Phonology **14** (1997) 315–370. Printed in the United Kingdom

*Licensing Inheritance: an integrated theory of neutralisation**

John Harris
University College London

1 Introduction

If we compare the contrastive potential of different phonological contexts in any language, it usually does not take long to establish that the distributional spoils are unevenly divided. Each context typically displays its own subsystem of oppositions which may be bigger or smaller than those associated with other contexts (cf. Twaddell 1935). The traditional term NEUTRALISATION describes the relation between a defective subsystem and one that is distributionally better endowed.

The failure of a position to sustain a particular contrast can manifest itself in one of two ways, as Trubetzkoy was among the first to point out (1939: 209ff). Under ASSIMILATIVE NEUTRALISATION, the phonetic interpretation of the position with respect to the relevant contrast is determined by the melodic content of an adjacent position. This type of pattern is evident in vowel harmony, where the quality of a harmonising vowel is wholly or partially dependent on that of the dominant vowel within the domain. It is also to be seen in the assimilative suspension of consonantal contrasts. For example, in coda–onset interludes consisting of full or partial geminates, the phonetic interpretation of one position is wholly or partially dependent on that of the other.

REDUCTIVE NEUTRALISATION, on the other hand, refers to a situation in which restrictions on the melodic content of a position operate independently of contrasts in neighbouring positions. In vowel systems, for example, it is quite usual to find that the maximal inventory of oppositions is restricted to prosodically prominent nuclei, while shrunken subsystems of various shapes and sizes show up in weak positions. In its most extreme form, syncope, this results in a nuclear position being gutted of all melodic content. In the case of non-nuclear positions, contrastive potential can be curtailed by sonority sequencing constraints and by consonantal lenition processes which neutralise distinctions of manner (as in vocalisation and spirantisation) or place (as in debuccalisation).

Despite the apparent diversity of the contexts and segmental effects just reviewed, the fact that they all produce contractions in contrastive systems indicates that we are dealing with a unitary phenomenon. While the notion of system reduction helps sharpen the definition of what neutralisation is, it takes us only part of the way towards an understanding of why it occurs in the first place. For example, it offers little or no insight into the question of why certain phonological sites systematically favour neutralisation more than others – surely one of the cardinal issues confronting phonological theory. Most phonologists would probably agree that some underlying dimension of prosodic recessiveness is involved: that is, neutralisation targets weak positions within metrical or harmonic domains. However, the goal of providing a formal unification of the relevant contexts has proved somewhat elusive.

Any account of neutralisation, if it is to be more than just a catalogue of phonological contexts and the segmental effects that occur there, should ideally integrate the three criteria in (1).

(1) *Desiderata for an integrated theory of phonological neutralisation*
  a. unify the set of neutralising contexts,
  b. unify the set of segmental effects that occur in (a), and
  c. supply some necessary link between (a) and (b).

Moreover, in line with the general move towards viewing phonological derivation as consisting in the parallel operation of output constraints rather than in the serial application of rewrite rules, we might impose the following additional requirement: the statements which make up an account of neutralisation should be expressible over phonological output. The main purpose of this article is to demonstrate that all of these goals are achievable.

In recent years, the bulk of work on neutralisation has focused on its impact on consonantal systems (see the references below). The first part of this article is specifically concerned with challenging certain traditionally held assumptions about how this aspect of neutralisation should be accounted for. Nevertheless we should not lose sight of the fact that any general theory of neutralisation must also encompass the recurrent patterns of contraction that affect vowel systems. This issue is taken up in the last part of the article.

The sites which individually or collectively promote a reduction in consonantal contrasts are traditionally described as (i) word-final, (ii) intervocalic and (iii) preconsonantal. According to a by now classic approach based on syllabic constituency, all three contexts can be subsumed under the coda. Underlying this view are the following assumptions: (i) a word-final consonant is syllabified in a coda; (ii) the consonant of a VCV sequence can be captured into the coda of the first syllable; and (iii) the preconsonantal context always corresponds to an internal coda–onset interlude. Below we will consider a number of reasons for rejecting this overall approach: it fails to account for the consistently extrarhymal behaviour of final consonants; in subverting core syllabification, it leads to an unwarranted loosening of syllable theory; and it ignores the fact that a significant proportion of alleged interludes contravene otherwise robust syllable-contact restrictions.

I will present an alternative theory of neutralisation which bypasses constituency and goes straight to the heart of the grammatical function that subserves it, namely phonological licensing (see McCarthy 1979, Itô 1986, Goldsmith 1989, 1990, Kaye *et al*. 1990, Harris 1994, Itô *et al*. 1995 and others). The licensing principle requires of each prosodic or melodic unit in a representation that it be bound in some way to some other unit in order to receive phonetic interpretation. By invoking licensing in its various forms, we are able to unify the set of neutralising contexts (goal (1a)) without resorting to resyllabification or compromising the extra-rhymal status of final consonants. Common to these sites is a configuration in which a position is a non-head at some level of prosodic structure.

The unification of the melodic effects accompanying neutralisation (goal (1b)) can be most straightforwardly achieved within a privative model of segmental structure that is entirely free of segmental redundancy. This allows the contrastive potential of a syllabic position to be directly reflected in the degree of melodic complexity it is able to license. It also allows dynamic processes of consonantal lenition and vowel reduction to be uniformly represented as the suppression of melodic material.

The goal of forging a link between the contexts and the segmental effects of neutralisation (1c) can be achieved by positing an intimate connection between the melodic and prosodic aspects of licensing. The fundamental notion to be developed here – LICENSING INHERITANCE – is that the ability of a syllabic position to license melodic material directly reflects its status within the prosodic hierarchy. *Ceteris paribus*, prosodic heads enjoy a greater degree of melodic licensing potential than non-heads. The asymmetry in the degree of licensing power invested in different positions, I will argue, percolates throughout the phonological hierarchy. Having non-head status at some level of prosodic structure compromises a position's ability to license melodic material. The segmental effects of neutralisation can then be considered to result from the withholding of licensing from particular melodic units under certain prosodic conditions.

§2 features representative data on consonantal neutralisation and sets out the main reasons for rejecting the view that the contexts in which it occurs can be subsumed under the coda. §3 reviews the mechanism of phonological licensing and introduces the principle of Licensing Inheritance. §4 outlines a privative model of segmental structure that permits neutralisation to be uniformly expressed as a diminution in melodic complexity. §5 shows how Licensing Inheritance accounts not only for the lenition effects that instantiate consonantal neutralisation but also for sonority-sequencing restrictions on consonant clusters. §6 extends the treatment to vocalic reduction. §7 summarises the main conclusions.

2 Coda analyses of consonantal neutralisation

2.1 Braces and codas

One of the earliest motivations for the habilitation of syllable structure into phonological representation stemmed from a dissatisfaction with the over-generating properties of brace notation in linear rewrite rules. Braces, employed as a formal means of conflating different environments within the same rule, suffered from a failure to evaluate a small set of recurrent combinations any more highly than an excessively large set of unattested combinations. The most frequently observed combination, the familiar conjunction {C, #} (consonant or word boundary), is widely regarded as a prime site for consonantal neutralisation. It is now generally acknowledged to be a cryptic characterisation of a context more perspicuously identified in terms of syllable structure (see, for example, the arguments in Vennemann 1972). The assumption that soon gained ground, largely as a result of work by Kahn (1976), James Harris (1983) and others, was that the relevant syllabic context could be identified as the coda. (Adopting a widely held view, I assume that CODA does not label an independent constituent node but is simply an informal term for a postnuclear rhymal position.) This view continues to be widely held and has achieved the status of something approaching textbook orthodoxy (see, for instance, Roca 1994: 134ff, Spencer 1996: 174ff).

In what follows, I will use the term PHONOLOGICAL EVENT as a general descriptive label for any type of segmental regularity, irrespective of whether it involves dynamic alternations or static distributional patterns or both. (This is in preference to theoretically loaded terms such as RULE or CONSTRAINT or the procedurally loaded term PROCESS.) The examples in (2) and (3) illustrate some of the best-known cases of neutralisation occurring in what was originally formulated as the {C, #} context but which have subsequently been reinterpreted in coda terms. The examples in (2) are drawn from James Harris's (1983) work on Spanish. Those in (3) crop up in many coda-based accounts (see for example Kahn 1976 and Halle & Idsardi 1997 on (3b)).

(2) *Spanish*

a. s-*debuccalisation*: [s] → [h]

costa	co[h]ta	'cost'
después	de[h]pué[h]	'afterwards'

b. *Liquid gliding* (*Cibaeño Spanish*): [r l] → [y]

revolver	revo[y]ve[y]	'to turn over'
papel	pape[y]	'paper'
algo	a[y]go	'something'
carta	ca[y]ta	'document, card'

c. *Lateral depalatalisation*: [ʎ] → [l]

be[ʎ]o	'beautiful'	be[l]dad	'beauty'
donce[ʎ]a	'lass'	donce[l]	'lad'

d. *Nasal depalatalisation*: [ɲ] → [n]
re[ɲ]ir 'to quarrel' re[n]cilla 'quarrel (N)'
desde[ɲ]ar 'to disdain' desdé[n] 'disdain (N)'

(3) a. l-*vocalisation* (*Brazilian Portuguese*): [l] → [w]

sal	sa[w]	'salt (N)'	sa[l]eiro	'salt cellar'
salgar	sa[w]gar	'to salt'		
papel	pape[w]	'paper'	pape[l]ão	'cardboard'
falta	fa[w]ta	'lack'		

b. r-*vocalisation/loss* (*prototype non-rhotic English*)

carry	ca[r]y	ca~~r~~	ca~~r~~d
rain	[r]ain	fea~~r~~	boa~~r~~d

Additional examples of languages displaying more or less the same events include: *l*-vocalisation to [y] (Austrian German: Rennison 1981); *l*-vocalisation to [w] (Serbo-Croat: Kenstowicz 1994: 90, London English: Wells 1982: 313ff); *r*-vocalisation (German, Danish). Other events occurring in the same context will be discussed below.

Some of the events in (2) and (3), such as vocalisation and *s*-debuccalisation, are traditionally described as weakenings or lenitions. That is, they involve historical developments of a type which, if allowed to proceed unchecked, can eventually culminate in segment deletion (Hyman 1975: 165). Below I will show how all of them, and indeed all neutralisation effects, can be formally expressed as reductions in melodic complexity.

Once the coda had become identified as a favourable environment for weakening, it was but a short step for researchers to reanalyse cases of consonantal neutralisation occurring in other contexts along the same lines.

Besides the preconsonantal and word-final contexts exemplified in (2) and (3), a third favourable site for neutralisation is formed by a consonant flanked by vowels. The core syllabification of such sequences, it is generally agreed, must satisfy the requirement that onsets be maximised. However, according to one view, this basic configuration can be subverted in order to accommodate neutralisation. The context can be unified with others in which neutralisation occurs by moving the C of a core V.CV sequence into the coda of the first syllable. This device has been most widely appealed to in cases where the relevant string forms a trochaic foot:

(4) *Coda capture*
V.CV → VC.V/foot-internal

Coda capture may be achieved with or without severing the consonant's affiliation to the onset of the second syllable. (The ambisyllabic alternative is favoured by, for example, Kahn 1976; for the opposing view see, for instance, Selkirk 1982.)

Examples of foot-internal phenomena in English which have been treated in terms of coda capture include those given in (5a) (Kahn 1976, Selkirk 1982) and (5b) (Borowsky 1986).

(5)

	Foot-internal		*Foot-initial*
a.	*Tapping*: [t] → [ɾ]		
	pity	pi[ɾ]y	re[t]áin
	metre	me[ɾ]re	bou[t]íque
b.	*Defective* h		
	ve~~h~~icle		ve[h]ícular
	pro~~h~~ibítion		pro[h]íbit

(For further English examples analysed in these terms, see Gussenhoven 1986 and Wells 1990.) Tapping illustrates the typically neutralising nature of segmental effects in this context: the event results in a collapse of the distinction between [t] and [d].

Other languages displaying foot-internal intervocalic neutralisation of a type that is amenable to a coda-capturing analysis include Danish and Ibibio. As shown in (6a), original short plain plosives in Danish (still recorded in the orthography and evident in cognate languages such as Norwegian) have been subject to vocalisation and/or deletion in this context. ([ð] here indicates a dental approximant. The glide reflex of historical *g* takes on the quality of the preceding vowel.)

(6) *Danish*

a. *Foot-internal*			b. *Foot-initial*		
peber	pé[w]er	'pepper'	bebude	be[p]úde	'to foretell'
modig	mó[ð]ig	'brave'	bedyre	be[t]ýre	'to proclaim'
koge	kó[(w)]e	'to cook'	igen	i[k]én	'again'

In Ibibio (Benue-Congo), [p t k] weaken to [β r ɣ] respectively in the same context (Urua 1990); see (7a).[1]

(7) *Ibibio*

a. *Foot-internal*

[diβe]	'hide oneself'	cf.	[dip]	'hide'
[were]	'be written'	cf.	[wet]	'write'
[fʌɣɔ]	'cover oneself'	cf.	[fʌk]	'cover'

b. *Foot-initial*

[utaŋ]	(*[uraŋ])	'plaiting'
[ukʌp]	(*[uɣʌp])	'covering'

In both languages, foot structure can be independently identified on prosodic grounds. In Danish, it manifests itself as a stress trochee (see (6a)). In Ibibio, it consists in a trochaic frame which defines a template for verbal morphology and a domain for certain phonological events (Akinlabi & Urua 1992). Proof of the fact that lenition is indeed sensitive to foot structure in both cases is provided by the examples in (6b) and (7b). These show unlenited plosives in a context that is intervocalic but not foot-internal. (The laryngeal properties of the Danish plosives are discussed immediately below.) The failure of lenition in (7b) is due to the fact that the Ibibio nominalising prefix *u*- lies outside the verbal foot template. In both languages, the neutralising behaviour of the weak syllable of the foot is also evident in vowels: the second nucleus supports a smaller array of vocalic contrasts than the first.

Seeking to unify the intervocalic site with the preconsonantal/domain-final site appears to be on the right track, confirmed by the fact that neutralisation can affect both simultaneously. (The emergence of this possibility is inevitably dependent on the full range of relevant syllabic conditions being made available in a given language.) Cross-linguistic variation in the suspension of laryngeal contrasts in obstruents provides an illustration. A two-way laryngeal distinction is supported word-initially in, for example, Wolof (8a), Ibibio (9a) and Danish (10a). The contrast is neutralised in all three languages but in different combinations of contexts. (Danish *stød*, which is quite independent of the contrast in question, is omitted from the transcriptions in (10).)

(8) *Wolof*

a.	[pare]	'to be ready'	[ba]	'to abandon'
	[taal]	'to light'	[dof]	'to be crazy'
	[ceep]	'rice'	[ɟi]	'to plant'
	[kər]	'house'	[gan]	'guest'
b.	[up]	'to close'	[ubbi]	'to open'
	[təc]	'to close'	[tiɟɟi]	'to open'
	[dok]	'to die'	[doggali]	'to close a dead person's eyes'

(9) *Ibibio*

a.	[kpa]	'die'	[ba]	'exist'
	[ta]	'chew'	[da]	'stand
	[ke]	'at'		

b.	[diβe]	'hide oneself'	c.	[dip]	'hide'
	[were]	'be written'		[wet]	'write'
	[fʌɣɔ]	'cover oneself'		[fʌk]	'cover'

(10) *Danish*

a.	pil	[pʰ]il	'arrow'	bil	[p]il	'car'
	tale	[tʰ]ale	'to speak'	dale	[t]ale	'valleys'
	ko	[kʰ]o	'cow'	god	[k]od	'good'
b.	næppe	næ[p]e	'hardly'	ebbe	e[p]e	'low tide'
	sætte	sæ[t]e	'put'	bredde	bre[t]e	'width'
				sidde	si[ð]e	'to sit'
	frakke	fra[k]e	'coat'	lægge	læ[k]e	'to lay'
c.	lap	la[p]	'patch'	lab	la[p]	'paw'
	sæt	sæ[t]	'set (IMPER)'	mad	ma[ð]	'food'
	læk	læ[k]	'leak (N)'	lag	la[ɣ]	'layer'
d.	faktor	fa[kt]or	'factor'	fagter	fa[kt]er	'gestures
	scepter	sce[pt]er	'sceptre'	gabte	ga[pt]e	'yawned'

In Wolof, laryngeal neutralisation only occurs word-finally (Ka 1994). In (8b), the medial geminates in the forms on the right show the lexically voiced identity of the stops which undergo final devoicing. (The coronal series is neutralised under a tap finally; compare [teer] 'to arrive' with [teddi] 'to depart'.)

In Ibibio, the laryngeal contrast in oral stops collapses both foot-internally and word-finally (Urua 1990). In the former context, the neutralised series falls prey to the lenition already illustrated in (7) (see (9b)). Word-finally, it emerges as voiceless and unreleased (see (9c)). (The labial-velar stop is excluded from both positions.)

In Danish, the range of sites displaying laryngeal neutralisation is even more extensive. Danish plosives contrast on the basis of whether or not they are aspirated (spelt ⟨p t k⟩) or plain (spelt ⟨b d g⟩). Plain stops are generally voiceless; intervocalically, however, they are optionally and gradiently voiced, resulting from the interpolation of vocal-fold vibration from the flanking vowels. (For a full phonetic description, see Jessen 1997.) While the laryngeal opposition is maintained word- and foot-initially (see (10a)), it is suspended foot-internally (10b), word-finally (10c) and preconsonantally (10d). Evidence that a laryngeal contrast once existed in the foot-internal and word-final contexts is preserved in the orthography (see (10b) and (10c)). A vestige of it survives as a manner distinction between a plain stop and a vocalised reflex in forms such as *bre*[t]*e vs. si*[ð]*e* in (10b) and *sæ*[t] *vs. ma*[ð] in (10c). In other forms, it has been subject to merger under the plain stop series.

The representation of neutralisation raises two issues which are in principle quite distinct. One concerns the specific segmental consequences of individual neutralising events such as those just reviewed. The other centres on the more general question of why neutralisation occurs at all. Much of the relevant literature focuses on the first of these issues to the exclusion of the second. Any detailed treatment of neutralisation must of course get to grips with its segmental effects, and I will have certain proposals to make in this regard below. However, it is the second and more fundamental issue that will be spotlighted here.

Representing the segmental specifics of neutralisation has been a major preoccupation of phonological theory over the years. It is a goal common to such apparently diverse devices as *SPE* rewrite rules, feature co-occurrence filters (e.g. Selkirk 1982, Clements & Keyser 1983), coda conditions (e.g. Itô 1986, Prince & Smolensky 1993: 156ff), cluster conditions (e.g. Yip 1991), persistent delinking rules (e.g. Myers 1991) and feature licensing constraints (e.g. Itô *et al.* 1995, Lombardi 1995). I will argue below that such details are best expressed in terms of the ability of individual syllabic positions to license particular units of melody. Whatever the relative merits and demerits of these various formal devices might be, they are similar to the extent that they stipulate matters of segmental detail. This is probably a fair reflection of the degree to which grammars are offered genuine choices in the specifics of segmental form.

It is not even clear whether particular segmental effects are preferred in certain neutralisation-favouring contexts more than others. The suspension of laryngeal contrasts, for example, is classically associated with obstruent devoicing in word-final position, the pattern encountered in Wolof (as illustrated in (8)), Polish, Catalan, Dutch and many other languages (see Rubach 1990 and Lombardi 1995 for recent coda-oriented analyses). However, laryngeal neutralisation can also occur intervocalically within a foot, as demonstrated by the Ibibio and Danish examples in (9) and (10). Conversely, tapping is perhaps usually thought of as an intervocalic phenomenon. This is certainly true of English and Ibibio for example (see (5a) and (7a)). But there are also languages in which a tap occurs word-finally to the exclusion of coronal plosives. Examples include Lardil (Hale 1973: 426) and, as mentioned above, Wolof.

A fully elaborated taxonomy of neutralisation events may well reveal that particular segmental effects are preferred in particular contexts. However, any statement of the segmental details of neutralisation, irrespective of how it is formalised, in and of itself provides no answer to the question of why distributional restrictions hold unequally over different phonological contexts. It is this more fundamental issue that the Licensing Inheritance proposal is primarily designed to address.

To summarise the discussion so far, the classic contexts which individually or collectively promote consonantal neutralisation are traditionally identified as in (11).

(11) *Prime sites for consonantal neutralisation*
   a. word-final
   b. intervocalic within the foot
   c. preconsonantal

It is now widely taken for granted that these three contexts, apparently disparate when viewed in linear terms, are unifiable when viewed from the perspective of syllable constituency. One implementation of this insight is the assumption that each of the contexts identified in (11) is a coda. This claim is in turn dependent on the assumptions in (12) (listed in the order in which they will be discussed below).

(12) *The coda view*
   a. a word-final consonant is syllabified in a coda;
   b. the consonant of a VCV sequence can be captured into the coda of the first syllable;
   c. in any word-internal heterosyllabic C_1C_2 sequence, C_1 occupies a coda.

In the rest of this section, we will consider the validity of each of these claims.

Taken on its own, the general view expressed in (12) begs the question of why it should be the coda rather than any other constituent context that enjoys a privileged position with respect to neutralisation. This issue is independent of whether the view is implemented in a theory based on serial rule application or in one based on the parallel operation of output constraints. On its own, neither a rewrite rule nor an output constraint, even when specified in syllabic terms, supplies a necessary connection between codas and neutralisation. In the absence of some explanatory matrix, there is nothing in principle to prevent us from formulating a rule or constraint that would favour neutralisation in, say, onsets. This is empirically unfounded, at least as an exhaustive characterisation of the contexts in which such events occur. Any neutralisation that might be attested in onset position will either also be responsive to the location of the onset in foot or word structure or will also apply in other environments. Below I will attempt to provide an answer to the fundamental issue that is at stake here: why does neutralisation take place where it does?

A set of more specific problems surrounds the coda view. Recent work on syllable structure, including much which has no particular axe to grind on the issue of neutralisation, has yielded results which undermine all of the claims in (12). In what follows, we will examine each assumption in turn, reviewing evidence which shows that (12a) is untenable and that (12b) and (12c) are suspect.

2.2 The extrarhymal status of final consonants

Assumption (12a), that a word-final consonant occupies a coda, sits uneasily with the observation that this position systematically fails to display characteristics associated with codas which can uncontroversially be identified as occurring word-internally. Let us briefly note three respects in which this has been widely acknowledged to be true.

Firstly, for the purposes of stress assignment, a word-final consonant typically or consistently (depending on your view) fails to contribute to the quantity of the preceding rhyme to which it supposedly belongs. (Under one approach, this behaviour is recognised by allowing the consonant to be parametrically marked as extrametrical (see for example Hayes 1982). Below I will present arguments for considering it to be universally extrarhymal.) Secondly, the same quantitative independence is manifested in the failure of a final consonant to trigger closed-syllable shortening (see Myers 1987). Thirdly, in languages with final consonant clusters, the alleged coda clusters frequently contravene otherwise general sonority sequencing constraints (as noted by Levin 1985, among others). In all three of these respects, final consonants pattern with internal onsets rather than internal codas, a point we will return to presently.

There have been two quite different responses to this non-coda-like behaviour. One invokes a notion that can be broadly glossed as EXTRA-SYLLABIC (more generally, EXTRAPROSODIC). The other treats a final consonant as occupying the onset of a syllable lacking an audible nucleus. Let us briefly compare the two approaches.

The first response comes in two main variants, which differ on the formal status accorded extrasyllabicity and on the question of whether it persists in phonological output. Under a serial-derivational approach, the special status of final consonants is acknowledged by designating them as extrasyllabic during the early stages of derivation. Extrasyllabicity is a supplementary licensing mechanism which temporarily immunises domain-edge segments against stray erasure, the fate otherwise suffered by unprosodified material (McCarthy 1979, Steriade 1982). According to Itô (1986), extrasyllabicity holds obligatorily during the lexical phonology and optionally at word level. It is obligatorily disengaged at the postlexical level, where all segments must be prosodically licensed. At this point, a word-final consonant must either by syllabified (into a preceding rhyme or a following onset) or stray-erased.

Under a more recent optimality-theoretic version of this approach, extraprosodicity has no independent formal status but is instead an effect that derives from a specific interaction between two types of constraint (Prince & Smolensky 1993: ch. 4). One of these calls for the right edge of a syllable to be aligned with the right edge of a word. Violation of this constraint can be forced by some higher-ranked constraint (labelled $\mathbb{C}$ in (13)), which causes the end of a final syllable to be moved off the end of the word. (The exact nature of the weightier constraint is not crucial here.) This is illustrated in the tableau in (13), where, as a result of syllable–word misalignment, the final consonant in the optimal candidate form (13b) is unsyllabified (syllable edges represented here by parentheses, word edges by square brackets). The beaten candidate (13a), with perfect alignment, recapitulates the traditional analysis under which the final consonant is syllabified in coda position.

(13)

/…VC/	$\mathbb{C}$	Align
a. (VC)]	*!	
☞ b. (V)C]		*

There are several empirical reasons for rejecting both the final-coda view and extraprosodicity. One serious flaw in the assumption that a final consonant and an internal coda share the same syllabic affiliation arises from the prediction that any language possessing one will automatically also possess the other. This is demonstrably false. As shown by Kaye (1990), whether or not a language possesses internal codas is entirely independent of whether or not it sanctions domain-final consonants. (14) provides examples of the four-way typology that results from the intersection of these two binary options.

(14) VC *typology: internal coda* vs. *final* C

	Medial VC.?			
Final VC]?	*no*		*yes*	
no	-V.CV]	(Zulu)	-V(C).CV]	(Telugu)
yes	-V.CV(C)]	(Luo)	-V(C).CV(C)]	(German)

Invoking extraprosodicity allows the two contexts in question to be teased apart. Telugu, for example, might be characterised as a language which permits codas but eschews final extraprosodic consonants. However, extraprosodicity fails to alleviate another empirical problem which the final-coda view suffers from: neither approach is able to account for the phonotactic properties of word-final consonant clusters.

According to the extraprosodicity approach, the first in a final cluster of two consonants is syllabified in coda position, while the second is designated as extrasyllabic. As depicted in (15), the independence of an extraprosodic consonant (x_2) from the preceding coda (x_1) implies that the two positions should not display the kind of phonotactic restrictions that typically hold over adjacent pairs of fully prosodified consonants (for example, in interludes or complex onsets).

(15) $$\begin{array}{ll} \ldots x_1)_\sigma & x_2]_\omega \\ \quad | & \; | \\ \quad C_1 & C_2 \end{array}$$

This too is incorrect. There are very strict distributional dependencies operating in this context. And significantly they closely parallel those holding of internal interlude and/or internal branching-onset sequences, depending on the language. In English, for example, the dependencies are routinely of the coda–onset type. Compare the main distributional patterns that are evident in medial and final two-consonant clusters (examples from Harris 1994: 74):

(16) *English medial* -C.C- *and final* -CC] *clusters*

	medial	*final*		*medial*	*final*
stop–stop	chapter	apt	*sonorant–*	pamper	damp
	vector	sect	*stop*	winter	flint
				wrinkle	rink
				filter	guilt
				scalpel	scalp
fricative–	mister	mist	*sonorant–*	cancer	manse
stop	after	raft	*fricative*	dolphin	golf
	whisper	wisp			
	whisker	brisk			

How close you take the phonotactic fit between the medial and final contexts to be depends on your view of internal heterosyllabic consonant clusters. If it is assumed that these unfailingly constitute coda–onset interludes, then final clusters form a proper subset of those appearing medially. According to an alternative analysis, to be defended in §2.4, some of these alleged clusters are bogus and are in fact separated by an empty nucleus. Under this account, the phonotactic parallels between medial and final contexts in English are more or less complete. (The only significant difference has to do with [mb] and, in some dialects, [ŋg]. An independent development bars these clusters from appearing domain-finally; hence medial [mb] in *clamber* but not in *climb*.) This disagreement does not detract from the phonotactic affinity between final CC] and medial -C.C- clusters.

We might try to formulate the phonotactic restrictions on the final clusters illustrated in (16) in terms of an interaction between an extraprosodic position and a preceding coda. This would amount to treating extraprosodicity as a constituent node in its own right. But distributional statements couched in these terms would simply duplicate statements relating to internal coda–onset clusters. Nor is the problem resolved if the assumption is made that extraprosodicity shuts off at some point in derivation and the final consonant gets incorporated into the coda of the preceding rhyme. This syllabification creates novel coda clusters; that is, sequences exhibiting phonotactic dependencies which are not catered for by initially established constraints on internal codas.[2] In short, under this type of analysis, the parallel between the phonotactic conditions on internal -C.C- and final -CC] clusters is entirely accidental.

The problems of the final-coda and extraprosodicity approaches are further compounded when we turn to languages such as French and Polish. Here the challenge is that one set of final -CC] clusters displays the same distributional properties as internal -C.C- sequences (just as in the English examples in (16)), while the rest are distributionally identical to initial and internal and branching onsets. (For presentations and analyses of the relevant facts, see Dell 1976, Charette 1992 and Cyran & Gussmann in press.) On the onset parallel, compare the final and medial clusters in French examples such as *table* [tabl] 'table' – *tableau* [tablo] 'picture', *vitre* [vitr] 'window-pane' – *vitrine* [vitrin] 'shop-window'. Under an extraprosodicity analysis, the final liquids in forms such as [tabl] and [vitr] would be marked extraprosodic and the preceding consonants syllabified in coda position. The phonotactic dependency between the two positions is unaccounted for, as is their relation to internal branching onsets.

The simplest alternative is to assume that the distributional parallels between final and internal consonant clusters reflect parallel constituent structures. In other words, where a final -CC] cluster displays the same distributional properties as an internal -C.C- sequence, we assume that the second consonant occupies an onset position (as proposed by Kaye *et al*. 1990 for example). On the traditional assumption that there can be no such thing as a stray onset, we are then forced to conclude that the final consonant of a form such as *guil.t* must be licensed by a following nucleus, albeit one that is not phonetically expressed. (More on this point presently.) In this way, we succeed in unifying the statement of phonotactic restrictions on medial and final consonant clusters. In the case of forms such as *guilty* and *guilt*, the clusters in question are uniformly of the coda–onset type:

(17)
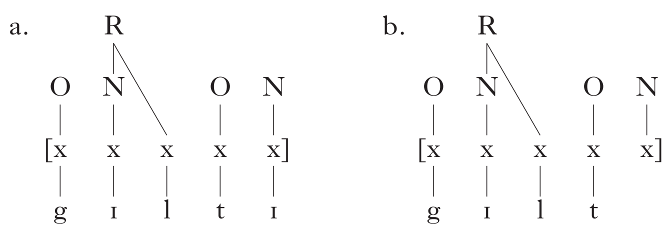

By the same token, any final two-consonant sequence which is distributionally identical to an initial or internal branching onset will be syllabified as such. French, for example, thus displays a branching [bl] onset in [tabl] no less than in [tablo] or *blé* [ble] 'wheat'.

Completing the parallel between medial and final syllabification, we may further assume that a singleton final consonant also occupies an onset. An onset [t] thus occurs not only in, for example, *city* but also in *sit*. This places us in a position to explain the other apparent peculiarities of final consonants outlined at the beginning of this section. We account for the fact that consonants in this position behave exactly like internal onsets for the purposes of calculating rhyme weight. The quantitative independence that a single final consonant displays in relation to a preceding rhyme follows directly from the fact that it never forms part of the rhyme; a rhyme preceding a final onset is just as free to support a light–heavy contrast as one preceding an internal onset. Moreover, the failure of final consonant clusters to respect otherwise general sonority sequencing constraints on codas follows trivially from the conclusion that they are not codas at all but either coda–onset or branching-onset clusters.

The combination of quantitative and phonotactic evidence just reviewed provides empirical support for the conclusion that a final consonant occupies an onset rather than a coda. Drawing the additional but independent conclusion that this position must be licensed by a following empty nucleus is initially motivated by theory-internal considerations. External confirmation that this move is on the right track comes from two main sources.

Firstly, there are good grounds for assuming that null-vowel syllables are parametrically metrifiable in the same way as weak-vowel syllables (cf. Giegerich 1985, Segundo 1990, Kiparsky 1991, Burzio 1994). In Spanish, for example, the typically final stress pattern of consonant-final words (e.g. *papél*) reduces to the penultimate pattern typically found in vowel-final words (e.g. *patáta*), if we assume that both types of form contain a final trochaic foot. The fact that the weak nucleus of the foot is sounded in *pa*(*táta*) (foot parenthesised) but not in *pa*(*pél*∅) is for metrical purposes irrelevant. In Polish, on the other hand, a final null-vowelled syllable does not project into metrical structure: words consistently show penultimate stress irrespectively of whether they are vowel-final or consonant-final. English accommodates both patterns. For example, amongst forms containing a metrified final empty nucleus, so-called superheavy-final stress reduces to penultimate stress: a final heavy–light foot is no less discernible in, say, *pre*(*vént*∅) than in *a*(*génda*) (Burzio 1994: ch. 3).

A second way in which final empty nuclei betray their presence is when some constraint forces their phonetic interpretation. It has been argued that the apparently epenthetic vowel separating alveolar obstruents of like manner in English suffixed forms with -*ed* or -(*e*)*s* is the phonetic expression of a domain-final nucleus which otherwise remains silent, for example in [[wɛd∅]d] *wedded* and [[mɪs∅]z] *misses* (Kaye 1990).

To conclude this part of the discussion: both the extraprosodicity and final-onset approaches acknowledge the non-coda-like behaviour of final consonants. Both accounts challenge claim (12a), one of the pillars supporting the coda view of consonantal neutralisation. Acceptance of extraprosodicity increases doubts about the claim. Acceptance of the final-onset account leads to its rejection.

2.3 Foot-internal consonants

Let us now turn our attention to (12b), the claim that an intervocalic consonant can be captured into the coda of the preceding syllable. Whether this is assumed to involve a severing of the consonant's link to the following syllable or not (thus giving rise to ambisyllabicity) is irrelevant to the following discussion. Resyllabification by coda capture is often invoked as a means of deriving one of the contexts in which consonantal neutralisation operates, such as the English tapping and defective-*h* examples in (5).

Applied to consonantal weakening, the term INTERVOCALIC is only a very loose characterisation of the context in question. There are indeed cases in which the vocalic nature of neighbouring segments can be construed as a significant conditioning factor, but this is usually insufficient to nail down the precise context. At least some aspect of the prosodic milieu is also typically involved. For example, the spirantisation of voiced stops in Spanish, often described as occurring intervocalically (e.g. Hyman 1975: 63), is sensitive to domain structure at the level of the phonological phrase or intonation group and to the presence of adjacent non-nuclear positions (James Harris 1969: 38–40). Essentially the same pattern is observed in Murut (Prentice 1971) and in the spirantisation of voiceless plosives in Tuscan Italian (Giannelli & Savoia 1980).

The influence of prosodic conditioning on intervocalic weakening is most evident in instances where VCV forms a trochaic foot, the context for which coda capture has been extensively invoked (e.g. Kahn 1976, Selkirk 1982, Borowsky 1986, Gussenhoven 1986, Wells 1990, Rubach 1996). There is of course no suggestion that every language will necessarily exhibit distributional or lenition asymmetries between different consonantal positions within a foot. Spanish spirantisation, for example, is completely blind to foot structure (compare, say, [sáβe] 'knows' with [saβér] 'to know'). However, where such imbalances do exist, it always seems to be the foot-internal position that gets a raw deal. Compare Spanish spirantisation with the corresponding event in Ibibio, for example, where intervocalic stops undergo lenition foot-internally (as in [ɲek-e] → [ɲeɣe] 'shake oneself' – cf. (7a)), but resist it when foot-initial (e.g. [[u][kʌp]] → [ukʌp] (*[uɣʌp]) 'cover (N)' – cf. (7b)) (Urua 1990). Thus, while there are languages such as English, Danish and Ibibio in which particular distinctions are maintained foot-initially but not foot-internally (see (5)–(7)), to the best of my knowledge there are none exhibiting precisely the opposite pattern.

Unlike the final-coda view rejected in §2.2, the main argument against coda capture and ambisyllabicity in particular and against resyllabification in general is not primarily an empirical one. It is rather of a general methodological nature: there is available to us a much simpler alternative. It has long been known, if not widely acknowledged, that coda capture can be dispensed with by referring to the relevant position's location within the foot (e.g. Kiparsky 1979, Harris & Kaye 1990, Jensen 1993). For example, rather than siting English *t*-tapping or *h*-deletion in a captured coda, we need simply refer to a foot-internal onset.

In what follows, I will use the term RESYLLABIFICATION to refer to any analysis in which syllable structure in phonological output is allowed to deviate from core syllabification. This usage is most readily associated with serial-derivational approaches employing transformational rules which alter initially established syllable structure. But it applies equally to any output-oriented analysis in which candidate forms with non-core syllabification are adjudged more optimal than those with core syllabification.

Under an approach which eschews resyllabification, the effect of onset maximisation cannot be overturned in output: a parse of VCV as VC.V is universally excluded. This restrictive view represents a closer approximation to the null hypothesis that basic syllable structure is immutable. Resyllabification, whether it be cast in rule-based or constraint-based terms, forms part of a research hypothesis which in the first instance has to be weighed up against the null hypothesis. As in any such enterprise, the preferred course of action is to abandon the null hypothesis only when the empirical balance tilts decisively in favour of an alternative.

The question then is this: can the more restrictive theory cope with data which has previously been treated in terms of resyllabification? More immediately, can it cope with the sort of data cited as motivating coda capture which, as expressed in (12b), constitutes another of the pillars supporting the coda account of neutralisation? The contention to be made in §5 is that it can.

2.4 Bogus consonant clusters

On the face of it, the third of the assumptions underpinning the coda account of consonantal neutralisation, that outlined in (12c), is the least controversial. This is the claim that any internal two-consonant cluster which does not form a complex onset, as in (18a), automatically constitutes a coda–onset interlude, as in (18b).

(18)
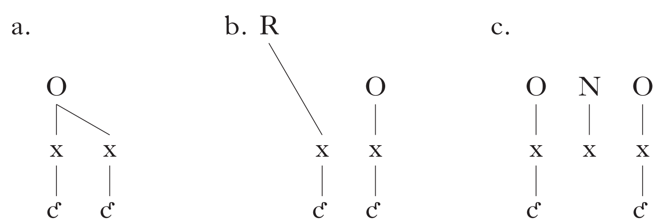

The main point of this subsection is to argue that the correctness of this assumption cannot be taken for granted. There is little doubt that at least some internal CC sequences must indeed by syllabified as in (18b). However, a significant proportion of heterosyllabic consonant sequences, it can be demonstrated, are not genuine clusters at all but should instead be treated as in (18c); that is, the members of such consonant pairs occupy independent onsets separated by a nucleus which does not necessarily have phonetic expression. The relevance of this point to the overall theme of §2 is that this type of sequence sometimes plays host to consonantal neutralisation events involving the recurrent {C, #} context.

The following discussion proceeds on the premise that the existence of systematic phonotactic dependencies between positions is proof that they are adjacent at some level of projection. By the same token, lack of phonotactic dependency can be taken as a sign of lack of adjacency. Recognition of the relevance of this principle to consonants in coda–onset sequences is inherent in such notions as sonority sequencing and syllable contact laws (e.g. Jespersen 1904, Selkirk 1982, Vennemann 1988, Clements 1990, Harris 1990, Kaye *et al*. 1990, Rice 1992). The phonotactic generalisation in this context is that the distributional latitude of the coda is narrower than that of the following onset.

Apparent counterexamples to this generalisation take the form of cases where heterosyllabic consonant sequences not only contravene otherwise general syllable contact laws but also show no systematic phonotactic interactions whatsoever. The cases typically involve a pair of consonants flanking a site which displays a vowel–zero alternation. The following Turkish example is one of a number of events occurring in the {C, #} context which Clements & Keyser (1983) analyse in coda terms. A vowel–zero alternation is observable in stems such as those in (19).

(19)		i. ACC	ii. NOM	ABL	
	a.	[vakt-i]	[vakit]	[vakit-ten]	'time'
		[koyn-u]	[koyun]	[koyun-dan]	'bosom'
		[karn-ɨ]	[karɨn]	[karɨn-dan]	'abdomen'
		[kɨsm-ɨ]	[kɨsɨm]	[kɨsɨm-dan]	'division'
		[azm-i]	[azim]	[azim-den]	'resolution'
		[hükm-ü]	[hüküm]	[hüküm-den]	'judgement'
	b.	[akl-ɨ]	[akɨl]	[akɨl-dan]	'intelligence'
		[kabr-i]	[kabir]	[kabir-den]	'tomb'
		[fikr-i]	[fikir]	[fikir-den]	'idea'
		[devr-i]	[devir]	[devir-den]	'transfer'

According to a classic type of resyllabification analysis, the lexical representation of such stems terminates in a CC cluster, the form they take in the accusative. The cluster is then assumed to be broken up by vowel epenthesis under specific conditions. (The quality of the vowel, which is in the main harmonically predictable, is not relevant to the point at hand.)

Expressed in linear terms, the conditions amount to the familiar {#, C} conjunction, the boundary referring to forms such as the nominative, the C to forms such as the ablative.

In Clements & Keyser's (1983) syllabic reanalysis, the first of two consonants in a final cluster is initially syllabified in a coda, while the second is designated as extraprosodic. The epenthesis of an intervening nucleus then creates a new syllable which attracts the first consonant into its onset and the formerly extraprosodic consonant into its coda, thus rescuing it from stray erasure. A derivation such as $dev.\langle r \rangle_{ep} \rightarrow dev.i\langle r \rangle_{ep} \rightarrow de.vir.$ illustrates the manner in which epenthesis feeds resyllabification.

For all its unfashionable serialism, this account incorporates certain fundamental assumptions about syllabification which have been carried over more or less unchanged into more recent optimalist treatments of epenthesis. Inkelas (1996), for example, proposes an analysis of Turkish in which epenthesis is achieved by ranking a constraint which rejects complex codas such as [kr] above a faithfulness constraint which demands that output segments have input correspondents. The effect of resyllabification is simulated by allowing the syllabic affiliation of a stem consonant to vary in output according to whether it appears in an epenthetic or a non-epenthetic alternant. The [v] of the stem [devr], for example, occurs in an onset in [devir] (the output of /devr/) but in a coda in [devri] (the output of /devr-i/).[3]

An immediate problem with this overall approach relates to the assumption that consonants flanking the vowel–zero site constitute coda–onset clusters in non-epenthetic alternants, such as those occurring in the accusative (19.i). Significantly, there are no systematic phonotactic dependencies between these segment pairs. At a pinch, a few of them might be considered to coincide with those well-formed coda–onset clusters that can be identified independently of the vowel–zero alternation. (Those in (19b) might have made good complex onsets, were it not for the fact that Turkish lacks any motivation for recognising such a structure independently of the vowel–zero site.) But any such correspondence must be seen as entirely fortuitous in view of the fact that most of the alternating sequences violate the sonority sequencing constraints which Turkish otherwise respects in the coda–onset context. As illustrated in the forms below, inseparable CC clusters fall into three basic patterns: the onset must be a plosive, while the preceding coda can be a liquid (20a), a homorganic nasal (20b) or [s] (20c).

		i. NOM	ii. ACC	
(20)	a.	[alt]	[altɨ]	'bottom'
		[kulp]	[kulpu]	'handle'
		[mülk]	[mülkü]	'property'
		[tort]	[tortu]	'sediment'
		[sarp]	[sarpɨ]	'rough terrain'
		[terk]	[terki]	'abandon'

b.	[kent]	[kenti]	'town'
	[genč]	[genǰi]	'young'
	[kamp]	[kampɨ]	'camp'
	[deŋk]	[deŋgi]	'equal'
c.	[dost]	[dostu]	'comrade'
	[gasp]	[gaspɨ]	'confiscate'
	[kösk]	[köskü]	'kiosk'

According to both of the analyses just outlined, final -CC] in column (i) of (20) forms a complex coda, while internal -CC- in column (ii) forms a coda–onset sequence. As with the similar English pattern illustrated in (16), this completely misses the exact distributional equivalence between the two contexts. The parallel is precisely what is expected under an alternative analysis which assumes a uniform coda–onset syllabification, as argued in §2.2.

Summarising the Turkish facts just reviewed, we can identify two types of internal -CC- sequence: inseparable clusters in forms such as [kenti] (20) display regular phonotactic patterning, while consonant pairs straddling an alternating vowel–zero site in forms such as [azmi] (19.i) do not. Whether implemented in serialist or constraint-based terms, any analysis which assigns the same coda–onset syllabification to both types of sequence fails to explain this systematic distributional discrepancy.

According to the phonotactic principle referred to above, the lack of any systematic distributional interaction between consonants flanking the vowel–zero context must be taken as evidence that they are not adjacent. Under the resyllabification account, this non-adjacency is only acknowledged in epenthetic forms (and even then only in output). A parallel state of affairs obtains in analyses of vowel–zero alternations which invoke vowel syncope. In this case, the non-adjacency of consonants straddling the vowel–zero site is only evident in initial representation, where the phonotactic independence of the consonants is reflected in their occupancy of positions separated by a vowel. Syncope of the vowel renders the consonants adjacent in output. (In a rule-based analysis, this is achieved by deletion of the vowel and resyllabification of the two consonants. In an optimalist analysis, the two consonants find themselves syllabically adjacent in the successful candidate by virtue of the fact that the vowel is underparsed and thus unsyllabified.)

In short, treatments of vowel–zero alternations which insert or delete syllabic positions fail to acknowledge that, no matter where in derivation we care to sample, consonants surrounding the alternation site consistently demonstrate phonotactic independence.

According to a quite different approach, the abiding phonotactic independence of consonants flanking a vowel–zero context is captured by assuming them to be consistently separated by a nucleus, including in output, as presented in (18c). 'Syncope' then consists in the suppression of melodic material associated with a nuclear position, while the position itself remains intact as an empty nucleus. 'Epenthesis' takes the form of the phonetic interpretation of a stable nuclear position, which under different circumstances remains unexpressed. It then remains to specify the conditions under which the nucleus fails to be made phonetically manifest – the phonological implementation of the Empty Category Principle. According to one proposal, the conditions are defined in terms of PROPER GOVERNMENT: briefly, an empty nucleus is subject to the ECP if it is licensed by an adjacent nucleus which is itself not so licensed (see especially Kaye 1990, Kaye *et al*. 1990, Charette 1991).[4]

Of course, not all word-internal heterosyllabic consonant pairs which fail to display the sonority sequencing expected of interludes enclose an actively syncopating vowel. English, for example, has non-alternating forms such as those in (21a).

(21)	a.	atlas	b.	bottling	bo[t(ə)l]ing
		chutney		fattening	fa[t(ə)n]ing
		kidney		maddening	ma[d(ə)n]ing
		Henry		general	ge[n(ə)r]al
		comely		family	fa[m(ə)l]y

We should be doubly suspicious of any analysis under which such sequences are treated as interludes: not only do they show an upward sonority slope but also they are matched by words which contain the same consonant pairs and yet do exhibit an optional vowel–zero alternation (see (21b)). As a comparison of the parallel consonant pairs in (21a) and (21b) reveals, the distinction between the categorically non-alternating and optionally alternating forms is purely lexical and is not at all dependent on the identity of the potential sequences involved. (Indeed some forms switch between alternation and non-alternation over time or across dialects, e.g. *evening* (*ev*[(ə)]*ning* (N)), *athlete* (*ath*[(ə)]*lete*).)

Under the alternative account outlined above, the failure of heterosyllabic consonant pairs to respect sonority sequencing constraints reflects the fact that the segments occupy separate onsets in phonological output, as depicted in (18c). This configuration remains constant irrespective of whether the intervening nucleus is phonetically unexpressed, either categorically (as in (21a)) or optionally (as in (21b)).

From the viewpoint of the present discussion, a significant point about presonorant consonants in forms such as those in (21) is that they are susceptible to segmental reduction. A similar effect is attested in other languages (as we will see in §5 below). For example, [t] in this position is subject to a loss of audible release or to full debuccalisation to a glottal stop, depending on the dialect. This affects optionally and categorically syncopated forms alike, e.g. *battery* (*ba*[ʔ]*ry*), *atlas* (*a*[ʔ]*las*). If we deny that [t] occupies a coda in such examples, we evidently have to attribute weakening to the following empty-nucleus context. This is the view to be argued for below.

Summarising the discussion in this subsection, we must conclude that not all demonstrably heterosyllabic consonant pairs can automatically be assumed to constitute coda–onset sequences. Some of the relevant strings display an alternative syllabification in which the first consonant occupies an onset rather than a coda. The fact that this position can act as a locus of lenition further undermines the ability of the coda view in (12) to provide a unified treatment of consonantal neutralisation.

2.5 The coda view: the verdict

The work on syllable structure reviewed in this section has led to an unravelling of three strands in the fabric of arguments supporting the coda analysis of consonantal neutralisation. Having examined each of the central claims of the coda view in (12), we conclude that one (12a) should be rejected, that it is premature to accept another (12b) and that the third (12c) is at best unsafe. On (12a), the evidence supports the assumption that a final consonant occupies an onset rather than a coda. On (12b), we should resist accepting coda capture until a simpler alternative which dispenses with the device is proved to be empirically underpowered. On (12c), only a subset of heterosyllabic consonant clusters can be assumed to constitute coda–onset sequences; others form independent onsets separated by a potentially empty nucleus.

Now we seem to be threatening to throw the baby out with the bathwater. Having undermined the central tenets of the coda account of neutralisation, we are in danger of losing the original insight that the arbitrary {C, #} conjunction is an evasive way of referring to a single context more transparently identified in terms of syllable structure. The purpose of the latter part of the paper is to argue that, in order to unify the relevant contexts, it is necessary to delve beyond syllable constituency by hacking into the licensing code that underlies it.

3 Autosegmental and prosodic licensing

3.1 Phonological licensing

The phonetic interpretability of any unit within a representation depends on its being integrated into the phonological hierarchy. Licensing is an asymmetric function which enables this integration by binding each unit in one way or another to some other unit. For example, in order to receive phonetic interpretation, a melodic unit must be associated to a syllabic position; a position must itself belong to a syllabic constituent; and this in turn must be incorporated into some larger prosodic domain, such as the foot or the word. Any unit which for one reason or another fails to be licensed in phonological output receives no phonetic interpretation.

A variety of terms is employed in reference to the different aspects of phonological licensing. Prosodic licensing describes the legitimisation of a whole segment by the syllabic position to which it is attached (McCarthy 1979, Itô 1986). Autosegmental licensing refers to the legitimisation of particular feature specifications by syllabic constituent nodes (Goldsmith 1989, 1990: 123ff, Itô & Mester 1993, Itô *et al*. 1995). Various proposals exist for defining licensing relations that are internal to the prosodic hierarchy, expressed in terms of dependency (Anderson & Jones 1977, Anderson & Ewen 1987), government (Kaye *et al*. 1990), or the strong/weak labelling of metrical and syllabic structure (e.g. Liberman & Prince 1977, Kiparsky 1979).

In what follows, it will be convenient to be able to distinguish between the melodic and prosodic dimensions of licensing. Adapting Goldsmith's notion of autosegmental licensing, I will use the term A-LICENSING to refer to the sanctioning of subsegmental units of melody by the syllabic positions to which they are attached. P-LICENSING, on the other hand, refers to relations that hold within the prosodic hierarchy, ranging from the skeletal tier through the successively higher domains of the syllabic constituent, the foot and the prosodic word (Selkirk 1981, McCarthy & Prince, 1986, Nespor & Vogel 1986).[5] It should be stressed that the A- and P- prefixes do not imply distinct mechanisms but simply refer to different facets of what is a single fundamental principle – that of phonological licensing.

Licensing relations between units in a representation conform to the grammatical principles of locality and headedness (Kaye *et al*. 1990). Locality requires that units in a licensing relation be adjacent on their projection. For example, the p-licensing relation between the two positions of a branching constituent (an onset, say) satisfies this condition, as does the relation between nuclear positions on the foot projection. Headedness manifests itself in the asymmetry of licensing relations. The head of a licensing domain sanctions the presence of any other unit that occurs in that domain. Within the domain of a branching nucleus, for example, the lefthand position licenses its sister by virtue of the fact that it is the head of that domain and is thus projected to the next level of structure (the rhyme). (On the reasons for assuming branching nuclei to be left-headed, see for example Levin 1985 and Harris 1994: 149ff.) Within an a-licensing domain, a skeletal position licenses its associated melodic expression. The inherent asymmetry of these relations is reflected in the following formulation of phonological licensing (after Kaye 1990):

(22) *Phonological Licensing Principle*

Within a domain, all phonological units must be licensed save one, the head of that domain.

The unlicensed head of a domain is itself licensed at some higher level of projection.

The direction of the p-licensing asymmetry can be seen to vary systematically according to the level of projection at which it operates (Kaye *et al*. 1990). Syllabic constituents (onsets, nuclei, rhymes) are universally left-headed. Between constituents (i.e. in onset–nucleus and coda–onset domains), the relation is universally right-headed. In both types of domain, the location of the head is revealed in the phonotactic imparities between adjacent positions. For example, in line with the left-headedness of intraconstituent licensing, the distributional latitude of the righthand position of branching onsets and nuclei is much narrower than that afforded the head position on the left. This is the opposite of the situation obtaining in the interconstituent relation contracted by an onset and a preceding coda; in this context, it is the righthand position which enjoys the greater degree of distributional freedom.

Nuclear head positions are projected to higher levels of the prosodic hierarchy, where they form domains including the foot and the word. At least at the word level, the location of the head is subject to parametric variation, allowing for cross-linguistic differences in such matters as stress placement and the directionality of harmonic spreading.

Interconstituent licensing involves the following sub-clauses (see Kaye 1990 on (23b)):

(23) a. *Onset Licensing*: an onset head position must be licensed by a nuclear position.
b. *Coda Licensing*: a postnuclear rhymal position must be licensed by an onset position.

(23a) is simply a restatement of the traditional notion that it is the nucleus that licenses a preceding onset rather than *vice versa* (expressed as the undominated constraint Nuc in Optimality Theory). Onset–nucleus sequences do not display anything like the degree of phonotactic dependency that is evident in coda–onset clusters (on which point more in §6). Indeed this is one of the primary motivations for recognising the onset–nucleus split in the first place. However, assuming right-headedness in this instance is justified by the pivotal role played by the nuclear head and its projections in the phonological hierarchy (again see Levin 1985 and Harris 1994).

It is via Onset Licensing that we derive the result whereby the onset occupied by a word-final consonant must be sanctioned by the presence of a following empty nucleus, as argued in §2.2. The difference between the type of language which permits final consonants and that which does not (see the typology in (14)) thus involves a choice between whether a final nucleus is allowed to be melodically empty (the former type) or not (the latter type).

By means of Coda Licensing, we derive the effect of onset maximisation. The exclusion of VC.V syllabification follows from the requirement that a coda be licensed by a following onset. A syllabification such as *[pɪt.i] *pity* violates (23b), since the coda [t] is unlicensed. Coda Licensing also plays a role in determining the syllabification of domain-final consonants as onsets rather than codas. In addition to the empirical motivation for this conclusion outlined in §2.2, we now see that the coda syllabification of a final consonant is in any case problematical on theory-internal grounds: there is no following onset to license a coda in this position.

This last point raises a quite general question as to whether constraints associated with phonological licensing are violable under any circumstances (for example, through being bested by some other type of constraint). There seems to be broad agreement that at least some of the constraints just mentioned define rigid design properties of phonological representation. It is relatively uncontroversial to assume that the general principle in (22), the very cornerstone of licensing, merits this status: stray bits of representation are uninterpretable. The same can be said of the requirement that an onset be sanctioned by a nucleus, acknowledged in Onset Licensing (23b) and in the indomitability of Nuc in Optimality Theory.

In work which explicitly espouses Coda Licensing, this too is generally considered to be indefeasible (e.g. Kaye 1990, Kaye *et al*. 1990, Charette 1991). On the other hand, analyses which permit effects that systematically violate Coda Licensing through coda capture or final-coda syllabification typically make no mention of the constraint or of anything directly akin to it (see for example Blevins 1995 and the literature summarised there). (Such analyses also make no attempt to accommodate the kinds of facts identified in §2.2 as posing quite fundamental problems for this overall approach.) The difference in generative capacity between the two approaches is most evident in the syllabification of VCV strings. With a hard interpretation of Coda Licensing, onset maximisation is universally the only possible parse. Without it, both V.CV and VC.V are possible. In what follows, we will encounter no good reasons for departing from the more restrictive approach.

To summarise: the Phonological Licensing Principle manifests itself in a variety of melodic and prosodic domains. In the case of p-licensing, it operates within syllabic constituents, between constituents and between the projections of nuclear head positions. In the case of a-licensing, it holds between melodic units and the syllabic positions to which they are attached. For each language, universal licensing mechanisms, supplemented by particular choices regarding such matters as whether individual constituents may branch, define a set of well-formed prosodic templates. The grammaticality of a phonological representation is then dependent on its being parsable in terms that satisfy these templates.

3.2 Licensing Inheritance

The a-licensing power of a syllabic position is reflected in its ability to support a particular set of melodic contrasts. Discrepancies in the contrastive capacity of different positions, resulting in neutralisation in particular contexts, thus reflect inequalities in the apportionment of a-licensing power. Imbalances of this sort are most vividly illustrated in the phonotactic asymmetries that hold between adjacent positions. For example, the second slot of a binary onset has significantly less distributional elbow room than its sister position. (Typically, as we will see below, the position on the right is restricted to a liquid or glide, while that on the left can support a full range of obstruents.) In coda–onset clusters, it is the first slot that gets distributionally short-changed. A typical situation within a binary foot is to find a maximal vocalic inventory in the dominant nucleus but a reduced inventory in the recessive nucleus.

As already mentioned in §2.1, capturing asymmetries of this sort involves two partially independent issues. On the one hand, we need to be able to identify the specific melodic units that fail to be sanctioned in positions with weak a-licensing power. In §4, I will demonstrate how this is straightforwardly achieved with uniformly privative features (or their equivalent). On the other hand, there is also the more fundamental and thus more pressing problem of explaining why such asymmetries exist at all.

In response to the second issue, I will argue that distributional imbalances can be accounted for by positing an intimate connection between the melodic and prosodic aspects of licensing: the ability of a position to support melodic contrasts depends crucially on its place in the prosodic hierarchy. Specifically, inequalities in the degree of distributional freedom exercised by adjacent positions reflect asymmetries in the p-licensing relations they contract with one another. All other things being equal, the paradigm of melodic contrasts supported by a p-licensing position is potentially greater than that of a p-licensed position. This basic idea is outlined in Harris (1994: 205ff) and has been exploited and further developed by, among others, Brockhaus (1992, 1995), Takahashi (1993), Marotta & Savoia (1994), Savoia (1994), Backley (1995), Nasukawa (1995) and Bafile (1997).

The Phonological Licensing Principle establishes licensing paths which extend throughout phonological representation, running from the a-licensing relation between a skeletal position and its melodic content through to ever higher p-licensing levels involving positions and their projections in the prosodic hierarchy. The head of a given phonological form can be identified as that position which is not licensed at any level of the hierarchy. This position, the ultimate licensor of all units within the form, corresponds to the DESIGNATED TERMINAL ELEMENT of earlier arboreal representations of metrical structure (Liberman & Prince 1977). This is the position, recall, from which a path can be traced through successively higher levels of projection without intersecting any node labelled weak.

By way of illustration, consider the partial representation of the word *pantry* in (24). To allow us to focus on the licensing relations in this form, constituency is represented here in terms of labelled bracketing rather than arboreally.

(24)

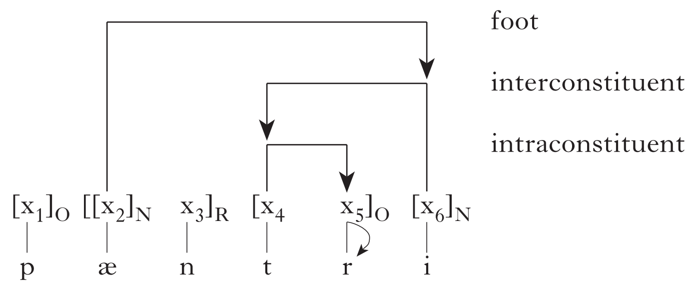

(24) contains a fragment of a map showing the licensing paths (indicated by arrow-headed lines) that traverse this representation. Our attention is directed towards the p-licensing path which terminates in position x_5, the weak (righthand) position of a foot-internal branching onset. This position a-licenses the melodic expression symbolised by [r] and is itself p-licensed on the intraconstituent projection by the onset head x_4. The latter is in turn p-licensed by the nuclear position x_6 on the interconstituent projection. Position x_6, for its part, is p-licensed by x_2, the dominant nucleus on the foot projection (and, in this particular form, on the word projection).

In order to develop the idea that the place of a given position within the p-licensing hierarchy is directly reflected in its ability to a-license melodic material, it will be useful to refer to the position's A-LICENSING POTENTIAL. By this I mean the position's ability either (i) to directly a-license a melodic unit or (ii) to confer a-licensing potential on another position. The proposal to be developed here is that a fundamental asymmetry exists in the a-licensing potential of licensed as opposed to licensing positions, a reflection of the headed nature of p-licensing relations. Specifically, within a given domain, a p-licensing position has a greater degree of a-licensing potential at its disposal than a p-licensed position. Two assumptions will help give this notion some formal substance. Firstly, a licensed position acquires its ability to a-license melodic material from its p-licensor. Call this LICENSING INHERITANCE:

(25) LICENSING INHERITANCE
A licensed position inherits its a-licensing potential from its licensor.

The inheritance of a-licensing potential is exemplified in (24) by the relation between onset position x_5 and its associated melodic expression [r]. This position acquires its a-licensing power from x_4, the head of the onset, which in turn acquires it from nuclear position x_6, its p-licensor on the inter-constituent projection. The ultimate source of a-licensing potential in this form lies with position x_2, the head nucleus of the word *pantry*. This position directly a-licenses its associated melodic expression [æ].

Let us make the further assumption that the stock of a-licensing potential invested in a position is depleted through transmission via an intervening position. A p-licensed position can thus be thought of as a resistor whose effect is to attenuate the a-licensing charge delivered by its p-licensor.[6] This notion, in conjunction with Licensing Inheritance, derives the reduced distributional leeway of a licensed position as compared to that of its licensor. The defective distributional capacity of the righthand position of a branching onset, for example, thus reflects the fact that its potential is diluted as a result of being acquired from another position, namely the licensor on the left.

Below we will explore the impact that Licensing Inheritance has on the a-licensing potential of different positions. Before going any further, however, it is necessary to settle on a model of melodic representation that will allow for a straightforward implementation of the proposal.

4 Defining melodic complexity

4.1 Elements of melody

Thus far I have been speaking of the melodic consequences of licensing in terms of the paradigms of segmental contrast that may appear in different positions. According to this essentially phonemic way of thinking, the greater the a-licensing potential a position possesses, the larger the array of segments it should be able to support. One not particularly satisfactory aspect of this notion, and indeed of any segmental conception of distinctiveness, is that a-licensing potential is not directly coded in phonological representations but has to be calculated by extrinsic reference to segment inventories. Couched in these terms, the relation between a phonological position and the magnitude of the segmental subsystem it hosts is quite arbitrary. Inventories of all shapes and sizes are predicted in principle to be freely distributed across phonological structure, something we know to run counter to the facts. (For a detailed critique of segmental-phonemic approaches to distinctiveness, see Archangeli & Pulleyblank 1994: ch. 2.) In short, brute segmentalism fails to meet any of the goals set out in (1). What is required is a theory that allows a position's a-licensing potential to be directly read off the representation itself. To fulfil this criterion, a-licensing statements have to make reference to the specifics of subsegmental content.

It is not difficult to express the statements in terms of traditional *SPE* features or some direct derivative thereof. However, such formulations are not particularly revealing, mainly because of an intrinsic design property of this type of framework, namely segmental redundancy. Of all the feature values that appear in the output representation of a segment, only those which are distinctive are relevant to licensing, since these are the units of melody that are independently sponsored by the position in question. Redundant values tell us nothing about the ability of a position to license melodic material, since they are present primarily to secure the phonetic interpretability of a segment.

Underspecification can be exploited to provide representations in which only distinctive feature values are visible to licensing statements. This has the desirable effect of defining a level at which the relative licensing potential of each position is transparently represented in terms of the number of different feature values it can support. Unfortunately, the transparency can only be guaranteed in lexical representation; the full specification of redundant feature values in output has the effect of equalising the feature load borne by different positions.

One immediate response to this problem is to ditch bivalency in favour of a privative construction of segmental contrasts – not just for a subset of melodic primes (as in Itô *et al*. 1995 and Lombardi 1995, for example) but for all (as in van der Hulst & Smith 1985, Kaye *et al*. 1985, Cole & Kisseberth 1994 and the further references given below). Uniform privativeness immediately eliminates anything resembling redundant complement values for representations. However, pursuing the goal of ensuring that phonological output is not cluttered with details which reveal nothing about the melodic licensing properties of different positions, we need to go one step further than this and also rid representations of intrinsically non-distinctive features (such as [voice] in vowels). In this way, we banish segmental redundancy from phonology altogether.[7]

This goal can be achieved by endowing melodic primes with fully autonomous interpretability. That is, each prime – ELEMENT is the term I will adopt here – can be conceived of as being phonetically realisable without requiring the support of any other prime. One consequence is to allow for the definition of segments which consist of single elements; each such segment is thus the independent phonetic embodiment of the element it contains. Other segment types can be constructed by compounding different elements. By definition, no redundant properties are needed to support the interpretability of a stand-alone element. The only elements appearing in the output representation of a phonological form are those that are lexically pertinent. By the same token, these elements are the only units that can be referred to in melodic licensing constraints (Harris 1996).

The approach just outlined is either wholly or partially implemented in models which construct vowel space on the basis of the private elements (A), (I) and (U), independently manifested as [a], [i] and [u] respectively (e.g. Anderson & Jones 1974, Schane 1984, Goldsmith 1985, van der Hulst & Smith 1985, Kaye *et al*. 1985, Anderson & Ewen 1987). (I follow the practice of symbolising each element in a way that evokes its independent phonetic interpretation.) Other vowel qualities are derivable by combining elements in the by-now familiar manner. The compound expression (A,I), for example, yields [e], while (A,U) yields [o]. The same type of arrangement can be extended to consonantal representations (cf. Smith 1988, Kaye *et al*. 1990).

One of the advantages of this conception of melodic form is that it permits a simple characterisation of a-licensing potential.[8] In particular, the a-licensing capacity of a position is directly reflected in the complexity of the melody units that can attach to it. For most purposes, complexity is straightforwardly calculated in terms of the number of elements contained in a melodic expression. A difference in licensing potential between two positions will be revealed by the fact that one is able to support a higher degree of elemental complexity than the other.

By way of a very simple preliminary example (similar to a real one we will explore in more detail below), suppose we have a language with a maximal inventory of five canonical vowels, all of which are free to appear in stressed nuclei but only three of which, [a i u], occur in unstressed nuclei. The greater a-licensing potential of the stressed position is revealed in the fact that it can sponsor two-element compounds (the configuration of [e] and [o]), whereas unstressed positions are limited to single-element expressions. Panning out across a representation, we are able to discern contours of melodic complexity which reveal the relative licensing potential of different positions. Moreover, the absence of anything resembling redundant feature values means that these contours stand out in stark relief in phonological output. They are not flattened out by any filling in of missing values.

In principle, Licensing Inheritance can probably be fitted to any privative model of melodic form. The closer any such model gets to zero segmental redundancy, the less call there will be for adaptor mechanisms. In any event, the validity of the proposal does not stand or fall by the particular model that is selected. The one adopted here for demonstration purposes is in most essentials that outlined in Kaye *et al*. (1985), Harris (1994), Harris & Lindsey (1995) and elsewhere.

4.2 Neutralisation as element suppression

From the viewpoint of orthodox *SPE*-type features, neutralising events such as vowel reduction, consonant lenition and sonority sequencing appear as an unruly collection of unrelated melodic effects. Any effect of vocalic or consonantal weakening, for example, has to be expressed as the essentially random substitution of one set of feature specifications by another. In a privative model incorporating autonomously realisable elements, in contrast, both types of event are directly and uniformly represented as a diminution in melodic complexity.

By way of a preliminary illustration, consider the elemental representation of the various lenitions that can befall a labial plosive, including the spirantisation and vocalisation effects displayed by Danish (6) and Ibibio (7). (For a more detailed exposition of the element-based approach to lenition, see Harris 1990, 1994: 119ff.) As shown in (26a), the plosive comprises (U) (labiality), (ʔ) (stop) and (h) (noise). The acoustic expression of (ʔ) is an abrupt and sustained drop in overall amplitude, executed by an articulatory gesture which occludes the oral cavity. (h) contributes aperiodic energy to the speech signal, achieved by a narrowed articulatory stricture which produces turbulent airflow; in the case of a plosive, it manifests itself as the noise burst that occurs on release. We may abstract away from the laryngeal dimension; in the absence of a privative voice element, a stop is interpreted as plain unaspirated (Harris 1994: 133ff, cf. Lombardi 1995). (Angled brackets indicate underparsing of an element.)

(26)	a. Labial plosive		(U,ʔ,h)
	b. Spirantisation	[p] → [f]	(U,⟨ʔ⟩,h)
	c. Stop debuccalisation	[p] → [ʔ]	(⟨U⟩,ʔ,⟨h⟩)
	d. Spirant debuccalisation	[p] → [h]	(⟨U⟩,⟨ʔ⟩,h)
	e. Vocalisation	[p] → [w]	(U,⟨ʔ⟩,⟨h⟩)

Spirantisation consists in the suppression of (ʔ), the remaining elements (U, h) defining a labial fricative (see (26b)).[9] Stop debuccalisation (26c) involves the exclusion of (U) and (h); here the glottal articulation of the residual (ʔ) reflects the fact that this is the only way of executing a stop gesture in the absence of any place-defining element. In the case of spirant debuccalisation (26d), the remaining element is (h), with glottal location this time automatically implementing bare friction. The lone (U) that remains when (ʔ) and (h) are stripped away defines [w], the outcome of vocalisation (see (26e)).

These mappings illustrate the autonomous interpretability of elements: the outcome of each event immediately defines a fully realisable segment. There is no call for supplementary fill-in operations or linking conventions of the sort that are required when the same events are expressed in terms of *SPE* features. As we will see later (§6), the same general point can be made in relation to the representation of vowel reduction.

Let us now consider how this general account extends to the treatment of the other consonantal neutralisations introduced in §2.1. As shown in (27a), the element (R) is stifled in Spanish *s*-debuccalisation (recall the examples in (2a)); the outcome, [h], is the solo interpretation of the remaining element, (h).

(27)	a. *s*-debuccalisation	[s] → [h]	(⟨R⟩,h)
	b. Liquid gliding	[r] → [y]	(⟨R⟩,I)
		[l] → [w]	(⟨R⟩,U,⟨ʔ⟩)
	c. Lateral depalatalisation	[ʎ] → [l]	(R,⟨I⟩,ʔ)
	d. Nasal depalatalisation	[ɲ] → [n]	(R,⟨I⟩,ʔ,N)

Liquid vocalisation and sonorant depalatalisation, exemplified by Spanish in (2b–d) and by Portuguese in (3a), involve similar patterns of elemental reduction. In the case of the various gliding events affecting liquids, we may assume that the vocalic outcome reflects a segment's secondary resonance characteristic. Thus vocalisations to [y] and [w] indicate respectively a clear (palatalised) and a dark (labial-velarised) source. In element terms, a palatalised [r] is composed of (R) and (I). As shown in (27b), gliding of the type that occurs in Cibaeño Spanish involves the muting of (R), the residual (I) element being independently manifested as [y]. Coronal laterals are made up of (R) and (ʔ), to which can be added a secondary resonance component. Vocalisation to [w] involves the suppression of all but an (U) component. As shown in (27c–d), depalatalisation of laterals and nasals (illustrated by the Spanish forms in (2c–d)) results in the delinking of (I). ((N) is the nasal element.)

The neutralisation events exemplified in (3b) and (5) are represented in (28). The defective distribution of [r] and [h] consists in the exclusion of (R) (28a) and (h) (28b) respectively.

(28)	a. *r*-vocalisation/loss	[r] → ∅	(⟨R⟩)
	b. Defective *h*	[h] → ∅	(⟨h⟩)
	c. *t*-tapping	[t] → [ɾ]	(R,⟨ʔ⟩,⟨h⟩)

As shown in (28c), suppression of stopness (?) and noise release (h) in a coronal plosive leaves (R), which is independently interpreted as a tap (Harris & Kaye 1990, Harris 1990).[10]

Completing our element-based round-up of the phonological events featured in §2.1, let us consider the effects of laryngeal neutralisation in obstruents. In languages lacking a voice contrast, such as native Finnish, Quechua and most Australian languages, plosives are typically plain (Maddieson 1984). As mentioned above, a plain stop bears no laryngeal element. (On the elemental representation of source elements, see Harris 1994: 133ff and the references there.) Systems with a two-way laryngeal distinction fall into two basic types. On the one hand, there are languages such as French, Polish and Dutch, in which plain stands in contrast with a fully voiced series, the latter containing the slack vocal cords element (L). On the other, there are languages such as English, Danish and Icelandic, in which plain is opposed by a voiceless aspirated series, the latter containing the stiff vocal cords element (H). Three-way systems such as Thai and Sesotho show all three possibilities – plain *vs.* (L) *vs.* (H). In element terms, the suspension of a laryngeal distinction is in principle no different from that affecting manner or place: it takes the form of element suppression. Muting either (L) or (H) in an obstruent gives rise to a plain segment.

As shown in (29a), obstruent devoicing consists in the effacement of (L) (Brockhaus 1995). The failure of aspiration to appear on a stop in a given context is expressed as the masking out of (H) (29b).

(29)	a. Suppression of voicing	[b] → [p]	(U,?,h,⟨L⟩)
	b. Suppression of aspiration	[p^h] → [p]	(U,?,h,⟨H⟩)

The Danish facts in (10) submit to a treatment in terms of (29b). While the distinction between aspirated and plain plosives is maintained word- and foot-initially, only the plain series appears in other contexts. (H) is thus supported in the contrastive site but is excluded elsewhere. (The susceptibility of plain obstruents to optional and gradient phonetic voicing in vocalic contexts is consistent with their being devoid of any independent laryngeal specification. A similar spontaneous voicing effect accompanies tapping; in this case, it happens to be part of the independent interpretation of (R), the sole element that remains when (?) and (h) are suppressed in a coronal plosive (shown in (28c).) In Thai, while both (H) and (L) appear word-initially, neither is sanctioned word-finally, resulting in a collapse of the three-way laryngeal distinction under the plain series. (Word-final plain stops in Thai are unreleased, indicating that the noise element (h) too is barred from this context. Ibibio is identical in this respect (cf. (9c)).)

The statements in (26)–(29) target the pieces of melody that particular grammars elect to suppress in the relevant neutralising contexts. The statements are inevitably stipulative to the extent that they deal with details of segmental form in which Universal Grammar allows for cross-linguistic variation. On the other hand, any theory of melody worth its salt must account for the fact that the set of observed events to which a segment is susceptible is comparatively small. In the privative framework outlined in this section, this set is delimited (i) by the number of elements of which the segment is composed and (ii) by the assumption that neutralisation is uniformly represented as a contraction in melodic complexity. In this way, element theory directly addresses one of the challenges thrown down at the outset of this article, that of integrating the segmental effects of neutralisation (1a). We will see below how the same basic notion of elemental complexity extends to the treatment of sonority sequencing constraints, with the difference that in this case relative complexity is calculated over pairs of neighbouring segments rather than over the mapping between input and output.

In the following section, we switch our attention to another of our initial goals, that of uniting the contexts in which consonantal neutralisation occurs (1b). At the same time, I will try to demonstrate that the Licensing Inheritance proposal, in conjunction with element theory, allows us to tackle the third and most challenging task, that of supplying a causal link between the segmental effects and contexts in question (1c). To the extent that the attempt is successful, it promises an answer to what is perhaps the most fundamental question we can ask about neutralisation: why does it occur in the first place?

5 The melodic consequences of Licensing Inheritance

5.1 Introduction

The Licensing Inheritance proposal is an attempt to explain the variations in melodic complexity that are symptomatic of neutralisation. Complexity differentials follow from the assumption that a-licensing potential emanates from prosodically strong positions and is dissipated through transmission to weaker positions.

In this section, we will examine the workings of Licensing Inheritance in the consonantal neutralisation sites listed in (11). We start in §5.2 by focusing on the melodic complexity effects that are observable in branching onsets and coda–onset interludes, two contexts where consonantal neutralisation is traditionally described in terms of sonority sequencing constraints. What unifies these sites is that, in each case, a relation of direct licensing holds between two positions. According to the Licensing Inheritance proposal, one of these positions, the licensor, should display an ability to sustain a greater degree of melodic complexity than the other, the licensee.

An earlier instantiation of the notion that elemental asymmetries are attributable to licensing relations within complex onsets and interludes is to be found in Kaye *et al*. (1990) and in the Complexity Condition proposed in Harris (1990). (For a similar idea couched in terms of features, see Rice 1992.) This specific proposal, we will see, can be subsumed under a generalised theory of Licensing Inheritance.

In order to show that Licensing Inheritance provides a unified account of consonantal neutralisation, it will be necessary to demonstrate that it can be extended to domain-final and foot-internal contexts (§5.3). What distinguishes these two sites from complex onsets and interludes is that the segmental effects which occur there cannot be attributed to the direct licensing influence of an adjacent position.

It will also be necessary to say something about the status of Licensing Inheritance in two types of domain containing nuclei. In one case, a nuclear position licenses a non-nuclear position, either a preceding onset or a following coda. In the second, one nuclear position licences another, either at the foot or word level. Both types involve direct licensing relations between positions which are adjacent on some projection. However, it is not immediately obvious that they are subject to the same sort of distributional asymmetries as those associated with non-nuclear domains. Nuclei are to a large extent phonotactically independent of neighbouring onsets and codas – not at all the situation that Licensing Inheritance would lead us to expect. Moreover, although it is not difficult to find examples of a distributional imbalance between nuclei at the foot or word level, this pattern is evidently not adhered to in all languages. These issues are taken up in §6.

5.2 Melodic complexity differentials under direct licensing

5.2.1 *Onsets*. The righthand position of a branching onset is predicted to be a neutralising context by virtue of the fact that it inherits its a-licensing power from its p-licensor, the onset head on the left. This is consistent with the distributional restrictions placed on this position by generally recognised sonority sequencing constrains. Whether the restrictions constitute hard universals or soft preferences is a matter we take up presently.

A cross-linguistically typical pattern in onset clusters is to find that a full set of obstruents is supported in the head position. In the weak position, in contrast, there is barely room to swing the distributional cat. Here, melodic content is often restricted to a glide or liquid. In terms of the melodic model outlined in §4, the narrowed distributional latitude of the weak slot is reflected in an ability to support only one element or at best two. Thus in this context, we find either (i) a glide represented by the element (I) (independently manifested as [y]) or (U) (manifested as [w]) or (ii) [r] (the independent exponent of (R)) or (iii) [l] (composed of (R) and (?)). In some languages, a maximum of one element is tolerated in the weak onset position. Witness, for example, the historical reduction of [l] in this position to a simplex segment in some Romance languages – [r] (= (R)) in Portuguese and [y] (= (I)) in Italian, cf. French *plat* [pla], Portuguese *prato* [pratu], Italian *piatto* [pyatto] 'dish'. As illustrated by the well-formed onset clusters in (30), this distributional defectiveness stands in marked contrast to the onset head, which directly a-licenses obstruents – expressions displaying a greater degree of complexity than liquids and glides. (Elements appear here and below in matrix-shaped packages only for the purposes of saving space. I assume they are deployed on independent tiers in classic autosegmental fashion. Whether or not they are also gathered under geometric class nodes has no bearing on the present discussion.)[11]

(30) a. [br] b. [tw] c. [fl]

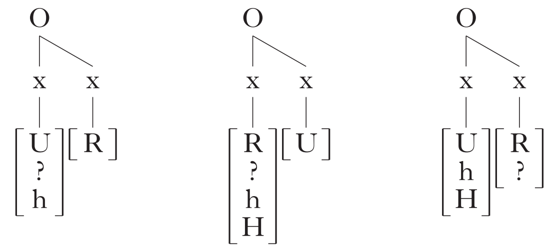

This is an appropriate point at which to ask whether or not Licensing Inheritance can be violated in phonological output, the question we posed in relation to the other aspects of licensing discussed in §3.1. To put it in theoretically loaded terms: does Licensing Inheritance constitute an indefeasible universal principle, or is it subject to parametric variation? Alternatively: is Licensing Inheritance a universally undominated constraint (perhaps part of Gen), or can it be violated on a grammar-specific basis through being outranked by other, conflicting constraints? Although we will return to this question at several points below, it is important to stress that the answer has little direct bearing on the validity of the present proposal. The issue it raises is largely a derivational one, and Licensing Inheritance is about representation. In fact both hard and soft interpretations of the condition are to be found in the literature. Harris (1994), Brockhaus (1995) and Bafile (1997), for example, assume it to be an inviolable principle. Nasukawa (1995), on the other hand, affords it the status of a rankable constraint within Optimality Theory.

The response to the violability question determines whether or not we are prepared to countenance the onset syllabification of clusters, such as [sp zd ft md], which do not display the rising sonority profile otherwise associated with this constituent. The cluster [sp], for example, sports an upward complexity slope (viewed from the left) – (R, h) succeeded by (U, h, ?). Accommodating both consonants in a branching onset would require a greater degree of a-licensing power to be invested in the licensed position than in its licensor. This is sufficient to damn onset [sp] in the eyes of Licensing Inheritance. Whether it is also sufficient to exclude the structure in principle from phonological output depends on whether the condition is deemed defeasible. Under a hard interpretation, this pattern cannot constitute a branching onset under any circumstances, even in word-initial position (the view assumed in Harris 1990, for example). In this particular instance, independent evidence indicates that such clusters form coda–onset interludes (see Borowsky 1984, Gussmann & Kaye 1993, Kaye 1996 and §5.2.2 below). A soft interpretation of Licensing Inheritance goes hand in hand with a more promiscuous view of syllabification, according to which any consonant string can potentially form an onset cluster, even if it flouts otherwise quite general sonority sequencing constraints.

5.2.2 *Interludes*. Turning now to the a-licensing properties of codas, we must resist any temptation to let our eyes stray to domain-final consonants. In accordance with Coda Licensing, a coda position will always be domain-internal, followed as it is by an onset. Moreover, we should remember to exclude bogus clusters of the sort identified in §2.4.

Within authentic interludes, the distributional space of codas is systematically cramped, an effect that manifests itself in two ways. Firstly, the independent ability of a coda to sponsor contrasts is typically less than that of a following onset. Secondly, while different melodic properties may be interpreted in a coda position, these are often wholly or partially assimilated from the following onset.[12] Under such circumstances, it is thus the onset rather than the coda that bears the contrastive load. In other words, the onset is afforded a degree of a-licensing power that is denied to the coda. The effect of a rising complexity slope across an interlude is guaranteed by the coda's failure to license place (yielding homorganicity) or stopness or noise or voice (yielding laryngeal assimilation) or some combination of these.

This overall imbalance manifests itself in a recurrent pattern whereby a coda is only permitted to sponsor liquids, or [s], or a segment that is homorganic with a following onset consonant, or some combination of these. Variations on this general theme are recorded in languages such as Axininca Campa, Diola Fogny, Hausa, Italian and Lardil (for referenced discussion and further exemplification, see Itô 1986, Goldsmith 1990: 128ff, Kaye *et al*. 1990, Rice 1992 and Itô & Mester 1993). Even languages in which the coda is allowed a much freer distributional hand, such as English, show clear evidence of a righthanded bias in interlude phonotactics (Harris 1994: 66ff).

In terms of the present proposal, the diminished a-licensing potential of a coda stems from the fact that it is inherited from the following onset position, the coda's p-licensor within the interconstituent domain. The curtailed distributional freedom of codas has elsewhere been expressed in terms of filters such as coda conditions (Itô 1986) or cluster conditions (Yip 1991). As noted in §2.1, such devices in and of themselves do not explain why these restrictions systematically target particular positions rather than others. For example, they offer no reason why codas are consistently kept on a much tighter contrastive leash than the onsets that follow them. Within a framework incorporating Licensing Inheritance, the restricted distributional freedom of codas follows straightforwardly from the right-headedness of the inter-constituent licensing relation: the onset, *qua* p-licensor, enjoys a greater degree of a-licensing power than the coda, *qua* p-licensee.

The diminished a-licensing power of codas is most strikingly illustrated in 'Prince' languages, those in which coda–onset clusters are restricted to geminates and/or partial geminates (Prince 1984, Goldsmith 1989). Adopting a widely held view, we may assume that the occurrence of adjacent identical melodic expressions is prohibited (under the OCP or some other such constraint) and that the melodic content of a true geminate is lexically specified in the onset and spreads automatically into the coda (Kaye *et al*. 1990, Itô & Mester 1993). In other words, the coda does not itself license the melody in question, that role being assumed by the onset. Spreading here is thus to be understood as an instruction for the licensed position to be interpreted in terms of melodic material that is sponsored by the p-licensor (Cole & Kisseberth 1994, Harris 1994: 164ff). Assimilation under this view does not result in the copying of melodic content from one position to another, as is typically assumed in linear-segmental traditions. A fully non-linear conception of spreading is crucial to the achievement of a unified treatment of both assimilative and reductive neutralisation. The linear-segmental scenario in which an assimilating position clones melody from some other source can only be understood as fortition. Representationally, this immediately divorces the effect from segmental reduction, thereby obscuring the fact that both types of event result in neutralisation. Under the present account, in contrast, both assimilative and reductive neutralisation involve a diminution in the degree of melodic complexity that a position is able to license. Whether or not the position is phonetically interpreted in terms of melody that is licensed by some other position is a separate matter.

In a full geminate, the coda position fails to a-license any elements of its own, and its phonetic expression depends entirely on its identification with the melodic content of its onset licensor. Partial geminates reflect the failure of a coda to a-license all but those elements that are responsible for defining nasals and/or liquids. These structures are illustrated in (31) (where ≪ denotes spreading).

(31) a. *Geminate* [pp] b. *Partial geminate* [mb]

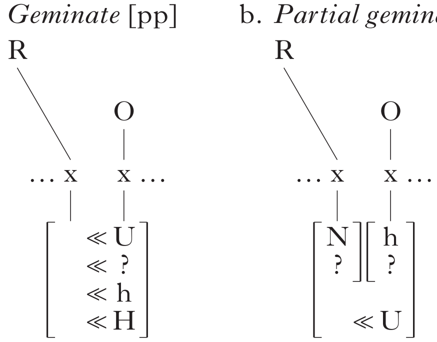

The melodic complexity differential that is evident in geminates is also to be observed in non-Prince languages, i.e. those that permit obstruent clusters which are not necessarily homorganic. In such sequences, we must conclude, a place-defining element is distinctively present in the coda. In English, for instance, we find a contrast before [t] between [s] and

[f] (e.g. *castor vs. after*), as well as between [p] and [k] (e.g. *chapter vs. factor*). Nevertheless, the diminished a-licensing potency of codas is also evident in obstruent clusters and is reflected in an inability to a-license independent laryngeal elements (as exemplified by the Danish pattern in (10d)). It is also discernible in a preference for the first of two oral stops in an interlude to remain unreleased; here the licensed coda position lacks a counterpart to the licensing onset's (h) element. An upward complexity gradient (viewed from the left) is thus maintained in interludes such as [ft] (where (U, h) < (R, h, ?)) and [pt] (where (U, ?) < (R, h, ?)).

5.2.3 '*Sonority*'. The phonotactic regularities observable in both of the contexts just reviewed, branching onsets and coda–onset interludes, are traditionally characterised in terms of sonority sequencing generalisations. With orthodox features, sonority is not directly coded in representations but is gauged by reference to an extrinsic hierarchy which is little more than a stipulative ordering of segment types. With elements, in contrast, 'sonority' is merely a descriptive label for a formal property that is intrinsic to representations, namely melodic complexity. (Translating loosely between theories, we might say that the more elements a consonant has the less sonorous it is.) This allows us to dispense with the sonority hierarchy as a theoretical construct. Instead, the phonotactics of onset and interlude clusters are transparently displayed as inequalities in the degree of melodic complexity that can be supported by adjacent positions. The defective distribution of melodic units in the righthand slot of a branching onset and in the lefthand slot of an interlude reflects the fact that each position is p-licensed by an onset head. In both cases, the position's a-licensing potential is thus restricted through being inherited from some other source.

'Sonority' is in any event only a partial description of the empirical ground which the present account can be made to cover. As we will see below, the same basic notion of elemental complexity can be extended to contexts which are not traditionally deemed to involve sonority.

5.3 Indirect licensing

5.3.1 *Introduction*. Of the three contexts identified in (11) as prime sites for consonantal neutralisation, we have now established that one, the coda, favours neutralisation on the grounds that it inherits weak a-licensing potential. The question now is whether this account can be extended to the other two contexts – foot-internal intervocalic position and domain-final position. I will now try to demonstrate that it can. It is possible to discern a licensing configuration that is common to all three of the contexts in question. We will first consider how the domain-final VC] site can be shown to reduce to the foot-internal site (§5.3.2) and then go on to see how these two unify with the coda (§5.3.3).

A coda–onset interlude forms a domain in which two consonants stand in a direct licensing relation and thus engage in an unequal competition for distributional resources. It is perhaps not immediately obvious how licensing mechanisms might be invoked to account for the distributional properties of single final or intervocalic consonants, where the conditions for a two-sided contest of this kind do not exist. The neutralising tendency of the two sites must evidently be related indirectly to relations involving neighbouring nuclear positions.

Before going any further, it is worth reminding ourselves that, in seeking to unify the contexts in question, we are not aiming to collapse them altogether. That would be no more desirable in an account based on licensing than in one based on constituency. Internal coda C and domain-final C, for example, are certainly broadly similar in their diminished ability to support contrasts, but they are by no means always identical in the segmental details that implement this overall effect. Even under a uniform coda approach, it is still often necessary to invoke additional conditioning factors in order to tease the two sites apart. For example, Selkirk's (1982) coda-based treatment of tapped and unreleased reflexes of English [t] requires that a distinction be drawn between the two contexts on the basis of whether or not a consonant or pause follows.

The type of dynamic alternation featured in (2) and (3) is only one instantiation of the simultaneous impact of neutralisation on internal coda C and final C]. The same overall effect can be manifested in a static distributional manner. In some languages, the segmental inventories sponsored by the two contexts are either identical or nearly so. Selayarese is one example (Mithun & Basri 1986, Goldsmith 1990: 131ff). Final consonants in this language are confined to [ʔ] and [ŋ], which can reasonably be taken to indicate that place is not licensed in this context (Rice 1996). The same restriction applies to internal codas, where we find either [ʔ] or the first part of a full geminate or a nasal that is place-linked to the following onset. Another example is Lardil, where both positions only tolerate liquids and nasals, the slight difference being that final C] is restricted to apicals, while consonants in internal codas are subject to certain constraints binding their place specification to that of the following onset (Hale 1973, Itô 1986: 84ff).

Nevertheless, it is by no means true that complete isomorphism between the contrastive profiles of internal and final consonants represents the cross-linguistic norm. In many languages, internal coda C draws a shorter contrastive straw than final C]. In Turkish, for example, internal syllable-closing consonants are restricted to [l r s] and nasals that are place-linked to the following onset (see (20)). Sonorants and [s] can also appear word-finally. So can nasals, although here they are contrastively specified for place. Moreover, final C] also supports plosives and affricates. A similar pattern is found in native Malay, where obstruents are free to occur domain-finally but not in an internal coda (Maris 1980).

Pondering the significance of total or close distributional resemblances between internal coda C and final C] brings us back to the question considered in §2.2: should these similarities be taken as proof of syllabic identity? To answer this in the affirmative, by treating both positions as codas, is to ignore the otherwise robustly non-coda-like behaviour of final

C]. An alternative response might be to assume that the syllabification of final C] varies cross-linguistically, such that it is deemed extrarhymal unless distributional considerations suggest otherwise (e.g. Piggott 1993) – a general approach similar in spirit to that in which intervocalic C is assumed to be coda-captured in some languages but not others. However, before agreeing to loosen syllable theory in this way, we need to assure ourselves that there is no other possible basis for the distributional parallel. As I will now try to demonstrate, by identifying a common pattern of licensing in the two contexts, we are able to capture the parallel without compromising the extrarhymal status of final consonants or deviating from the principles of basic syllabification.

5.3.2 *Foot-internal* (*including domain-final*). A union of intervocalic and domain-final sites can be straightforwardly achieved without resorting to coda-capturing resyllabification. On the strength of the empirical evidence and theory-internal arguments rehearsed in §2.2 and §3.1, we proceed on the assumption that a final consonant occupies an onset supported by an empty nucleus. Syllabically, this configuration is of course identical to that in which an intervocalic consonant finds itself. The difference between the two contexts boils down to the fact that the postconsonantal nucleus is melodically empty in one case but filled in the other.

The two contexts also cohere in terms of their p-licensing properties, at least in the case where the intervocalic context forms a trochaic foot. As pointed out in §2.2, there are good grounds for assuming that a final empty nucleus is licensed by the preceding nucleus, just as a weak vowel is licensed by the dominant nucleus of its foot. The reasons for reaching this conclusion are quite independent of any considerations involving consonantal neutralisation. For one thing, the metrical properties of final V́C] are straightforwardly unifiable with those of V́CV̌], if we make the reasonable assumption that a final empty nucleus may be metrified in the same way as a final weak vowel. This view is consistent with the claim that all feet are minimally binary and that the word in many languages must consist minimally of a foot (McCarthy & Prince 1986). Under a final-coda analysis of English, [(C)VC] forms such as *sit* are deemed to satisfy prosodic minimality by virtue of containing a single bimoraic syllable in the same way as [(C)VV] forms such as *see*; [(C)VCV] forms such as *city* fulfil the requirement by containing two morae in separate syllables. Under the null-vowel view, in contrast, the 'bimoraic' prosodic frame is defined in terms of nuclear licensing: a foot/word minimally comprises a domain in which one nuclear position licenses another. In this way, *sit* ([CVC∅]) lines up with *city*: both feature foot-level licensing between two independent nuclei. Note how this unification is achieved without having to renege on a commitment to the non-rhymal status of final consonants.

The question now is why foot-internal onsets, both intervocalic and domain-final, make poor melodic licensors. We will see below (§6) how the asymmetry of licensing relations explains the distributional discrepancies between nuclei within a foot. According to the Licensing Inheritance account, we expect this imbalance to be potentially mirrored in the onsets the nuclei license. That is, just as the distributional latitude of the recessive nucleus of a foot is more tightly squeezed than that of the dominant nucleus, so the onset licensed by the recessive nucleus should be afforded less distinctive potential than that of the onset licensed by the dominant nucleus. This expectation is indeed consistent with the observed tendency for distinctions holding in foot-initial onsets to be neutralised foot-internally (recall the examples in (5)–(7), (9)–(10)). The asymmetry can be seen to follow from Licensing Inheritance, once we compare the relevant licensing paths which terminate in the two contexts:

(32) a. *Foot-initial* C b. *Foot-internal* C (intervocalic/domain-final)

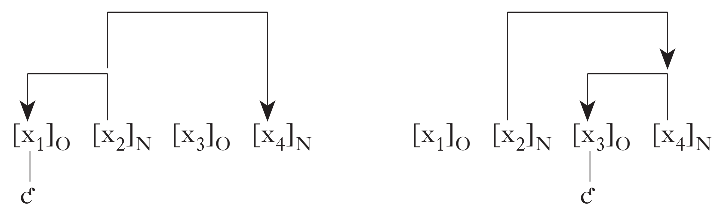

As shown in (32a), the source of a-licensing potential exercised by the foot-initial onset (x_1) lies at one remove from the position itself – specifically in its p-licensing nucleus x_2, the dominant nucleus of the foot. By contrast, the diminished a-licensing potential of the foot-internal onset (x_3) in (32b) reflects the fact that it is inherited at two removes – immediately from nucleus x_4, ultimately from nucleus x_2.

One of the significant advantages of this conception of foot-internal onsets as weak a-licensors is that it enables us to dispense with coda resyllabification. By way of illustration, consider how the defective distribution of [h] in English (5b) can be reanalysed without resorting to the coda-capture device employed in other accounts. [h] can occur foot-initially (e.g. *ve*[h]*ícular*) but not intervocalically within a foot (e.g. *vé~~h~~icle*). Nor can it occur domain-finally; hence the ungrammaticality of forms such as *[vi:h]. In terms of the present account, a lone (h) element (individually manifested as [h]) is a-licensed by a foot-initial onset but not by a foot-internal onset, irrespective of whether this is intervocalic (*VhV) or domain-final (*Vh∅]).

Note that, although we have identified a common licensing configuration in VC] and VCV, we have not gone so far as to merge the contexts altogether. We still need to preserve a distinction between the locations, given that not every neutralising event occurring in one necessarily also occurs in the other. A case in point is Spanish *s*-debuccalisation; as illustrated in (2a), this affects VC] but not VCV. In coda-based theory, these contexts remain distinguishable (as coda *vs.* onset) in a grammar that lacks coda capture. In a theory in which both contexts potentially form a foot, the distinction can only be drawn by referring to the melodic content of the second nucleus: [s] lenites before a melodically empty nucleus ([mes∅] → [meh∅]) but retains buccality before a nucleus with vocalic content. The prosodic status of the filled nucleus is irrelevant: it can be the weak member of a foot (as in [kása] 'house') or the strong (as in [kasár] 'marry').

The facts of *t*-lenition in English and Ibibio provide further confirmation that, where a distinction needs to be drawn between intervocalic and domain-final instantiations of the foot-internal configuration, this must be accomplished by reference to melody. The contexts in which tapping takes place in both languages are exemplified in (33a–c) (relevant details of foot structure parenthesised). The conditions for tapping can be summarised as follows: the target segment must occur (i) before a vowel ('the vocalic condition') and (ii) either word-medially within a foot (as in (33a)) or word-finally (as in (33b–c) ('the domain condition')).

(33)		*English*	*Ibibio*	
tap	a.	(bítter)	[(boro)]	'be created'
	b.	(bít) of	[(bor) owo]	'create someone'
	c.	(bit) óff		
plosive	d.	bou(tíque)	[u(taŋ)]	'plaiting'
unreleased stop	e.	(bit)|	[(bot)]|	'create'
	f.	(bit) me	[(bot) mmoto]	'create a vehicle'

In the event of the vocalic but not the domain condition obtaining, [t] remains unlenited (see (33d)). If on the other hand the domain condition holds but not the vocalic, [t] is realised as an unreleased stop; this is the situation encountered both utterance-finally (indicated by |, see (33e)) and preconsonantally (see (33f)).

Coda-based treatments of this phenomenon in English typically have to rely on various forms of resyllabification in order to manoeuvre [t] into the tapping site (e.g. Kahn 1976, Selkirk 1982, Borowsky 1986). As in §2.3, my intention here is not to question the descriptive adequacy of this overall approach. Rather it is to show that the facts in question can be accounted for without recourse to resyllabification and without losing sight of the non-coda-like behaviour of final consonants.

According to an earlier analysis, which provides a springboard for a Licensing Inheritance treatment of the same facts, the relation between tapped and unreleased reflexes of English [t] is secured by defining the generalised lenition site as consisting of an onset followed by a weak nucleus (Harris & Kaye 1990, Harris 1990, Harris 1994). Under that account, the weakness of the nucleus in question stems from the fact that it is licensed either by the head of its foot (as in *city*) or, in the case of its being empty, by parameter (as in *pit*∅). Under the alternative proposal anticipated in §2.2, the empty nucleus is also incorporated into a foot. The generalised lenition site is now simply definable as a foot-internal onset.[13]

In both Ibibio and the type of English represented in (33), the foot-internal context thus accommodates not only the tapped but also the unreleased reflexes of [t]. Tapping only occurs if the following nucleus has melodic content. This condition may be satisfied on the word cycle (as in (33a)) or at phrase level (as in (33b–c)). Otherwise the unreleased stop appears (as in English *bi*[t$^{\neg}$] and Ibibio [bot$^{\neg}$] in (33e–f)). The independent nature of the melodic conditioning is confirmed by the fact that, as in Spanish *s*-debuccalisation, the prosodic status of the following nucleus is irrelevant: in English, tapping occurs word-finally before both an unstressed vowel (33b) and a stressed vowel (33c). The reason tapping fails before a word-internal strong nucleus (as in *boutíque* and [u-taŋ] (33d)) is that [t] in this position never finds itself in the foot-internal site that plays host to lenition.

The prosodic conditioning of *t*-lenition in English and Ibibio illustrates the general distributional imbalance between foot-initial and foot-internal onsets. Being p-licensed by the dominant nucleus in the foot domain (x_2 in (32a)) imbues the initial onset (x_1 in (32a)) with a relatively high degree of a-licensing power: it is able to support the full melodic complexity of plosive [t] (composed of (R, ʔ, h) and, in English, (H)). The attrition of elements that occurs in a foot-internal onset results in loss of release (yielding (R, ʔ)) or tapping (to (R)). These effects reflect the fact that the a-licensing capacity of the foot-internal onset (x_3 in (32b)) is diminished through being inherited via the weak nucleus in the foot (x_4 in (32b)).

If it is correct to attribute the defective contrastivity of domain-final consonants to the weak a-licensing capacity of a following empty nucleus, we should expect similar behaviour to be displayed in the other context where an onset is licensed by an empty nucleus, namely before a word-internal syncope site (see §2.4). The a-licensing potential of both types of nucleus is weakened through being inherited from some other nucleus – the foot head in the domain-final case, a proper governor in the internal case. There is evidence that the two contexts can indeed define a single conditioning site for neutralisation.

One example is German final obstruent devoicing: (L), lexically specified in an onset, fails to be awarded an a-licence when the following final nucleus is empty, as in *Han*[t]∅ 'hand' *vs.* *Hän*[d]*e* 'hands' (Brockhaus 1992, 1995). Empty nuclei also occur word-internally in German, in a syncope site which is otherwise occupied by schwa; compare *Hand∅lung* 'act' with *Hand*[ə]*l* 'trade'. As Brockhaus points out, some dialects treat the internal and word-final contexts identically for the purposes of devoicing, giving rise to alternations such as *Han*[d]*el* – *Han*[t]∅*lung*, *e*[b]*en* 'level' – *E*[p]∅*nung* 'levelling'.

The same pattern of internal and final empty nuclei operating in tandem is evident in English *t*-lenition. Dialects which display unreleased or debuccalised reflexes word-finally when a consonant or pause follows, as in *get∅ by*, show the same reflex in the parallel internal site illustrated by forms such as *bott∅ling* and *at∅las* in (21) (Harris & Kaye 1990, Harris 1990, 1994: ch. 4).

The coda-based and licensing-based approaches come up with opposite responses to the weak status of foot-internal consonants. The coda-capture approach merges the foot-internal and domain-final sites under the latter, while the present account merges them under the former. Having identified the coda as the common neutralisation site, a constituent-based approach must find ways of manipulating a foot-internal consonant into it, but it can do so only at the expense of violating independently established principles of basic syllabification. Moreover, the directionality of the unification in this case does violence to the observed extrarhymal status of domain-final consonants. Under the licensing account, treating a domain-final consonant as the onset of a null-vowelled syllable allows us to unify this site with the foot-internal context. Not only does this respect the extrarhymal status of domain-final consonants but, by dispensing with the need for resyllabification, it is also derivationally leaner.

5.3.3 *Foot-internal and coda.* Having combined the intervocalic and domain-final contexts under a single foot-internal licensing configuration, we turn now to the question of whether this can for its part be unified with the other consonantal neutralisation site identified in (11), true (i.e. internal) codas. The key to an answer is to be found when we compare the licensing paths that terminate the positions in question, coda x_1 in (34a) and onset x_2 in (34b):

(34) a. *Coda* C

$[x_1]_R \ [x_2]_O \ [x_3]_N$

c

b. *Foot-internal* C

$[x_1]_N \ [x_2]_O \ [x_3]_N$

c

Although the two sites clearly differ in terms of constituency, they are located in similar licensing milieux. In both instances, the neutralising position is licensed (at the interconstituent level) by a position which is itself licensed from some other source. In (34a), the coda (x_1) is licensed by an onset (x_2) which is interconstituent-licensed by a following nucleus (x_3). In (34b), the onset (x_2) is licensed by a following nucleus (x_3), which is in turn licensed on the foot projection by the preceding nucleus (x_1). In each case, therefore, the source of the a-licensing potential inherited by the neutralising position is separated by at least one other licensed position from the point at which it is discharged. Being depleted at two stages on a licensing path, the a-licensing capacity invested in both positions is expected to be correspondingly diminished, with adverse consequences for their ability to support melodic complexity. In short, it is the lowly p-licensing status of codas and foot-internal onsets that marks them both out as favourable neutralisation sites.

It is this unified licensing environment that hosts wholesale consonantal neutralisation of the sort which affects Danish laryngeal contrasts. As seen in (29b), the particular element involved in the Danish case is (H), defining plosive aspiration. While this element is supported in a foot-initial onset, a strong melodic licensor, its a-license is withheld under the conditions specified in (34). The effect, illustrated by the forms in (10b–d), is that aspiration cannot be sponsored by a foot-internal intervocalic onset or by a word-final consonant or by a coda.

In the case of the examples in (2) and (3), neutralisation takes place in coda and word-final position but not intervocalically. Melodically, each of these events involves a reduction in elemental complexity, as shown in (27) and (28a). In all instances, the primary trigger is prosodic: specifically, the a-licensing of particular elements is potentially withheld under the p-licensing conditions defined in (34). The melodic context defines a secondary condition: what unites the neutralising coda and domain-final positions and distinguishes them from the non-neutralising intervocalic position is the absence of a following melodically filled nucleus.

Word-final position provides the alternating consonantal site in which a lexically present element is contingently unlicensed (Goldsmith's 1989 term) in its onset. That is, any element involved in an alternation in this context is phonetically interpreted only if it receives an a-licence as a result of the requisite licensing conditions being made available by morphosyntactic concatenation; otherwise it is suppressed. In Spanish *s*-debuccalisation, for example, a word-final (R) forming part of a lexical expression defining [s] remains unlicensed as long as the following nucleus is empty. As shown in (27a), underparsing of (R) exposes the residual (h) element, independently realised as [h]; hence forms such as *las*∅ → *la*[h]∅ 'the (FEM PL)'. However, should the following nucleus come to be filled through the lexical insertion of a vowel-initial form, the (R) in question gains an a-licence and contributes to the interpretation of [s], as in *la*[s] *once* 'eleven o'clock'.

Coda position, being necessarily domain-internal (per Coda Licensing), is permanently shielded from the domain-edge conditions which trigger morphophonological alternation. Failure of an element to be licensed in this context will thus only manifest itself as a static distributional effect, unless some kind of stylistically conditioned variation is in play, as is typically the case with the Spanish lenitions in (2a–b) (see Amastae 1989). In *s*-debuccalisation, (R)'s a-licence is optionally revoked in this context to yield variants such as *co*[s]*ta*–*co*[h]*ta* 'cost'.

6 The licensing properties of nuclei

6.1 Internuclear relations

At the foot and word levels, dominant nuclei are identified as powerful a-licensors by virtue of the fact that they are unlicensed within their domain. A recessive nucleus, *qua* p-licensee, possesses correspondingly less a-licensing power. This is consistent with the recurrent pattern whereby a maximal inventory of vocalic contrasts manifests itself in dominant nuclei, while reduced inventories show up in recessive nuclei. System contraction of this general type is observable in, for example, Bergün Romansch (Kamprath 1987), Bulgarian (Pettersson & Wood 1987), Catalan (Palmada Félez 1991) and Chumash (Applegate 1971). Vowel reduction is expressible as the suppression of particular elements, which occurs when the inherited a-licensing power of the position they would otherwise be attached to is insufficient to ensure their interpretation.

Having sketched a hypothetical example of this scenario in §4.1, let us now consider a real case. As described by Bafile (1997), on whom the following account is based, Neapolitan Italian has a basic seven-vowel system. However, the full set of vocalic contrasts only shows up in tonic nuclei (typically penultimate in the word). In unstressed nuclei, the inventory shrinks to three or four vowels phrase-medially and further to two phrase-finally:

(35)									
	Tonic		i	e	ɛ	ɔ	o	u	a
	Atonic	phrase-medial		i/ə			u		a
		phrase-final				ə			a

The less drastic degree of contraction is evident in the following alternations, which allow us to compare tonic and pretonic contexts:

(36)	a.	[díʃə]	'he says'	[diʃítə]/[dəʃítə]	'say (IMP PL)'
	b.	[péʃkə]	'I fish'	[piʃkatórə]/[pəʃkatórə]	'fisherman'
		[sečča]	'cuttlefish'	[siččətɛ́lla]/[səččətɛ́lla]	'small cuttlefish'
	c.	[vékə]	'he sees'	[vərítə]	'see (IMP PL)'
		[tɛ́nə]	'he keeps'	[tənítə]	'keep (IMP PL)'
	d.	[kórrə]	'he runs'	[kurrítə]	'run (IMP PL)'
		[pɔ́rta]	'he brings'	[purtátə]	'bring (IMP PL)'
		[pɔ́kə]	'a little'	[pukuríllə]	'a very little'

Whether [i] retains its quality pretonically or centralises to [ə] is apparently a matter of free phonological variation (see (36a)). In this context, mid front vowels optionally raise to [i] only under secondary stress (i.e. when they occur in the dominant syllable of a non-final foot (see (36b)); otherwise they reduce to [ə] (see (36b–c)). Raising of pretonic back mid vowels is quite regular (see (36d)).

Expressed in terms of elements, the raising and centralisation that produce reduction converge on a single result: they banish complex melodic expressions from atonic nuclei. (For an element-based treatment of similar effects in English and Catalan, see Harris 1994: 108ff.) Of the vowels that occur under reduction, [i], [u] and [a] each embodies a single element ((I), (U) and (A) respectively). The other, [ə], is the phonetic interpretation of a nucleus stripped of all elemental content (see Schane 1984, Kaye *et al*. 1985, Harris & Lindsey 1995). Tonic nuclei, in contrast, are able to support compound expressions (mid vowels) in addition to simplex (high and low). Raising consists in the suppression of (A) from mid vowels, centralisation in the suppression of all melodic content. This is set out in (37), where (37b) records the melodic differences between tonic (37a) and pretonic (37c) nuclei. (Underlining identifies the element which acts as the head of its expression (cf. note 9). I assume that tense [e]/[o] are headed by (I)/(U) respectively, while lax [ɛ]/[ɔ] are headless, cf. Anderson & Jones (1974), Anderson & Ewen (1987) and, for a fully referenced discussion, Harris (1994: 108ff).)

(37) a. *Tonic*	i	e	ɛ	a	ɔ	o	u
b.	<$\underline{\mathrm{I}}$>	<A> <$\underline{\mathrm{I}}$>	<A> <I>	$\underline{\mathrm{A}}$	<A> U	<A> $\underline{\mathrm{U}}$	$\underline{\mathrm{U}}$
c. *Atonic*		ə		a		u	

Note that the correspondence between tonic and atonic vowels is represented elementally without the need for supplementary fill-in or linking devices of the sort that are required by any approach which represents raising and centralisation as the substitution of certain feature values by others. In the event of (A) being underparsed in [ɔ], for instance, the residual element (U) is independently available for phonetic interpretation as [u].

As Bafile (1997) shows, the distributional asymmetries between tonic and atonic nuclei in Neapolitan fall out from Licensing Inheritance. A tonic nucleus, being the ultimate licensor of the word domain, can dispense a-licensing favours with greater largesse than its atonic counterparts; it is free to support expressions comprising one element (defining [i a u]) or two ([e ɛ o ɔ]). Atonic nuclei cannot afford to be so generous; being p-licensed on either the foot or word projection, they possess correspondingly less a-licensing power and can sustain at most one element.

Against this background, alternations arise whenever the vowel of a morphological root finds itself in different forms with differing licensing configurations. In [tɛ́nə], for instance, both elements of the compound (A, I), defining the root vowel [ɛ], are supported by a nucleus which, by virtue of being head of both its foot and the word, is a strong a-licensor. In [tənítə], on the other hand, the root vowel occupies a nucleus in which a-licensing potential is diluted as a result of being inherited from some other source, specifically the head nucleus of the word (occupied by [i]). In this case, the a-licence for (A, I) is revoked; stripped of all elemental content, the nucleus is interpreted as [ə].

Although the contraction of vowel contrasts in prosodically recessive positions is widely attested across the world's languages, we only have to compare Neapolitan with standard Italian to see that the effect is not universally observed. The high–mid vowel sequence found in the final feet of standard Italian forms such as *fin*(*íre*) 'to finish' and *ven*(*úto*) 'come

(PAST PART)', for example, displays a configuration in which an atonic nucleus supports a more complex vocalic expression than its tonic p-licensor. The inevitable conclusion is that Licensing Inheritance, at least as formulated in (25), is violable – or perhaps just not equally enforceable in all phonological contexts. Whether this should be construed as a matter of parametric variability or as indicative of constraint ranking has no direct bearing on the validity of the Licensing Inheritance proposal. A systematic difference such as that between Neapolitan and standard Italian certainly lends itself to a treatment in terms of ranked constraints. According to this type of analysis, vowel reduction reflects a situation in which Licensing Inheritance overrides any pressure to maximise the correspondence between lexical input and phonological output. In Neapolitan, satisfaction of Licensing Inheritance forces the underparsing of, for example, (A, I) in the first nucleus of /tɛn-ítə/ = [tənítə]. In standard Italian, the opposite ranking holds; input (A, I) in the final nucleus of /fíne/ 'fine, thin', for instance, is parsed in output [fíne] only at the expense of violating Licensing Inheritance.

Returning to the main representational issue at hand, we may note that, under an element-based account, vowel reduction to schwa and vowel syncope are two sides of the same coin. Both involve a nucleus being drained of all melodic content. What distinguishes the two cases is whether this type of nucleus must nevertheless receive phonetic interpretation. Like many languages, Neapolitan requires all nuclei to be made phonetically manifest. Other languages, as argued in §2, allow empty nuclei to remain unexpressed under specific conditions. Those researchers who have explicitly adopted the notion of empty nuclei consider the difference to be a matter of parametric choice (Kaye 1990, Kaye *et al*. 1990, Charette 1991, Burzio 1994).

Vowel harmony too demonstrates the effects of Licensing Inheritance within internuclear domains, typically at the level of the word. The assimilatory nature of this type of event reflects a situation in which the phonetic interpretation of a harmonically recessive nucleus is partially or wholly determined by the melodic content of a harmonically dominant nucleus. Just as in the geminate-consonant case discussed in §5.2.2, this implies an asymmetry in a-licensing power: as sponsor of the assimilating melody, the p-licensing position supports at least one more element than p-licensed positions.

6.2 Nuclei as phonotactic islands

Given a completely free rein, Licensing Inheritance would be expected to impose melodic complexity differentials on any two positions that make up a licensing domain. As demonstrated above, this delivers the desired result in the case of complex onsets and coda–onset interludes (§5.2) and internuclear domains at the level of the foot and word (§6.1). This, however, is not an exhaustive inventory of domains in which pairs of positions stand in a direct licensing relation. The two remaining contexts are onset–nucleus and nucleus–coda sequences. Neither of these displays anything like the degree of distributional dependency associated with the contexts already reviewed. It is surely significant that these are the only two licensing domains which involve a relation between a nuclear and a non-nuclear position. The others, in contrast, involve relations between positions of like status – between non-nuclear positions (onset and interlude clusters) or between nuclear positions (foot or word).

There is, to be sure, a quantitative relation between a nucleus and a following coda, evidenced in weight-sensitivity and closed-syllable shortening. Moreover, some languages exhibit qualitative dependencies between a nucleus and a preceding onset, although these are typically limited to place contrasts, as in the palatalisation of CV sequences in Japanese, Polish (Gussmann & Kaye 1993) and Irish (Ní Chiosáin 1994, Cyran 1997). However, the general pattern is for the sets of qualitative contrasts in a nucleus and in an adjacent non-nuclear position to be largely independent of one another. Where nuclei are systematically subject to external distributional influences, these are most likely to emanate from other nuclei within the same foot or word domain, producing the effects of umlaut or harmony.

This distributional segregation is evidently linked to the broadly complementary nature of the roles played by nuclear and non-nuclear positions in phonological representation. Non-nuclear positions shoulder the main burden of qualitative contrast, localised to melodic associations with the skeletal level. This is reflected in the fact that a language's consonant inventory is typically larger than its vowel inventory. The primary role of nuclei is, in contrast, prosodic: by virtue of being projected to higher levels of phonological structure, they form the backbone of the prosodic hierarchy. The distinction between the two types of position happens to be most dramatically underlined in languages with templatic morphology, where non-nuclear positions bear lexical-category distinctions, while the vocalic content of nuclear positions is either restricted to affix material or completely predictable (McCarthy 1979, Kaye 1990).

The phonotactic isolation of nuclei leads to the conclusion that the writ of Licensing Inheritance must somehow be prevented from running to relations contracted between a nuclear head and any non-nuclear positions it licenses. One response might simply be to add a rider to Licensing Inheritance deflecting it from the licensing relations in question. Another might be to try to express more directly the competition between the melodic demands made by Licensing Inheritance and the prosodic responsibilities borne by nuclear heads – precisely the sort of tension that Optimality Theory is designed to capture. We might propose that Licensing Inheritance is universally dominated by some family of nuclear constraints (presumably including NUC itself). In this way, the melodic effects of Licensing Inheritance would only be visible in contexts where the constraint is not in direct conflict with the prosodic duties placed on nuclei. In so far as the ordering of constraints on any universal dominance hierarchy is purely stipulative, the ranked-constraint approach is hardly any less of a brute-force solution than a simple amendment to Licensing Inheritance. On the other hand, a precedent for the ranking analysis was set in §6.1 above, where it was acknowledged that otherwise expected a-licensing effects can sometimes be short-circuited at foot and word level.

6.3 The p-licensing properties of nuclei

The focus of this article has been on the melodic consequences of Licensing Inheritance. However, inherent in the proposal is the possibility of its being extended to the prosodic licensing potential of a position. That is, we might ask whether asymmetries exist in the ability of positions to p-license other positions. The question is of most relevance to nuclei, given their role as prosodic anchors. Several considerations point to this being a promising line of future enquiry.

A generalised version of Licensing Inheritance would lead us to expect the prosodic status of a nucleus to impact on its capacity to support branching constituent structure. One phenomenon which is consistent with this expectation is quantity-sensitivity in stress placement. Under the weight-to-stress principle, the dominant rhyme of a foot either can or must (depending on the language) contain branching structure, while the weak rhyme cannot (Hayes 1981, Prince 1990). We can think of this as the recessive nucleus within the foot being denied the degree of p-licensing power that enables the head nuclear position to support another position within its rhyme.

We might also expect headship asymmetries between nuclei to be reflected in an unequal capacity to sustain branching structure in adjacent non-nuclear constituents, for example in a preceding onset. A sequence in which a nucleus follows a complex onset contains a two-stage p-licensing path: the nuclear position licenses the initial onset head position, which in turn licenses its following complement. Any diminution in the p-licensing power of the nucleus in this potential configuration, such as might be suffered by the recessive member of a foot, would place the second of these stages in jeopardy. This sets up the following implicational universal: while there should exist languages in which branching onsets are supported in the dominant syllable of a foot but not in the weak syllable, no language should display precisely the opposite distribution. There are indeed languages which conform to the predicted pattern. In southeastern Brazilian Portuguese, for example, branching onsets are free to occur in stressed syllables (e.g. *prato* [prátu] 'plate', *livreto* [livrétu] 'small book'); in unstressed position, however, the complex clusters found in other dialects (e.g. *pratinho* [pračíɲu] 'small plate', *livro* [lívru] 'book') are simplified ([pačíɲu], [lívu]). A localised variation on this theme is found in German hypocoristics: the template for the truncated form consists of a trochaic foot in which the onset of the weak syllable must not branch. This gives rise to forms such as *Gabriella* > *Gabi* (**Gabri*), *Andreas* > *Andi* (**Andri*) (see Itô & Mester 1997 for references and discussion). It is quite unlikely that there exist languages in which complex onsets are restricted to weak syllables.

This asymmetry generalises to word-final consonant clusters. Since the empty nucleus supporting the onset occupied by a final consonant can form the weak member of a foot (as per the arguments in §2.2), it too is expected to exhibit diminished p-licensing potential. This is consistent with the observation that, while some languages permit branching onsets in this context (as in French *table* [tabl] 'table'), the majority of languages do not. To the best of my knowledge, there are no languages in which branching onsets only occur in this position. Charette (1992) treats this cross-linguistic difference in terms of a parameter which controls the ability of a final empty nucleus to license a preceding onset head to govern its complement. In referring to the licensing paths that underpin final .C(C)∅], this notion has much in common with Licensing Inheritance. In terms of the present proposal, we can say that the onset head position inherits its ability to p-license its complement from its nuclear p-licensor.

7 Conclusion

One of the problems we started out with was that of trying to unify the range of phonological contexts which favour consonantal neutralisation. Viewed in linear terms, the sites form a disparate collection of segmental and boundary conditions. Viewed in terms of syllabic constituency, they can be made to reduce to coda position. However, this can only be achieved at the expense of invoking resyllabification devices of various sorts and riding roughshod over a significant body of facts which demonstrate the consistently non-rhymal behaviour of domain-final consonants.

Here I have presented a third alternative, which builds on the syllabic approach by accessing the language in which constituency is compiled, namely licensing. This allows us to forge a direct link between a syllabic position's ability to support melodic material and its place within the prosodic hierarchy. The approach dispenses with resyllabification and respects the extrarhymal status of final consonants. Moreover, it subsumes the melodic and contextual dimensions of consonantal weakening under a general theory of neutralisation which also covers sonority sequencing and vocalic reduction.

Under the Licensing Inheritance proposal, melody-bearing potential is transmitted from licensing to licensed positions and is progressively diluted the further from its point of origin it is discharged. The asymmetry that is inherent in this mechanism is responsible for variations in the complexity of melodic units that can appear in various positions. Neutralisation consists in a capping of the melodic complexity levels permitted in particular positions. What unifies the sites where this occurs is their location at some distance from the source which ultimately licenses their melodic content.

NOTES

* An early cut of this article was presented at the 20th GLOW Colloquium (Lisbon, April 1992) and appeared in *UCL Working Papers in Linguistics* 4 (1992). In preparing the present version, I have benefited greatly from an exchange of ideas with Wiebke Brockhaus. My thanks go to her as well as to the following for their valuable comments on earlier drafts: Phillip Backley, Phil Carr, Edmund Gussmann, Keren Rice, Leonardo Savoia, Toyomi Takahashi, Eno-Abasi Urua, Geoff Williams and three anonymous *Phonology* reviewers. Thanks also to the following people for supplying and discussing some of the language data presented here: Fritz Larsen and Jørgen Staun (Danish), Eno Urua (Ibibio), Yilmaz Vural (Turkish) and Thaïs da Silva (Brazilian Portuguese).

[1] Tones, which are not relevant to the weakening process, are omitted in all Ibibio forms.

[2] The question of whether there might be some phonotactic resemblance between final CC] and putative medial coda clusters in English does not arise, since most of the evidence supports the conclusion that internal codas can contain no more than one consonant anyway (see Harris 1994: 66ff for a summary).

[3] The possibility of a consonant appearing in a coda in one form of a morpheme and in an onset in another is not excluded by an approach which dispenses with resyllabification. For example, the [p] of *descriptive* clearly closes a syllable (hence the shortness of the preceding vowel). In contrast, as argued in §2.2, the [b] of *describe* occupies the onset of a null-vowelled syllable (where it can have no influence on the length of the preceding vowel). Although the two forms are part of the same morphological root paradigm, we are not obliged to consider them related via phonological derivation – especially since there are good phonological reasons for assuming that English root-level morphology is lexically listed (see Kaye 1995). Under this approach, there is thus no sense in which the root-final labial stop is phonologically transferred from one syllabic position to another.

[4] It might be objected that this analysis builds an undesirable amount of predictable information into representations: an empty nucleus is posited underlyingly precisely where an inappropriate sonority differential between otherwise contiguous consonants would result. This criticism reveals a concern with the shape of input representations that has no place in output-oriented theory. A certain proportion of the autosegmental associations in phonological output inevitably reflect predictable relations between melodic units and their syllabic affiliation – regularities which must be captured by output constraints, including those referring to sonority or its equivalent. It is irrelevant to these constraints whether a directionality inheres in such relations, such that the syllabification of a string might be derived from its sonority profile or *vice versa*. Seeking to establish directionalities of this sort is not only inimical to the spirit of output-oriented theory but is also in any event typically circular, since the dependencies in question can just as easily be considered mutual (a point acknowledged by Levin 1985 and Borowsky 1986, for example).

[5] I adopt an x-slot representation of the skeletal tier because it presents a single level over which it is possible to define not only quantitative relations but also, crucially for the treatment of neutralisation, phonotactic dependencies between adjacent positions. This is in preference to moraic theory, within which the statement of phonotactic regularities requires a hybrid scansion of moras (for rhymal segments) and feature-geometric root nodes (for onset segments).

[6] I owe this metaphor to Jonathan Kaye (*voce*).

[7] This argumentation applies only to the melodic content of representations. Output representations inevitably contain a certain amount of prosodically derivable information which is predictable and therefore, from a phonemic-contrastive perspective, redundant. This can include details pertaining to word stress, foot structure and segmental precedence determined by syllabic affiliation. Whether this type of information should be omitted from lexical representation has no bearing on the issue at stake here.

[8] Any additional advantages that zero segmental redundancy might appear to enjoy in the matter of lexical storage are probably spurious and are in any event irrelevant to the current proposal.

[9] Refining the definition of the labial fricative requires reference to the headship of the relevant elemental expression. With (h) as head, an obstruent is strident. A labial fricative is thus [f] when headed by (h) and [ʍ] when not. On the headed nature of privatively represented melodic expressions, see Anderson & Ewen (1987), van der Hulst (1989), Kaye *et al*. (1985).

[10] Other lenitions that can affect English [t] submit to the same general element-based treatments as the events represented in (26); see Harris (1990, 1994: ch. 3).

[11] The alphabetic transcription of the plosives in (30) is that conventionally used for English. [b] indicates a plain stop (as does [p] in [sp] below); [t] is voiceless aspirated.

[12] In order to sharpen the definition of the interlude context, it is perhaps worth mentioning an assimilation site that is not under discussion here, namely one where two consonants are separated by a phonological domain boundary. This too constitutes an interlude under the traditional final-coda view, but not under the assumption that a domain-final consonant occupies the onset of a null-vowelled syllable. Under the second view, the two consonants are not contiguous on the skeletal tier, since an empty nucleus intervenes, *viz*. … C∅]C … (stem + suffix, as in [[*lip*]*s*]) or … C∅][C … (word + word, as in [[*lip*][*salve*]]). Where assimilatory interactions occur in this context, they thus necessarily involve melodic adjacency on the relevant autosegmental tier ('minimal scansion' in Archangeli & Pulleyblank's 1994 terms). Since this configuration is quite distinct from that found in a (domain-internal) interlude, there is no reason to expect the same type of distributional asymmetries to occur there. While the contrastive potential of interludes is clearly skewed in favour of the righthand position, no such consistent directionality is associated with the cross-domain context, where perseverative as well as anticipatory patterns of assimilation are attested. Compare, for example, Ibibio [dip]-CV → [dippe] 'lift up' (Urua 1990) with Selayarese [ta]C-[pelaʔ] → [tappelaʔ] 'get lost' (Goldsmith 1990: 133).

[13] As noted in §2.2, languages vary according to whether a final null-vowelled syllable is metrified as the weak member of a trochaic foot (as in Spanish) or not (as in Polish). English accommodates both possibilities, which raises the question of what happens to [t] when the onset it occupies lies outside the trochaic frame, as in *edit∅*, *covet∅*, etc. Lenition occurs here no less than in *sit∅*, *bet∅*, etc. From the viewpoint of the present account it matters not whether we regard the final syllable of *edit∅* as unfooted (cf. Hayes 1995) or as the rightmost member of a ternary foot (cf. Burzio 1994). (In the former instance, the silent nucleus would presumably have to be p-licensed at the word level.) In both cases, the onset in question falls outside the head syllable of the foot, with the result that it fails to acquire the strong a-licence which would protect it from lenition.

REFERENCES

Akinlabi, A. & E. Urua (1992). Prosodic target and vocalic specification in the Ibibio verb. *WCCFL* **11**. 1–14.

Amastae, J. (1989). The intersection of *s*-aspiration/deletion and spirantization in Honduran Spanish. *Language Variation and Change* **1**. 169–183.

Anderson, J. M. & C. J. Ewen (1987). *Principles of dependency phonology*. Cambridge: Cambridge University Press.

Anderson, J. M. & C. Jones (1974). Three theses concerning phonological representations. *JL* **10**. 1–26.

Anderson, J. M. & C. Jones (1977). *Phonological structure and the history of English*. Amsterdam: North-Holland.

Applegate, R. (1971). Vowel harmony in Chumash. *Berkeley Papers in Linguistics* **1**. 3–12.

Archangeli, D. & D. Pulleyblank (1994). *Grounded phonology*. Cambridge, Mass.: MIT Press.

Backley, P. (1995). A tier geometry for vowel systems. *UCL Working Papers in Linguistics* **7**. 399–436.

Bafile, L. (1997). L'innalzamento vocalico in napoletano: un caso di interazione fra fonologia e morfologia. In L. Agostiniani (ed.) *Atti del III Convegno Internazionale della Società Internazionale di Linguistica e Filologia Italiana*. Naples: Edizioni Scientifiche Italiane. 1–22.

Blevins, J. (1995). The syllable in phonological theory. In J. A. Goldsmith (ed.) *The handbook of phonological theory*. Cambridge, Mass. & Oxford: Blackwell. 206–244.

Borowsky, T. J. (1984). On resyllabification in English. *WCCFL* **3**. 1–15.

Borowsky, T. J. (1986). *Topics in the lexical phonology of English*. PhD dissertation, University of Massachusetts, Amherst. Distributed by Indiana University Linguistics Club. Published 1990, New York: Garland.

Brockhaus, W. G. (1992). *Final devoicing: principles and parameters*. PhD dissertation, University College London.

Brockhaus, W. G. (1995). *Final devoicing in the phonology of German*. Tübingen: Niemeyer.

Burzio, L. (1994). *Principles of English stress*. Cambridge: Cambridge University Press.

Charette, M. (1991). *Conditions on phonological government*. Cambridge: Cambridge University Press.

Charette, M. (1992). Mongolian and Polish meet Government Licensing. *SOAS Working Papers in Linguistics and Phonetics* **2**. 275–291.

Clements, G. N. (1990). The role of the sonority cycle in core syllabification. In J. Kingston & M. Beckman (eds.) *Papers in laboratory phonology I: between the grammar and physics of speech*. Cambridge: Cambridge University Press. 283–333.

Clements, G. N. & S. J. Keyser (1983). *CV phonology: a generative theory of the syllable*. Cambridge, Mass.: MIT Press.

Cole, J. & C. Kisseberth (1994). An optimal domains theory of harmony. *Studies in the Linguistic Sciences* **24**. 101–114.

Cyran, E. (1997). *Resonance elements in phonology: a study in Munster Irish*. Lublin: Wydawnicwo Folium.

Cyran, E. & E. Gussmann (in press). Consonant clusters and governing relations: Polish initial consonant sequences. To appear in H. van der Hulst & N. Ritter (eds.) *The syllable*. Berlin: Mouton de Gruyter.

Dell, F. (1976). Schwa précédé d'un groupe obstruante–liquide. *Recherches Linguistiques* **4**. 75–111.

Durand, J. & F. Katamba (eds.) (1995). *Frontiers of phonology: atoms, structures, derivations*. London: Longman.

Giannelli, L. & L. Savoia (1980). L'indebolimento consonantico in Toscana (II). *Revista Italiana di Dialettologia* **4**. 39–101.

Giegerich, H. J. (1985). *Metrical phonology and phonological structure*. Cambridge: Cambridge University Press.

Goldsmith, J. A. (1985). Vowel harmony in Khalkha Mongolian, Yaka, Finnish and Hungarian. *Phonology Yearbook* **2**. 253–275.

Goldsmith, J. A. (1989). Licensing, inalterability, and harmonic rule application. *CLS* **25:1**. 145–156.

Goldsmith, J. A. (1990). *Autosegmental and metrical phonology*. Oxford: Blackwell.

Gussenhoven, C. (1986). English plosive allophones and ambisyllabicity. *Gramma* **10**. 119–141.

Gussmann, E. & J. Kaye (1993). Polish notes from a Dubrovnik café. *SOAS Working Papers in Linguistics and Phonetics* **3**. 427–462.

Hale, K. (1973). Deep–surface canonical disparities in relation to analysis and change:

an Australian example. In T. A. Sebeok (ed.) *Current trends in linguistics*. Vol. 11. The Hague: Mouton. 401–458.
Halle, M. & W. J. Idsardi (1997). *r*, hypercorrection, and the Elsewhere Condition. In I. Roca (ed.) *Derivations and constraints in phonology*. Oxford: Oxford University Press. 331–348.
Harris, James W. (1969). *Spanish phonology*. Cambridge, Mass.: MIT Press.
Harris, James W. (1983). *Syllable structure and stress in Spanish*. Cambridge, Mass.: MIT Press.
Harris, John (1990). Segmental complexity and phonological government. *Phonology* **7**. 255–300.
Harris, John (1994). *English sound structure*. Oxford: Blackwell.
Harris, John (1996). Phonological output is redundancy-free and fully interpretable. In J. Durand & B. Laks (eds.) *Current trends in phonology: models and methods*. Vol. 1. Salford: University of Salford Publications. 305–332.
Harris, John & J. Kaye (1990). A tale of two cities: London glottalling and New York city tapping. *The Linguistic Review* **7**. 251–274.
Harris, John & G. Lindsey (1995). The elements of phonological representation. In Durand & Katamba (1995). 34–79.
Hayes, B. (1981). *A metrical theory of stress rules*. PhD dissertation, MIT. Published 1985, New York: Garland.
Hayes, B. (1982). Extrametricality and English stress. *LI* **13**. 227–276.
Hayes, B. (1995). *Metrical stress theory: principles and case studies*. Chicago: University of Chicago Press.
Hulst, H. van der (1989). Atoms of segmental structure: components, gestures and dependency. *Phonology* **6**. 253–284.
Hulst, H. van der & N. Smith (1985). Vowel features and umlaut in Djingili, Nyangumarda and Warlpiri. *Phonology Yearbook* **2**. 277–303.
Hyman, L. M. (1975). *Phonology: theory and analysis*. New York: Holt, Rinehart & Winston.
Inkelas, S. (1996). The interaction of phrase and word rules in Turkish: an apparent paradox in the prosodic hierarchy. *The Linguistic Review* **13**. 193–217.
Itô, J. (1986). *Syllable theory in prosodic phonology*. PhD dissertation, University of Massachusetts. Distributed by Indiana University Linguistics Club. Published 1988, New York: Garland.
Itô, J. & R. A. Mester (1993). Licensed segments and safe paths. *Revue Canadienne de Linguistique* **38**. 197–213.
Itô, J. & R. A. Mester (1997). Sympathy Theory and German truncations. Available as ROA-211 from the Rutgers Optimality Archive.
Itô, J., R. A. Mester & J. Padgett (1995). Licensing and redundancy: underspecification in Optimality Theory. *LI* **26**. 571–613.
Jensen, J. T. (1993). *English phonology*. Amsterdam: Benjamins.
Jespersen, O. (1904). *Lehrbuch der Phonetik*. Leipzig & Berlin: Teubner.
Jessen, M. (1997). *Tense vs. lax obstruents in German. Arbeitspapiere des Instituts für Maschinelle Sprachverarbeitung* 3, Universität Stuttgart.
Ka, O. (1994). *Wolof phonology and morphology*. Lanham, MD: University Press of America.
Kahn, D. (1976). *Syllable-based generalizations in English phonology*. PhD dissertation, MIT. Distributed by Indiana University Linguistics Club. Published 1980, New York: Garland.
Kamprath, C. K. (1987). *Suprasegmental structures in a Räto-Romansch dialect: a case study in metrical and lexical phonology*. PhD dissertation, University of Texas, Austin.
Kaye, J. (1990). 'Coda' licensing. *Phonology* **7**. 301–330.
Kaye, J. (1995). Derivations and interfaces. In Durand & Katamba (1995). 289–332.

Kaye, J. (1996). Do you believe in magic? The story of $s+$C sequences. In H. Kardela & B. Szymanek (eds.) *A Festschrift for Edmund Gussmann*. Lublin: University Press of the Catholic University of Lublin. 155–176.

Kaye, J., J. Lowenstamm & J.-R. Vergnaud (1985). The internal structure of phonological elements: a theory of charm and government. *Phonology Yearbook* **2**. 305–328.

Kaye, J., J. Lowenstamm & J.-R. Vergnaud (1990). Constituent structure and government in phonology. *Phonology* **7**. 193–231.

Kenstowicz, M. (1994). *Phonology in generative grammar*. Oxford: Blackwell.

Kiparsky, P. (1979). Metrical structure assignment is cyclic. *LI* **10**. 421–442.

Kiparsky, P. (1991). Catalexis. Paper presented at the EUROTYP workshop, Salzburg.

Levin, J. (1985). *A metrical theory of syllabicity*. PhD dissertation, MIT.

Liberman, M. & A. Prince (1977). On stress and linguistic rhythm. *LI* **8**. 249–336.

Lombardi, L. (1995). Laryngeal neutralization and syllable wellformedness. *NLLT* **13**. 39–74.

McCarthy, J. J. (1979). *Formal problems in Semitic phonology and morphology*. PhD dissertation, MIT. Distributed by Indiana University Linguistics Club. Published 1985, New York: Garland.

McCarthy, J. J. & A. S. Prince (1986). *Prosodic morphology*. Ms, University of Massachusetts, Amherst & Brandeis University.

Maddieson, I. (1984). *Patterns of sounds*. Cambridge: Cambridge University Press.

Maris, M. Yunus (1980). *The Malay sound system*. Kuala Lumpur: Penerbit Fajar Bakti.

Marotta, G. & L. M. Savoia (1994). Vowel properties and nuclear constituents: evidence from Italian dialects. *Probus* **6**. 43–79.

Mithun, M. & H. Basri (1986). The phonology of Selayarese. *Oceanic Linguistics* **25**. 210–254.

Myers, S. (1987). Vowel shortening in English. *NLLT* **5**. 485–518.

Myers, S. (1991). Persistent rules. *LI* **22**. 315–344.

Nespor, M. & I. Vogel (1986). *Prosodic phonology*. Dordrecht: Foris.

Nasukawa, K. (1995). Nasality and harmony in Gokana. *UCL Working Papers in Linguistics* **7**. 511–533.

Palmada Félez, B. (1991). *La fonologia del català i els principis actius*. PhD dissertation, Universitat Autònoma de Barcelona.

Ní Chiosáin, M. (1994). Irish palatalisation and the representation of place features. *Phonology* **11**. 89–106.

Pettersson, T. & S. Wood (1987). Vowel reduction in Bulgarian and its implications for theories of vowel reduction: a review of the problem. *Folia Linguistica* **21**. 261–279.

Piggott, G. L. (1993). Phonological government and segmental licensing. Ms, McGill University.

Prentice, D. J. (1971). *The Murut languages of Sabah*. Canberra: Australian National University.

Prince, A. S. (1984). Phonology with tiers. In M. Aronoff & R. T. Oehrle (eds.) *Language sound structure*. Cambridge, Mass.: MIT Press. 234–244.

Prince, A. S. (1990). Quantitative consequences of rhythmic organization. *CLS* **26:2**. 355–398.

Prince, A. S. & P. Smolensky (1993). *Optimality Theory: constraint interaction in generative grammar*. Ms, Rutgers University & University of Colorado, Boulder.

Rennison, J. R. (1981). *Bidialektale Phonologie: die Kompetenz zweier Salzburger Sprecher*. Wiesbaden: Franz Steiner.

Rice, K. D. (1992). On deriving sonority: a structural account of sonority relationships. *Phonology* **9**. 61–99.

Rice, K. D. (1996). Default variability: the coronal–velar relationship. *NLLT* **14**. 493–543.
Roca, I. (1994). *Generative phonology*. London: Routledge.
Rubach, J. (1990). Final devoicing and cyclic syllabification in German. *LI* **21**. 79–94.
Rubach, J. (1996). Shortening and ambisyllabicity in English. *Phonology* **13**. 197–237.
Savoia, L. (1994). Parametrical variation in prosodic licensing: evidence from Italo-Albanian dialects. *Rivisti di Grammatica Generativa* **19**. 105–141.
Schane, S. A. (1984). The fundamentals of particle phonology. *Phonology Yearbook* **1**. 129–155.
Segundo, S. de Oliveira (1990). On empty nuclei and empty structure in Brazilian Portuguese. *SOAS Working Papers in Linguistics and Phonetics* **1**. 39–51.
Selkirk, E. O. (1981). On prosodic structure and its relation to syntactic structure. In T. Fretheim (ed.) *Nordic Prosody II*. Trondheim: TAPIR. 111–140.
Selkirk, E. O. (1982). The syllable. In H. van der Hulst & N. Smith (eds.) *The structure of phonological representations*. Part 2. Dordrecht: Foris. 337–383.
Smith, N. (1988). Consonant place features. In H. van der Hulst & N. Smith (eds.) *Features, segmental structure and harmony processes*. Part 1. Dordrecht: Foris. 209–235.
Spencer, A. (1996). *Phonology: theory and description*. Oxford: Blackwell.
Steriade, D. (1982). *Greek prosodies and the nature of syllabification*. PhD dissertation, MIT. Distributed by Indiana University Linguistics Club. Published 1990, New York: Garland.
Takahashi, T. (1993). A farewell to constituency. *UCL Working Papers in Linguistics* **5**. 375–410.
Trubetzkoy, N. S. (1939). *Grundzüge der Phonologie*. Göttingen: Vandenhoek & Ruprecht.
Twaddell, W. F. (1935). On defining the phoneme. *Language Monograph* **16**. Reprinted in M. Joos (ed.) (1957). *Readings in linguistics I*. Chicago: University of Chicago Press. 55–80.
Urua, E. (1990). *Aspects of Ibibio phonology and morphology*. PhD dissertation, University of Ibadan.
Vennemann, T. (1972). On the theory of syllabic phonology. *Linguistische Berichte* **18**. 1–18.
Vennemann, T. (1988). *Preference laws for syllable structure and the explanation of sound change*. Berlin: Mouton de Gruyter.
Wells, J. C. (1982). *Accents of English*. Vol. 2: *The British Isles*. Cambridge: Cambridge University Press.
Wells, J. C. (1990). Syllabification and allophony. In S. Ramsaran (ed.) *Studies in the pronunciation of English: a commemorative volume in honour of A. C. Gimson*. London: Routledge. 76–86.
Yip, M. (1991). Coronals, consonant clusters, and the coda condition. In C. Paradis & J.-F. Prunet (eds.) *The special status of coronals, internal and external evidence*. San Diego: Academic Press. 61–78.