both occur word finally, unlike short monophthongs (Fudge 1969 : 272f, Harris 1994 : 37, Gussmann 2002 : 20–23, see Prince 1984 for the same conclusion in Finnish for both vowels and consonants). Accordingly, there is a general consensus that long vowels ought to be represented as in (5a), long (i.e., geminate) consonants as in (5c). The representation of a diphthong and another consonant cluster is given in (5b) and (5d), respectively, for comparison.

(5) *The autosegmental representation of vowel and consonant clusters*

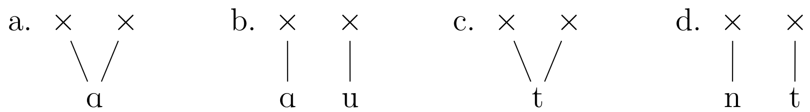

It is not only complete segments that may be linked to more than one skeletal position. The standard situation in fact is that features (autosegments) are multiply linked. Take, for example, the Hungarian word *különbség* [k̟ylø̃mpʃeːg] ‘difference’, depicted in (6). (The features only serve illustrative purposes, their exact identity and location is irrelevant here.)

(6) *Multiply linked features in the representation of* [k̟ylø̃mpʃeːg]

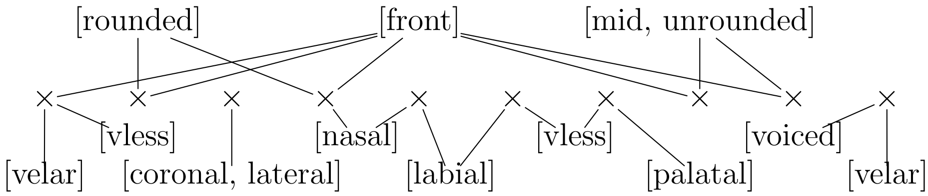

Chaotic as it seems, the diagram in (6) does not contain all the relevant features specifying the segmental content of the string [k̟ylø̃mpʃeːg], manner of articulation features, for example, are all missing. Nevertheless, it can clearly be seen that it is more common for a feature to be associated with several skeletal slots, than to be associated exclusively with one. In (6) this is because of voicing, place of articulation assimilations, vowel nasalization, consonant fronting, as well as vowel harmony taking place. The reasons for associating a single feature to successive slots, instead of associating separate instances of the same feature to each slot are elaborated in chapter 84 OCP.

1.3 One-to-zero and zero-to-one relations

As we have seen, many-to-one and one-to-many relations between the skeleton and melody are very common. Two further options are discussed in this section. It is possible that a skeletal position is not associated with any melodic material. The opposite case may also occur: features unlinked to any point on the skeleton.

French liaison exemplifies both of these possibilities. The phenomenon is well known: a word-final consonant is pronounced when the next word begins with a vowel, but not when it begins with a consonant. (The subtle syntactic conditions on liaison need not concern us here.) Thus in the phrase *petit garçon* 'little boy' the first element ends in a vowel ([pəti garsõ]), in *petit enfant* 'little child' a [t] is pronounced at its end ([pətit ɑ̃fɑ̃]). According to one analysis (e.g., Prunet 1987 : 226) *petit* comes with only four skeletal slots, but five segments, *enfant* on the other hand has an extra skeletal slot, it begins with an initial consonantal slot which is empty. The situation is shown in (7).

(7) *Liaison*

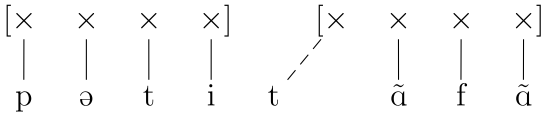

The [t] at the end of *petit* is not associated to the skeleton, it is said to be floating. Floating melody fails to be pronounced unless it gets some chance of associating to the skeleton. Vowel-initial words supply an empty skeletal position that the floating melody can associate to. The floating [t] at the end of *petit* must be lexically determined: there are other liaison consonants besides [t], their identity is unpredictable (e.g., *gros enfant* [groz ɑ̃fɑ̃] 'fat child', *mon enfant* [mõn ɑ̃fɑ̃] 'my child', *gentil enfant* [ʒɑ̃tij ɑ̃fɑ̃] 'nice child', *long article* [lõg artikl] 'id.', etc., where the consonant before the space appears only if the next word begins with a vowel). Therefore this consonant must be included in the lexical representation. It is also not unjustified to suppose that vowel-initial words carry an empty skeletal slot at their left side. It is true for all languages that at least some words (and syllables) begin with a consonant. For some languages this is not an option but an obligation, but, crucially, there are no languages where this could not occur. One may argue that a syllable-initial consonantal position is in fact obligatory in all languages, the optionality is whether this position may or may not be left empty (see, e.g., Kaye 1989 : 134). Thus consonant-initial words do not carry an empty skeletal slot at their left side, vowel-initial words do, and as a result, the latter can host the floating consonantal melody at the end of the preceding word. Apparently, even languages that allow syllable-initial consonantal positions to be empty aim at them being filled.

Hypothesizing that there is an empty skeletal position between two vowels in hiatus and that languages make an effort to fill it also explains the prevalent practice of hiatus filling. Unless a language manages to get rid of this consonantal position (often together with one of the neighbouring vowels), an intervocalic consonantal position is filled by some melody associating to it from one of the vowels.

(8a) illustrates this taking English *skier* and (8b) Hungarian *síel* [ʃiːɛl] 'he/she skies' as examples.

(8) *Hiatus filling*

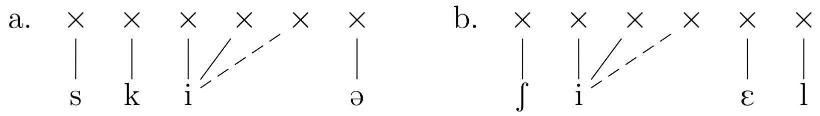

The hiatus between [iː] and [ə] or [ɛ] is filled by the melody of the first vowel, resulting in the forms [skiːjə] (Gimson 1989 : 215, 2001 : 213) and [ʃiːjɛl] (Siptár & Törkenczy 2000 : 283).

The possibility of vocalic positions being empty is surrounded by a significantly greater degree of suspicion, this issue will be taken up in §5.

2 The CV skeleton

The notion of the CV tier was originally developed for the analysis of the nonconcatenative morphology of Classical Arabic by McCarthy (1979, 1981). Like in other Semitic languages, a large part of morphological categories are not expressed by linking morphemes after one another, but by fusing individually unpronounceable components into one. A similar, but much less elaborate case is the ablaut found in Germanic languages, e.g., English *sing, sang, sung,* and *song,* where the consonants carry the lexical entry and the vowel the grammatical category. (Also see chapter 117 SEMITIC TEMPLATES.)

Paradigms in Arabic are classifed into groups traditionally called conjugations — or, as McCarthy refers to them, binyans. The prime phonological property of a binyan is the order in which consonants and vowels are arranged. Roots of three (sometimes two or four) consonants contribute a lexical field to the meaning, the vowels are often responsible for grammatical categories like tense and voice. A portion of McCarthy's (1979 : 244) table depicting the forms for the root $\sqrt{\text{ktb}}$ 'to write' is given in (9).

(9) *Some forms of* $\sqrt{ktb}$

binyan	perf-act	perf-pass
I	katab	kutib
II	kattab	kuttib
III	kaatab	kuutib
⋮		
IX	ktabab	—

The CV skeletons of the first three binyans are then CVCVC CVCCVC, and CVVCVC, respectively, that of binyan IX is CCVCVC. The root consonants and the vowels supplied by the grammatical category are mapped onto this skeleton more-or-less according to the association conventions elaborated by Goldsmith (1976). Three cases are shown in (10).

(10) *The construction of* katab, kattab, *and* ktabab

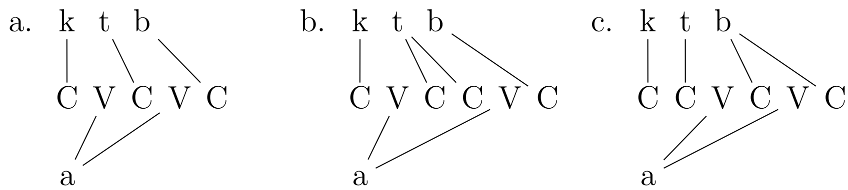

In (10a), the consonants are linked to the C slots of the skeleton, one by one. It is vital that the consonantal and vocalic skeletal slots be distinguished, since the linking of the root consonants and the vowel(s) can be done as required only thus. The case of (10c) shows that association takes place from left to right: with three consonants to four positions, the last consonant is linked to the surplus position (*ktabab*). (10b) poses a problem in this respect: either association is weirdly edge-in, or some extra mechanism is neeeded to tackle this case. McCarthy (1979 : 256) uses brute force here: he supposes the expected **katbab* in the first round with a later rule delinking the first linkage of [b] (*kat*•*ab*, where • represents the C slot from which the melody of the [b] was delinked), which is automatically followed by the spreading of the [t] (*kattab*), much like an instance of compensatory lengthening (see §1.2).

A slightly less powerful solution is proposed by Lowenstamm & Kaye (1985 : 117–118): they claim that association to the first position is inhibited from the start, thus each consonant of the root occupies its final position in the first round, as shown in (11a). The resulting configuration (empty C followed by filled C) is interpreted as a geminate, as in (11b). I have adapted the original to the previous diagrams of this chapter to ease comparison. We will see below (§3) that Lowenstamm & Kaye use a significantly different scheme.

(11) *The mapping of a geminate (*kattab*)*

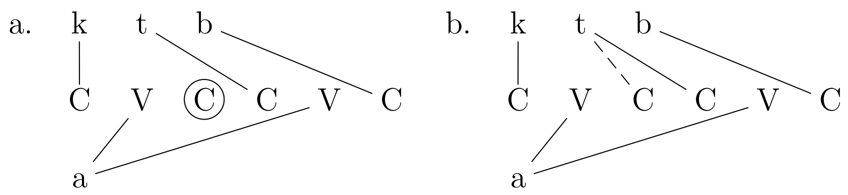

Note that McCarthy's second-round-spreading solution cannot apply after Lowenstamm & Kaye's first-round blocking, since that would yield the unattested form **kaktab*.[2]

In McCarthy's analysis, the CV skeleton of Arabic words is a morpheme (a prosodic template in his words), identifying the binyan of the word form, contributing to the semantic elements of the specific binyan (as if the Attic–Aeolic difference between [eːmi] and [emːi] represented a difference in morphological categories).

Clements & Keyser apply the CV skeleton as a universal phonological device, the mediator between the syllable and autosegments, its two types of members, C and V representing "the useful but

[2] Neither analysis gives a reason for delinking or inhibiting the association of the consonant encircled in (11a), so that the unattested form **katbab* is avoided. Following Hoberman (1988) we may assume that long-distance geminates (those separated by a vowel) are more marked (their inhibition is ranked higher) than local ones and that word-initial geminates are even more marked. This explains why *kattab* is preferred to **katbab,* but *ktabab* to **kkatab*.

ill-defined notion of 'phonological segment'" (1983 : 11). The C for them is an anchor for anything [−syllabic] and the V for [+syllabic] segments. Prince (1984) shows that such impoverished representations adequately capture the templates of, for example, verbal person endings in Finnish: they are -C in the singular and -CC*e* in the plural, with the melody [m] in the first, and [t] in the second person. The surface forms are thus 1-sg *-n* (by an independently motivated rule turning *-m* to *-n* word finally), 2-sg *-t,* 1-pl *-mme,* and 2-pl *-tte.*

3 The X skeleton

Simultaneously with the development of theories of the CV skeleton there evolved an alternative view that considered the distinction of C and V slots redundant, and argued that skeletal slots are uniform, usually marked with dots or x's (e.g., Lowenstamm & Kaye 1985, Levin 1985). Proponents of the X skeleton have put forward a number of arguments against skeletal positions predestined for syllabicity.

3.1 Reduplication in Mokilese

Levin (1985 : 35–41) shows some peculiar cases of reduplication from Mokilese, which, she believes, are analysable only with an X skeleton. The point is that the reduplicant is a copy of the first three segments of the first syllable of the stem, irrespective of their being consonants or vowels. So, argues Levin, the template of the reduplicant must also lack this information. The relevant data are given in (12).

(12) *Mokilese reduplication*

	stem	progressive	gloss for stem
a.	pɔdok	pɔdpɔdok	‘plant’
b.	kasɔ	kaskasɔ	‘throw’
c.	pa	paːpa	‘weave’
d.	wia	wiːwia	‘do’
e.	caːk	caːcaːk	‘bend’
f.	onop	onnonop	‘prepare’
g.	andip	andandip	‘spit’

Levin contends that the reduplicant must be a totally specificationless chunk, $[_{\sigma}\times\times\times]$, to which the copy of the melody of the stem is associated following universal conventions. The case of (12a, b, g) is now straightforward. When the stem is too short, as in [pa], (12c), the last melody is multiply linked. The fact that the reduplicant is a single syllable inhibits the second vowels of [wia] and [onop] from associating to the skeleton, (12d, f), as a result, the preceding vowel or consonant is lengthened again. There is a problem with the stem [caːk] though, (12e). The melody of the stem comprises three bits ‘c’, ‘a’, and ‘k’, therefore the expected reduplicated form is *[cakcaːk], instead of the attested [caːcaːk]. Levin has to stipulate that multiple melodic associations, like that of the long [aː], are transferred in reduplication. A further problem of this analysis lies in the interpretation of the reduplicant: it is specified as a syllable, but it is not one in [on.n-onop] or [an.d-andip] (where the dot indicates the syllable boundary, the hyphen is between the reduplicant and the stem), since a word-internal prevocalic consonant forms a syllable with the following vowel, as the universal onset maximalization principle requires. Yet, the constraint on the reduplicant being a syllable cannot be relaxed, because if the first three segments were copied without reference to a syllable, undesired results like *[wi.a-wia] or *[o.no-onop] would emerge. In fact, Moravcsik says that in her survey of reduplication types she has never come across formulations like “reduplicate the first two [or, in our case, three — szp] segments (regardless of whether they are consonants or vowels)” (1978: 307–308). If in a language reduplication copies the first CVC part of the stem for consonant-initial stems, it will copy VC (not VCV) of vowel-initial stems.

Actually, a simpler account is available for the data in (12). Theoretically it is not more plausible than Levin's, but needs less stipulations, and thus invalidates her analysis *as an argument* for the X skeleton. Suppose, as in §1.3 above, that syllable onsets are always represented on the skeleton, either as a filled or as an empty C position. (This immediately explains Moravcsik's observation.) The reduplicant then is a copy of the first CVC part of the stem, melody and skeleton included. The cases of (12a–c) are obvious. The third slot for (12c) is automatically filled by the vowel of the reduplicant, just like for Levin. The objection that vowels cannot spread onto a consonantal slot (also made by Broselow (1995 : 184)) is mistaken: a C slot is not meant to host consonants exclusively, but nonsyllabic segments. If a syllable has one syllabic segment, then a long vowel is hosted by a VC sequence on the skeleton, as Clements & Keyser (1983 : 12) argue.

In (12d), the empty intervocalic C position is involved in the copying, but being preconsonantal in the reduplicant it serves as an anchor for the preceding vowel, unlike in the stem, where it is prevocalic. This is shown in (13a). Prevocalic stems blindly copy the initial empty C position, thus only the first two "real" segments form the reduplicant. (12g) seems to cause a problem now: here the reduplicant appears to be [and-], i.e., VCC, instead of the expected VC. Raimy (1999) suggests an obvious solution: if [nd] is analysed as [$^{\text{n}}$dd], a geminate prenasalized stop, then the situation is identical to that in (12f). The stem-initial empty C must be filled to satisfy onset maximalization: it is impossible to have a coda consonant followed by an empty onset. This is illustrated in (13b). (The reduplicant and the stem are enclosed in brackets for easier identification.)

(13) *Reduplication and empty onsets*

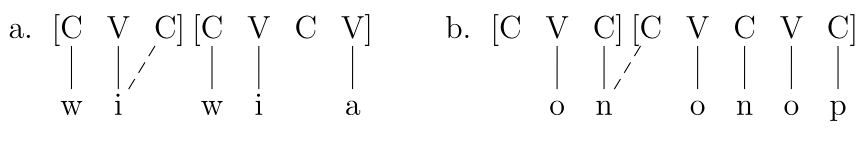

In fact, Levin herself suggests the empty-C-slot analysis as an escape hatch for the CV skeleton, but rejects the idea on the grounds that the vowel of the causative prefix [ka-] does not lengthen when prefixed to vowel-initial stems (e.g., [ka+adanki] > [kaːdanki] 'to name', [ka+uruːr] > [kauruːr] 'to be funny'). Vowels do not usually lengthen by filling a prevocalic empty C position (cf. Hayes 1989 : 281), what is more, it is hard to expect a long vowel or a diphthong to further lengthen. The conclusive test, the prefix [ak-], which is expected to geminate its consonant if prefixed to a vowel-initial stem if there was an empty consonantal slot, "was only found prefixed to C-initial stems" (Levin 1985 : 40). We can conclude that the hypothesis that vowel-initial stems carry an empty consonantal position at their left edge is not refuted by Levin's data.

3.2 Redundancy of C and V

A better argument against CV skeletons is that specifying syllabicity on the skeleton is redundant if the same information can be read off higher prosodic structures, like syllabic constituents, especially the nucleus. Lowenstamm & Kaye (1985) argue that simple syllable trees, like those in (14), adequately define the slots of the skeleton.

(14) *Syllable trees*

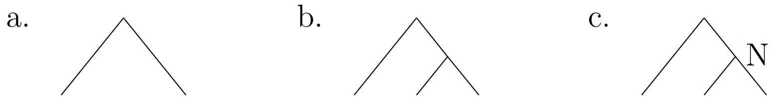

They suggest that labelling the trees is unnecessary since this information also follows from the configuration. Nevertheless, some minimal labelling is necessary to distinguish CVC, (14b), and CVV, (14c), syllables — think, for example, of the Arabic templates of binyans II (*kattab*) and III (*kaatab*), see (9).

Lowenstamm & Kaye (1985) raise the issue of whether the skeleton is an independent level in phonological representations, or merely a projection of higher prosodic structure, namely, syllable structure.

A consequence of this assumption is that the nodes representing syllabic constituents (like onset or nucleus) cannot be distinguished from the skeletal position(s) that they dominate. That is, it is impossible to conceive of skeletal positions not dominated by higher prosodic structure, or of a syllabic constituent that does not dominate a skeletal position.

Charette applies “pointless onsets” in an analysis of h-aspiré words in French (1991 : 90f). She claims that “normal” vowel-initial words begin with an onset that does not dominate any skeletal position, while those which contain h-aspiré — words that phonetically begin with a vowel, but phonologically behave as consonant-initial — begin with a regular “pointful” onset, dominating a skeletal position which is not associated with any melody. The vowel of the definite article is unpronounced before vowel-initial words, but it is pronounced before consonant- and h-aspiré-initial words. (15) illustrates the first part of the two cases using Charette’s examples: *l’amie* [lami] ‘the girlfriend’ and *la hache* [laaʃ] ‘the axe’.

(15) *Two types of empty onset*

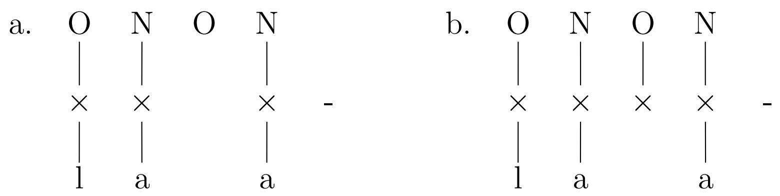

According to Charette’s analysis, the vowel of the article is deleted before a pointless onset as a result of the obligatory contour principle (see chapter 84 OCP), since the two nuclei are “adjacent” if the onset between them lack a skeletal slot, as in (15a). When such an onset is linked to a skeletal slot, it inhibits the deletion process, as in (15b). This analysis faces difficulties on several counts. On the one hand, the obligatory contour principle controls the appearance of identical *melodic elements* to adjacent skeletal positions. The nodes labelled nucleus do not qualify as such. On the other hand, liaison calls for the opposite representation of the two types of vowel-initial words. As mentioned in §1.3, some morphemes that are vowel-final preconsonantally exhibit a consonant when followed by a vowel-initial word. The plural of the definite article is an example: *les amies* [lez ami] 'the girlfriends' vs. *les haches* [le aʃ] 'the axes' (recall, h-aspiré-initial words behave like consonant-initial ones). Now the final [z] of the cliticized article is pronounced when there is no skeletal position for it to anchor to, and it is not pronounced when there is one, i.e., without further stipulations Charette predicts just the opposite of the attested liaison facts. The impoverished structures of (14) are also impossible if labels like "onset" and "nucleus" are treated separately from what they label: the skeletal slots.

To summarize, there is no compelling reason to distinguish skeletal points and the syllabic constituents containing them. Allowing pointless constituents or constituentless skeletal points makes unnecessary contrasts possible. But then, if prosodic nodes like onset and nucleus are not distinct from skeletal slots, then skeletal slots do carry the basic information of syllabicness: such a skeleton does contain Cs and Vs, irrespective of whether this is pencilled on paper as Cs and Vs, Os and Ns, or something else. The two levels must, nevertheless, be kept distinct if more than one skeletal slots can be associated with a single syllabic constituent, i.e., if branching onsets and nuclei are posited. §5 discusses a model, where even these are claimed not to exist.

4 Moras

As we have seen in the case of Mokilese reduplication (§3), preconsonantal empty C positions are available as targets for the spreading of a preceding vowel, intervocalic ones are not. In many languages a similar asymmetry characterizes these two consonantal positions.

Stress calculation, for example, may treat a preconsonantal consonant on a par with vowels, but prevocalic consonants are never so.[3]

Hock (1986) argues that the notion of mora must be (re)introduced into phonological theory. The mora, like the syllable, has been around in linguistic discussions from time immemorial, it is its theoretical status that is at stake here. Hock's proposal is to introduce the mora as an autosegment, rather similar to tones: in fact in his proposal tones are linked to moras. If compensatory lengthening could only lengthen a vowel in compensation for the loss of a tautosyllabic consonant, the "standard" CV or X skeleton would be fully capable of dealing with the process. We have seen, however, that compensatory lengthening also occurs at a larger distance: the loss of a vowel in the following syllable may lead to it, across an intervening onset consonant. Some of the relevant cases are: Greek glide loss (e.g., Proto-Greek [odwos] > Ionic [oːdos] 'threshold'; Steriade 1982 : 118) or Middle English schwa apocope (e.g., [talə] > [taːl] 'tale'; Minkova 1982).[4] In both cases the melody delinked and the vowel spreading is separated by a consonant that apparently remains linked to the skeleton.

(16) *Problematic cases of compensatory lengthening*

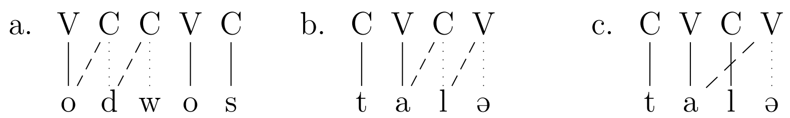

[3] It is common at this point to make a disclaimer to Everett & Everett 1984 (who claim that Pirahã is different in this respect) or to Davis 1988 (who collects cases where the quality of the onset seems to play a role in stress assignment). However, as Hayes correctly states: "I believe that the ability of moraic theory to account for wide-spread patterns of markedness should be given more weight in assessing the evidence than any particular awkwardness in the analysis of individual languages" (1989 : 303). This is probably true for any theory. Furthermore, some of the very few onset-sensitive systems were shown to be reanalysable so that they are not onset-sensitive (Goedemans 1996, Takahashi 1999).

[4] Despite its being debated, I keep Minkova's example because this is what features in the literature.

Actually, as (16a, b) show, the consonant standing in the way of compensatory lengthening is shifted to the right by one slot in both cases. This process was proposed by Steriade (1982 : 126–128) and named double flop by Hayes (1989 : 265–267). The Greek case — depicted in (16a) — can be explained by universal principles: the loss of [w] leaves us with an empty onset (provided that the syllabification is [od.wos]). The resulting [od.os] violates the onset maximalization principle, thus resyllabification ensues. But the skeletal position does not resyllabify, since there is an empty onset slot, recently vacated by [w]. It is to this slot that the [d] associates, leaving its original slot empty, triggering the lengthening of the preceding vowel.

The lengthening triggered by apocope, exemplified by the Middle English [talə] > [taːl] in (16b), is more problematic for a theory which lacks moras. The mechanism appears to be the same as in (16a), but now the consonant before the disappearing word-final vowel is supposed to flop to a vocalic position, to the nuclear slot of the last syllable. In addition, the position it leaves is not one that should cause lengthening of the preceding vowel. The alternative, whereby the vowel spreads out immediately to the vacated vocalic slot, as in (16c), is even worse, as it violates the axiomatic constraint inhibiting the crossing of association lines.[5]

In fact, with both CV and X skeletons it is hard to explain why the spreading of a vowel to some consonantal slots should cause lengthening, while in other cases an apparently similar vowel spreading does not. For example, the empty onset in Hungarian *pia* [pijɑ] 'drink' is filled by the spreading of the melody of the preceding vowel,

[5] This problem could be avoided by placing vowels and consonants on separate autosegmental planes (as in (10) and (11)), however, such a modification would loosen the theory beyond desirable limits: we would now find it hard to explain why so many processes deemed possible by the framework do not ever occur.

like in (8).[6] Yet the result is not a long vowel, which it is in *film* 'id.' for which the pronunciation [fiːm] is possible (Siptár & Törkenczy 2000 : 281). (Cf. Hayes 1989 : 281–283.)

Hock's (1986) proposal is to attach a mora (μ) to each weight-bearing position, that is, to each vocalic position, as well as to some consonantal positions, notably codas. The two cases now look as in (17).

(17) *Double flop with moras*

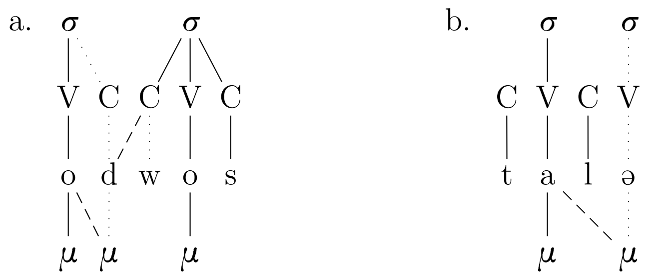

The moraic analysis of [odwos] > [oːdos] in (17a) is not significantly different form the moraless one, shown in (16a). It nevertheless suggests a reason for the asymmetry between onset and coda consonants: the former do not possess a mora, the latter do. The advantage of the mora analysis becomes clear in the lengthening of a vowel caused by apocope: [talə] > [taːl], (17b). The intervening onset consonant is not affected by the process at all, since it is not associated with a mora. Thus, the mora left floating after the final vowel is lost can associate to the stem internal vowel "above the head" (or rather "below the feet") of the intervening moraless consonant, much like in a vowel harmony process, where intervening consonants not possessing the relevant vocalic feature are transparent.

Hayes (1989) rearranges the relationship of the syllable and the mora by making the latter an integral part of prosodic structure,

[6] While it may be argued that *pia* is underlyingly [pijɑ], the question still holds why the same structure, the melody of [i] doubly linked to a V and a C slot, is [ij] in one case and [iː] in the other.

dominated by the syllable node. In a more radical innovation he also gets rid of the skeleton as it was conceived before. In his view the function of the skeleton is taken over by moras, moraless consonants are associated either directly with the syllable node, or they share a mora with the moraic segment. Accordingly, the two processes displayed in (16) and (17) would look as in (18).

(18) *Moras as the skeleton*

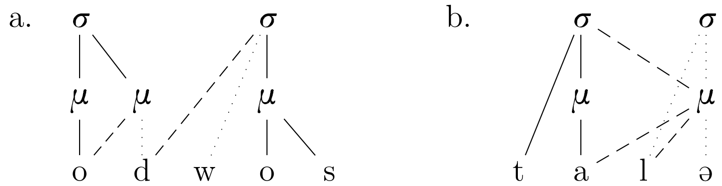

The simple double-flop case of Greek glide deletion in (18a) does not deserve much comment, the mechanism is again the same as before. For Middle English apocope, however, Hayes needs an extra stipulation called parasitic delinking: the loss of an overt nucleus in a syllable entails the dissolution of the whole syllable in his view. What is now left of the last syllable is joined to the first one, yielding the correct result. In Hock's analysis, on the other hand, the [l] remains in place, it does not have to be delinked and relinked, as can be seen in (17b).

Despite this complication, Hayes's model has definite advantages over Hock's use of moras. On the theoretical count, it is simpler in that it lacks the CV or X skeleton. On the empirical count, it predicts that compensatory lengthening of a vowel is only caused by the loss of a moraic segment that follows the vowel, never by the loss of one that precedes it. As (19a) shows, Hock's representations easily allow the latter case, which is not attested according to Hayes. His hypothetical example is [əla] > [laː].

(19) *Compensatory lengthening triggered by loss of preceding vowel*

In Hayes's model, (19b), the freed mora of the first syllable cannot be captured by the second mora, because the onset consonant inhibits this. The price to pay for this solution is the stipulative parasitic delinking mentioned above: if the moraic segment of a syllable is delinked, the onset consonant is also delinked, as in (18b). Without this an onset will always block the linking of a heterosyllabic mora. Note that in Hock's model not only the loss of a vowel, but also the loss of a moraic consonant could lead to the lengthening of a following vowel (e.g., Proto-Greek [esmi] > hypothetical *[emiː]). Such changes also seem to be unattested, as predicted by Hayes.

While theoretically attractive, dispensing with the skeleton has a serious repercussion. Recall that linking IPA symbols to elements of higher prosodic structure (slots of the skeleton, moras, syllables) is misleading since segments are not atomic. In partial trees like those in (20), where the Greek letters α–ε stand for (auto)segments, the temporal order of these autosegments is not specified. The string $\beta\gamma$ is usually referred to as a branching onset, δ is a moraless coda, which may occur word finally even in languages with moraic codas, like English.

(20) *Autosegmental representations without the timing tier*

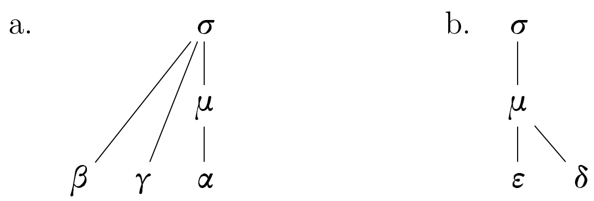

Accordingly, the order of two adjacent tautosyllabic or tautomoraic segments must be given by some stipulation. Kaye, for example, provides such a stipulation: "By universal convention the less sonorous of the two elements associated to the same point is produced first in the speech chain" (1985 : 289). It remains to be seen if this can be maintained. For syllable-initial consonants, (20a), this is exactly what the sonority sequencing principle dictates. In the domain of single segments, affricates follow this convention, but the existence of prenasalized stops casts some doubt on its validity. Apart from light diphthongs (like French [wa] in *trois* [trwa] 'three'), monomoraic rhymal sequences, like in (20b), obviously cannot be subject to this generalization, since they are invariably ordered in the opposite way, more sonorous (vowel) first, less sonorous (consonant) second. Be that as it may, without some similar (set of) principle(s) an autosegmental representation without a timing tier is uninterpretable.

To overcome this difficulty, one might wish to introduce root nodes, a notion familiar from frameworks organizing features into hierarchical structures, so-called feature geometries (Clements 1985, Sagey 1986, McCarthy 1988). The root node is the topmost node of such a hierarchy, containing all of the features making up the given segment, that is, the entirety of the segment. If the graphical order of root nodes specified their temporal order as well — which is assumed in the contour-segment model of affricates — then root node would be just another name for skeletal slot, that is, one would simply reintroduce the skeleton into the representation. The skeleton apparently is indispensable.

5 A return to the CV skeleton

The modern career of the mora was launched by the need to distinguish onset consonants from coda consonants. Only the latter are capable of contributing to the weight of a syllable, that is, of behaving like a vowel; onsets are not. Therefore, a mora is assigned to consonants in the rhyme, but not to those in the onset. Note, however, that the reasoning is circular: codas are equipped with a mora because we observe that they behave differently, and then refer to these moras to explain their difference. But we could just as well imagine an alternative world in which onsets were moraic and codas were not. There is no inherent property of coda consonants that predestines them to be moraic as opposed to others in the onset. To make things worse, as we will see it is not exactly true that onset consonants are never moraic, at least, that their loss never entails compensatory lengthening. It turns out to be an oversimplification to tie diverse phenomena like compensatory lengthening, stress assignment algorithms, the assignment of tone-bearing units, etc. to a single property of the representation, moras (Hayes 1995 : 299, Gordon 2004).

Incidentally, it is a version of the once rejected CV skeleton that might bring us closer to understanding this asymmetry in the behaviour of consonants at the two edges of the syllable. To distinguish it from the McCarthy and Clements & Keyser type of CV skeleton, I will refer to it by a widespread denomination, the strict CV skeleton. In §§1.3 and 3.1, we have seen why it is useful to suppose that some skeletal positions are empty. So far, we have only seen empty consonantal positions, but there is no particular reason why emptiness, that is, the state of not being associated to any melodic material, should be limited to consonantal positions. The claim that the host of the vowel (the nucleus) is the head of the syllable, therefore it cannot be missing, is not a very strong one. Syntactic heads, for example, the complementizer of a complementizer phrase, may remain empty (e.g., *I know* $[_{\text{CP}}\ [_{\text{C}}\ \emptyset\ ]$ *she'll come*]).[7] But other prosodic units like the foot may also exist without an overt head: in the previous sentence the first headed foot begins with *know,* the pronoun *I* before it is a headless, degenerate foot. Feet and syllables are similar types of prosodic units, headless syllables therefore are not the least unconceivable entities.

[7] In fact, in English it is by default empty in nonquestions, that is, there is an empty complementizer at the beginning of the matrix clause, too.

If nuclei may remain unpronounced[8] a very restricted syllable structure becomes available. Lowenstamm (1996) proposes that underlyingly all languages have the same skeleton, the simplest one available, comprising nonbranching onsets and nonbranching nuclei in strict alternation. Accordingly, no two consonants and no two vowels are adjacent on the skeleton, they are always separated by a position of the opposite type. (21) gives the four cluster types of (5) in the strict CV fashion.

(21) *The strict CV representation of vowel and consonant clusters*

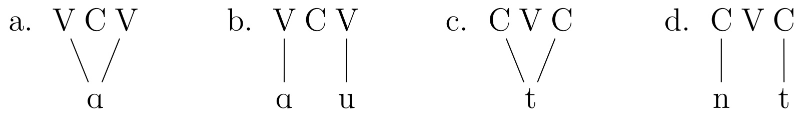

Recall the discussion in §3.2: the CV skeleton contains redundant information that can be read off higher prosodic structures. But this only holds *if* there is any higher prosodic structure. In fact, strict CV analyses generally do not call for the recognition of such structure, but certainly not of any further syllabic constituency.

The two well-known cases of compensatory lengthening—Proto-Greek [esmi] > Attic [eːmi] and Aeolic [emːi] — are illustrated in (22).

(22) *Compensatory lengthening in a strict CV skeleton*

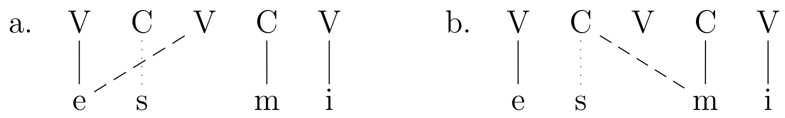

[8] Note that "empty" and "unpronounced" are not equivalent. In a privative feature framework, empty skeletal positions may be phonetically interpreted, as a sound maximally lacking any contrast, like e.g., [ə] or [ʔ]. Some empty skeletal positions may thus be pronounced, others may remain unpronounced if they satisfy certain conditions. See Kaye & al. 1985, 1990, Charette 1991, or Harris 1994 for details.

With the delinking of the [s], two skeletal slots are opened up for association: both the consonantal slot of the delinked coda and the vocalic slot enclosed within the original [sm] cluster. The choice is apparently controlled by a dialect-specific parameter, just like in any other theory of the skeleton.

The mora of moraic theories is an independent entity, which can be assigned to segments as the analyst needs it — it is only empirical considerations that stop them from assigning a mora to onsets. In the strict CV approach, moras are an inevitable consequence of the way the skeleton is built up (Scheer & Szigetvári 2005). A coda consonant is moraic, because it is followed by an unpronounced vocalic slot. That is, the moraicness of the coda is only apparent, it is the following vocalic slot that carries weight. In this view, it is exclusively vocalic slots that are moraic. The loss of an intervocalic consonant does not free any "buried" empty vocalic slot, as (23a) shows. The loss of a preconsonantal consonant, on the other hand, makes a so far unreachable vocalic slot available for spreading onto, as in (23b).

(23) *The loss of an intervocalic and a preconsonantal consonant*

The weight of closed syllables containing a short vowel is language specific. For example, in English and Cairene Arabic such syllables count as heavy, in Khalkha Mongolian and Yidiny they count as light (Zec 1995 : 89). This parametric variation is trivially encoded in moraic frameworks: coda consonants are now assigned a mora, now they aren't. In the strict CV model, the same fact is encoded by parameterizing whether an unpronounced vocalic slot is counted by the relevant process, or it is not. Crucially, however, since the shape of the skeleton is constant — it is always a strict alternation of vocalic and consonantal positions —, the uncounted vocalic slot is there even when it is not counted by a certain process (say, stress assignment). A prediction running counter to those of moraic theory follows from this fact: compensatory lengthening of a vowel should be possible even if coda consonants are not moraic in a language. Kavitskaya (2002) claims that at least two languages, Piro and Ngajan, are exactly like this. One could claim that the mora associated with the coda in such languages is one which does not contribute to weight, but does allow compensatory lengthening (as an anonymous reviewer points out). This then means that there are two types of mora, the "weight mora" and the "compensatory-lengthening mora." The strict CV model predicts exactly this: there are of two types of Vs. Pronounced Vs obligatorily contribute to weight, unpronounced ones are parameterizable.

In the strict CV framework, when an empty vocalic position enclosed between two consonants is "unearthed" compensatory lengthening may ensue, irrespective of whether this target of spreading is to the left or to the right of the vowel to lengthen. That is, the loss of an onset consonant may result in the lengthening of the vowel that followed it, as (24) shows.

(24) *Onset loss yielding compensatory lengthening*

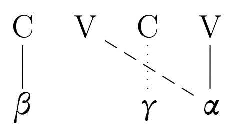

The theory dictates that this option is available only for postconsonantal onsets, not for intervocalic ones (see (23a)). Confirmation of this prediction comes from southwestern dialects of Finnish where gradated [k] is lost with compensatory lengthening. The data in (25) come from Kiparsky 2008, doubled vowels are long, as in standard Finnish orthography.

(25) *Compensatory lengthening in southwestern Finnish dialects*

input	SW dialect	standard	'gloss'
/jalka-t/	jalaat	jalat	'legs'
/nälkä-n/	nälään	näljän	'hunger-gen.'
/halko-t/	haloot	halvot	'logs'

In the Finnish data the lost consonant is always preceded by another consonant, it is never intervocalic. This is important because the empty vocalic slot is available *between* two consonants, but not after a vowel, as (23a) shows.

Samothraki Greek exhibits a similar type of compensatory lengthening. In this dialect, prevocalic [r] is lost, it is only retained in preconsonantal position — a mirror image of the distribution in the general nonrhotic dialects of English. The loss of postconsonantal [r] is illustrated in (26a). Intervocalic [r] is lost without trace, as in (26b), just as expected. (The data are from Topintzi (2006), who attributes them to Katsanis (1996)).

(26) *Loss of* r *in Samothraki Greek, 1*

	input	output	gloss
a.	/prótos/	[póːtus]	'first'
	/fréna/	[féːna]	'brakes'
	/xróma/	[xóːma]	'colour'
	/ɣráfo/	[ɣáːfu]	'I write'
b.	/léftirus/	[léftius]	'free'
	/varéʎ/	[vaéʎ]	'barrel'
	/méra/	[mía]	'day'
	/skára/	[skáa]	'grill'

To provide the missing mora, Hayes has to hypothesize an epenthesis stage before the loss of the [r]: [fréna] > [feréna] > [feéna] > [féːna] (1989 : 283). The strict CV analysis is rather similar, the only difference is a very important one though: the slot of the "epenthetic" vowel is lexically available, since any two consonants are always separated by such an empty slot. The relevance of this difference between the two analyses is that there is no empirical evidence for epenthesis in this case, furthermore this assumption creates paradox in the ordering of the historical events (Kavitskaya 2002 : 98), thus Hayes's hypothesis is not plausible. The strict CV skeleton, however, has the vocalic position where the vowel can spread without any extra process.

But even the strict CV model seems to be taken by surprise when it comes to the loss of word-initial [r]: this loss also triggers compensatory lengthening, as the words in (27) show.

(27) *Loss of* r *in Samothraki Greek, 2*

input	output	gloss
/rúxa/	[úːxa]	'clothes'
/réma/	[éːma]	'stream'

Scheer & Ségéral (2001) introduce the notion of coda mirror. Coda is a typical lenition environment, it is the position in the word that is not followed by a vowel, that is, preconsonantal and word-final position. Coda mirror is the opposite case: it is the position not preceded by a vowel, that is, postconsonantal and word-initial position, which is claimed to be the strong position, where lenition is not likely. Scheer & Ségéral's theory is built on the strict CV skeleton, for them "not followed by a vowel" means followed by an unpronounced vowel, and "not preceded by a vowel" means preceded by an unpronounced vowel. It is this empty vocalic position that causes the lengthening of the vowel in the Finnish and the Greek data discussed here. It is not only postconsonantal, but also word-initial consonants that are assumed to be preceded by an empty vowel, a proposal first argued for by Lowenstamm (1999). Accordingly, the loss of a word-initial consonant may also cause compensatory lengthening, as shown in (28).

(28) *Word-initial consonant loss yielding compensatory lengthening*

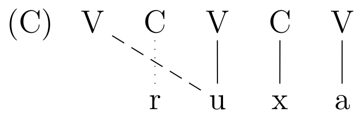

Since consonant loss is not common in the coda mirror position, compensatory lengthening is also rare in this environment. The peculiarity of Samothraki Greek then is that it exhibits [r] loss in the coda mirror position and not in the coda position. The ensuing compensatory lengthening is a consequence predicted by the strict CV skeleton.

6 Conclusion

The phonological skeleton evolved as a result of the autosegmental idea taken to its logical conclusion: segments, after having all their melodic content autosegmentalized, leave behind “traces” that encode their relative temporal order. The debates concerning the phonological skeleton are (i) whether skeletal slots specify any phonetic property (consonantalness vs. vocalicness) or none, that is, whether the skeleton contains Cs and Vs or uniform Xs, and (ii) whether the mora can replace skeletal slots, with moraless consonants linked directly to the syllable node. This chapter has argued that skeletal slots are Cs and Vs, not merely Xs, but there is no further prosodic constituency (like onsets, nuclei, or syllables). Furthermore, it has been claimed that the mora is not an independent element of the representation, but a consequence of parametrical settings on vocalic skeletal slots: pronouced V slots are universally moraic, unpronounced ones are moraic in some, but not in other languages. Consonants, on the other hand, are never moraic.