FLORIAN BREIT

FORMAL ASPECTS OF ELEMENT THEORY

FORMAL ASPECTS OF ELEMENT THEORY

FLORIAN BREIT

A dissertation submitted in partial fulfilment
of the requirements for the degree of
Master of Research
in Linguistics with specialisation in Phonology

Department of Linguistics
University College London

September 2013

SUPERVISOR: Prof. John Harris

WORD COUNT: 11,983

This document was typeset using the typographical look-and-feel `classicthesis` for LyX by André Miede. The type faces used are *Palatino* and *Euler* by Hermann Zapf, and *Bera Mono* based on Bitstream, Inc.'s *Bitstream Vera*.

"Ambulavitque cum Deo, et non apparuit:
quia tulit eum Deus."

(Gn 5:24)

In memoriam *Jacint* et *Magdalena*.

† Iunio 2013

ABSTRACT

As one of the main contenders of the theory of distinctive features, Element Theory (ET) has seen much change over the last decades. While this has vastly advanced its capability and accuracy as a theory of subsegmental phonology, on-going development has sometimes come at the price of explicitness in definition of the precise model assumed and consequently there are some unclarities and apparent contradictions in some of the current proposals of ET.

This dissertation first gives an outline of the current state of ET in light its historical development, highlighting the ways in which it differs from Feature Theory, Autosegmental Phonology and Government Phonology. The position that it functions as a relatively independent model of subsegmental representation is advanced.

On the basis of this, the dissertation proposes a concrete formalisation of ET, ground in basic mathematical set theory. It is argued that, somewhat analogous to the set-representation of syntactic treelets as ordered pairs $\{\alpha, \{\alpha, \beta\}\}$, segmental representations can be seen as partially ordered sets of the type $\{\{\alpha\}, \{\alpha, \beta, \gamma, \ldots\}\}$. It is illustrated how such sets can be used as the basis to formalise the model of ET set out here, including aspects such as composition, decomposition, well-formedness, and geometry.

Finally, based on the work of Reiss (2012), this model is used to compare the generative capacity and power of ET to that of classical feature theory. It is argued that a distinction needs to be made between generative capacity and generative power, and that the desirability or undesirability of overgeneration/powerfulness is more fine-grained and differs between the individual aspects of subsegmental phonology, specifically the set of all possible segments, the set of all possible inventories and the set of all natural classes. It is shown that ET is, converse to common assumption, actually more powerful than feature theory but its concrete capacity is reduced by the relatively small number of primes assumed.

ACKNOWLEDGMENTS

First and foremost I must express my gratefulness to my supervisor, John Harris, whose help, guidance and knowledge have been invaluable not only in the lead up to this work but have helped shape my overall conception of phonology over the past year. Special thanks for much guidance and encouragement are also due to Andrew Nevins, and to Marco Tamburelli who introduced me to Element Theory and encouraged me to follow up on my early ideas for the work presented in this dissertation by taking this MRes.

Collectively as a set (and thus in no particular order), I would also like to thank Charles Reis, John Coleman, Tobias Scheer, the audiences at both the 2012 Workshop on Melodic Representation at SOAS/UCL and APAP 2013 in Lublin, Poland and many of my classmates who have all helped me by providing valuable criticism, encouragement and pointers to other resources when they were confronted with earlier versions of the work contained herein. I must also acknowledge the contribution of Robert Wall's excellent (1972) Introduction to Mathematical Linguistics — had I not stumbled across this by chance, none of this would likely have happened.

Finally, I gratefully acknowledge an AHRC Research Preparation Masters Studentship and the continued support of my parents Eva and Heinrich, without either of which this would not have been possible.

> When I first heard his name, I said, just as you are going to say, "But I thought he was a boy?"
> "So did I," said Christopher Robin.
> "Then you can't call him Winnie?"
> "I don't."
> "But you said —"
> "He's Winnie-ther-Pooh. Don't you know what 'ther' means?"
> "Ah, yes, now I do," I said quickly; and I hope you do too, because it is all the explanation you are going to get.
>
> — Winnie-the-Pooh, A. A. Milne

CONTENTS

1 INTRODUCTION 1
2 ELEMENT THEORY 3
  2.1 Introduction 3
  2.2 Segmental Composition 3
  2.3 Elements, Compounds and Interpretation 5
  2.4 The Elements 7
  2.5 Government Phonology 9
  2.6 Element Calculus 11
  2.7 Composition and Decomposition 12
  2.8 Element Geometry 15
  2.9 Summary 16
3 A SET-THEORETIC MODEL OF ET 17
  3.1 Introduction 17
  3.2 Representations and Derivations 18
  3.3 The Vocabulary 20
  3.4 Segments 21
  3.5 The Set of all Segments 23
  3.6 Heads, Complements and Dependents 25
  3.7 Operations: Compose and Decompose 28
  3.8 Element Geometry 31
  3.9 Element Antagonism 33
  3.10 Summary 34
4 THE GENERATIVE CAPACITY OF ELEMENT THEORY 37
  4.1 Introduction 37
  4.2 Generative Capacity and Generative Power 37
  4.3 Variation and General Considerations 39
  4.4 Segments and Inventories of Segments 40
  4.5 Natural Classes 43
  4.6 Possible Segments 44
  4.7 Possible Inventories 46
  4.8 Possible Natural Classes 48
  4.9 Discussion 48
5 CONCLUSION 51

BIBLIOGRAPHY 53

LIST OF FIGURES

Figure 1 Spectral patterns of $|\mathrm{A}|$, $|\mathrm{I}|$, and $|\mathrm{U}|$: mAss, dIp, and rUmp. 6
Figure 2 Spectral patterns of the compounds $|\mathrm{A},\mathrm{I}|$ and $|\mathrm{A},\mathrm{U}|$. 6
Figure 3 Distribution of phonemic inventory sizes in UPSID. 42
Figure 4 Plot of $|\Sigma|$ for vocabulary sizes from 0 to 20, y-axis in millions. 46
Figure 5 Plot of $\log_{10}(|\Sigma|)$ for vocabulary sizes from 0 to 50. 46
Figure 6 Plot of $\log_{10}(|I|)$ for vocabulary sizes from 0 to 10. 47

LIST OF TABLES

Table 1 The six most commonly assumed elements of Element Theory (ET). 8
Table 2 Some further elements that have been proposed. 8
Table 3 Some values of $|\Sigma|$ for different sizes of V. 45
Table 4 Some values of $|I|$ for different sizes of V. 47

MATHEMATICAL SYMBOLS

$=$ Equality: $x = y$ means x equals y.

$\equiv$ Equivalence: $x \equiv y$ means x is equivalent to y.

$\doteq$ Definition: $x \doteq y$ means x is defined as y, or x is equivalent to y by definition.

$<$ Smaller than: $x < y$ means x is smaller than y.

$>$ Greater than: $x > y$ means x is greater than y.

$\leq$ Smaller or equal to: $x \leq y$ means x is smaller or equal to y; analogously $x \geq y$ means x is greater or equal to y.

$\neg$ Negation: $\neg p$ means '*not* p'. Also written as $/$ through a symbol, e. g. $\neq$ means 'does not equal', $\notin$ means 'is not a member of', &c.

$\wedge$ Conjunction, the logical connector 'and', $q \wedge p$.

$\vee$ Disjunction, the logical connector 'or', $q \vee p$.

$\veebar$ Exclusive disjunction, $q \veebar p$ means either p or q but not both.

$\exists$ Existential quantification: $(\exists x)(p)$ means there exists at least one x such that p.

$\forall$ Universal quantification: $(\forall x)(p)$ means p for all x.

$\rightarrow$ Implication: $p \rightarrow q$ means p implies q; but also rewriting: $x \rightarrow y$ means x is rewritten as y.

$\mapsto$ Mapping: $x \mapsto y$ means x maps to y.

$\setminus$ Subtraction of sets. $X \setminus Y$ is the set subtraction of Y from X.

$\cup$ $X \cup Y$ is the union of X and Y, $\bigcup X$ is the arbitrary union of X (i. e. the union of all the subsets of X).

$\cap$ $X \cap Y$ is the intersection of X and Y, $\bigcap X$ is the arbitrary intersection of X (i. e. the intersection of all the subsets of X).

$\in$ Set membership: $x \in X$ means x is a member of the set X.

$\subseteq$ Subset: $X \subseteq Y$ means X is a subset of Y (X and Y may be identical).

$\subset$ Proper subset: $X \subset Y$ means X is a proper subset of Y (X and Y may *not* be identical).

$\wp(X)$ The powerset of X, i. e. the set of all subsets of X.

$|X|$ The cardinality (number of subsets) of X. But note that $|\mathrm{A}, \mathrm{B}, \ldots|$ is used as a notation for element expressions.

$\{x, y, z, \ldots\}$ The set containing x, y, z, &c.

$\{x : p\}$ Set-builder notation: The set of all x such that p is true.

$\langle x, y\rangle$ The ordered pair where x precedes y.

$\varnothing$ The empty set $\{\}$.

$\bar{X}$ Indicates a set which is in some way related to a set X. Note that this does *not* indicate the complement of the set X.

$\times$ Multiplication: $x \times y$ is the product of x and y. This symbol is also used for timing slots in autosegmental representations.

$-$ Subtraction of numbers: $x - y$ is the difference between x and y.

$+$ Addition of numbers: $x + y$ is the sum of x and y.

$\circ$ The compounding function in element calculus, from Kaye et al. (1985).

$\square$ *Q.E.D.*, indicates the end of a proof.

ACRONYMS

AP Autosegmental Phonology

EC Element Calculus

ECP Empty Category Principle

ERP Empty Representation Principle

ET Element Theory

FT Feature Theory

GenP Generative Power

GenC Generative Capacity

GP Government Phonology

IIP Independent Interpretability Principle

IP Isomericity Principle

SGC Strong Generative Capacity

SOHC Single Optional Headedness Condition

SPE The Sound Pattern of English

SPP Strict Privativity Principle

UPSID UCLA Phonological Segment Inventory Database

UG Universal Grammar

UT Underspecification Theory

1SG First person singular

Simplify, simplify.

— Henry David Thoreau

1 INTRODUCTION

Much of the early work on linear models of phonology, such as Jakobson et al. (1952), Jakobson and Halle (1956) and especially Chomsky and Halle's (1968) The Sound Pattern of English (SPE), have assumed that the primary objects manipulated by the phonological component are bundles of binary features. These features covered not only categorial distinctions related to the phonetic form of the output, but also aspects such as stress, syllabicity, boundaries, &c. As such they made no explicit distinction between prosodic and melodic (i.e. segmental) information. Later research (e.g. Goldsmith, 1976; Halle and Vergnaud, 1982) however showed the need to formally distinguish these two levels of phonological representation. This led to the development of nonlinear models such as Autosegmental Phonology (AP, Goldsmith, 1976) which position prosodic and melodic information at different levels of representation and link them non-monotonically (cf. Harris, 2007).

One particular outgrowth of such nonlinear models is Government Phonology (GP, Kaye et al., 1985), which also developed a unique approach to melodic representation. In GP, the melodic primes manipulated by the phonology were assumed to be larger and further abstracted from articulatory mechanics than the SPE-style feature representations otherwise commonly assumed. Kaye et al. (1985) proposed that melodic representations consist of *elements* arranged on their own tiers below the skeletal tier. It was then assumed that an operation known as Element Calculus (EC) would convert these melodic representations into matrices of binary features that could be interpreted phonetically.

Research following this line of inquiry has however since shifted away from the idea that there may be different levels of phonological and phonetic interpretations linked by an EC, and instead have taken the stance that elements are the only level of melodic representations and can be mapped directly into the acoustic signal (e.g. Harris and Lindsey, 1993). As research on this paradigm of melodic representation, now principally known as ET, has progressed over the last decades, many assumptions have changed and a number of variations of ET have appeared.

A matter that has however often been neglected in the ongoing development of ET is that of its formalisation: while we have a formal account of AP with Kornai (1995), no notable attempts have been made to formally define the properties of current versions of ET. Of course a solid theory requires its framework to be defined as precisely as possible, in order to make precise predictions and avoid inconsistencies or vagueness. Additionally, a formal characterisation of a theory is what is required as the basis for comparison of different theories of melodic representation and a pre-requirement for efforts such as computational modelling.

In this dissertation, I will develop such a formal characterisation of ET grounded in the mathematical theory of sets and demonstrate how this can be used to compare the properties of different theories of melodic representation: Feature Theory (FT), Underspecification Theory (UT) and ET. For this I will first give a brief overview of the version of ET which I will assume as the basis for my formalism in the next chapter. This will be followed by a chapter in which I discuss the implementation of this model in set-theoretic terms, discussing different possibilities of definition where appropriate and giving evidence for choosing one over another definition where this is possible. In the subsequent chapter I present a study based on Reiss (2012) which compares the Generative Capacity (GenC) and Generative Power (GenP) of ET to that of FT and UT. To conclude, I will give a brief summary of what was achieved by the work presented in this dissertation and how this may foster further inquiry. I will assume that the reader is familiar with the basic assumptions and constructs of mainstream FT, UT, AP and feature geometry. Concerning the mathematical propositions in this dissertation, a basic understanding of elementary set theory and first order predicate logic is assumed[1].

1 The material presented in the first few chapters of Partee et al. (1990) goes far beyond what is required; a much more concise alternative is Houston (2009).

2 ELEMENT THEORY

2.1 INTRODUCTION

Element Theory (ET, Kaye et al., 1985; Kaye and Harris, 1990; Harris, 1994; Harris and Lindsey, 1993, 1995a; Backley, 2011) is an alternative model of segmental representation that has developed as part of Government Phonology (GP, Kaye et al., 1985, 1990; Kaye, 1992, 2000; Harris, 1994), which is itself based on the earlier work in Autosegmental Phonology (AP, Goldsmith, 1976). In this chapter, I will give a brief outline of ET and detail the principal ways in which this approach differs to that made popular by Chomsky and Halle's (1968) The Sound Pattern of English (SPE). In particular, I will focus on the version of ET advanced in Kaye et al. (1985); Harris (1994); Harris and Lindsey (1995a); Backley (2011). As this will be important for discussion, I will also include a short account of the now outdated Element Calculus (EC), the method employed in Kaye et al. (1985) to convert elemental into feature representations.

2.2 SEGMENTAL COMPOSITION

With regards to the composition of segments, ET stands in stark contrast to that of traditional Feature Theory (FT). One of the major differences is the nature of primitives assumed, which in ET are known as *elements*. Distinctive features in traditional FT are assumed to be throughout bivalent and equipollent (i. e. to express a two-way contrast), while in ET all elements are monovalent and privative — they only express a property through their presence, but not through absence. We can call this the Strict Privativity Principle (SPP):

(1) STRICT PRIVATIVITY PRINCIPLE: All segmental primes are monovalent and privative.

Whereas SPE-style features are linked to articulatory phonetics, elements are linked to the acoustic signal. While features are only meaningful to the interpretation component of the language faculty in fully specified matrices, ET posits that each element itself is an inherently meaningful, i. e. independently interpretable, cognitive prime. This means that, while for instance features such as [+low] or [nas] cannot be pronounced by themselves, elements such as |A|, |I| and |U| are pronounceable in isolation as [a], [i], and [u] respectively. In combination, these elements give rise to more complex segments. For instance |I, U| gives rise to a front rounded vowel [y], |A, I| to a mid front unrounded vowel [e], |A, U| to [o]; the combination |A, I, U| represents a mid front rounded vowel [ø]. We can call this the Independent Interpretability Principle (IIP):

(2) INDEPENDENT INTERPRETABILITY PRINCIPLE: All segmental primes are interpretable both independently in isolation and in combination with other primes.

A segment can also be completely underspecified, i. e. empty. This is usually realised as the central vowel schwa [ə], the 'unmodified default signal'. In analogy with the Empty Category Principle (ECP) in GP and syntax, which determines when a position (a category) may be empty, or forego phonetic interpretation (namely, when it is properly governed), let us call this the Empty Representation Principle (ERP):

(3) EMPTY REPRESENTATION PRINCIPLE: A segmental representation may be empty of primes.

ET also assumes that segments are inherently structured, in contrast to the assumption that feature matrices are unordered sets. A segment in ET is nowadays commonly assumed to have a single, optional head. I. e. one of the elements in the segment can be selected as a *head* (notated by underlining it) and it will then propagate its characteristics asymmetrically over the other elements, the *dependents*. Thus, while for instance |I, A| is interpreted as a mid front unrounded vowel [e̞], |A, $\underline{\text{I}}$| is interpreted as high-mid [e] and |$\underline{\text{A}}$, I| as a low-mid [ɛ]. As such, segmental structure in ET is not only relevant to phonological processes, as is the case with feature geometry, but it is phonologically meaningful — segments with identical content but different heads are conceptually distinct. Conversely, the order of elements does not make a meaningful differentiation, so that two segments $|X, Y|$ and $|Y, X|$ receive the same interpretation. Let us refer to these two principles as the Single Optional Headedness Condition (SOHC) and the Isomericity Principle (IP), respectively:

(4) SINGLE OPTIONAL HEADEDNESS CONDITION: A segment may have exactly one head or no head at all.

(5) ISOMERICITY PRINCIPLE: Two segments are phonologically distinct if and only if they are composed of different elements or have a different head.

Note however that the assumption of the SOHC is not universal among current work in ET, for instance Backley (2011) allows segments to be at least doubly headed. This concept of headedness is also different to that which was used earlier in Element Calculus; this is discussed further in section 2.6.

2.3 ELEMENTS, COMPOUNDS AND INTERPRETATION

A different aspect of the assumptions surrounding elements is a debate about the granularity of the phonological primes. Early proposals of elements, and specifically Kaye et al. (1985), did not contest that the phonological spell-out was principally arranged via features. The association of elements with the acoustic signal was made later (Lindsey and Harris, 1990; Harris and Lindsey, 1993). In fact, Kaye et al. (1985) proposed a specific mechanism called Element Calculus, which translated elements into feature representations[1] which could then be realised by the articulators (I discuss this further in section 2.6). While, it was soon realised that these translations did not produce desirable phonetic representations (Coleman, 1990a,b; Kaye, 1990), the core argument was about the type of primes that are manipulated by the phonology proper, and that these were more granular, privative objects, rather than the phonology directly manipulating highly resolute articulatory features. A major desire was also to be more restrictive in terms of the phonetic forms the system could generate (Kaye, 1990).

As an alternative to Kaye et al.'s (1985) proposal, Lindsey and Harris (1990), Harris and Lindsey (1993, 1995b), Harris (1994, 1996) and Harris and Urua (2001) have argued that there is no such distinct level of phonetic (feature) representation and that elements by themselves are mapped into the acoustic speech signal independently. This is what is now the commonly accepted view in ET. To illustrate this conception, Harris and Lindsey (1993) utilised visualisation of the spectral patterns of the three underlying vocalic primes of ET, |A|, |I|, and |U| — the corner points of the triangular vowel space. Their individual vowel realisations [a, i, u] are associated with particular types

1 In fact, they assumed that each element itself is a fully specified feature matrix.

of spectral patterns, which Harris and Lindsey termed *mAss*, *dIp* and *rUmp* respectively (cf. figure 1, based on Backley 2011). Their core observation is, that the spectral patterns of the hypothesised compound representations such as |A, I| for [e] or |A, U| for [o] are basically combinations of these primitive pattern. As can be seen from figure 2, |A, I| here shows the typical dIp pattern with two prominent formants, however the low energy concentration associated with F1 in the mAss pattern clearly modulates this and leads to lowering of F2 and smoothing of the dip between the two formants. In the |A, U| pattern, the F1 peak of the mAss pattern is lowered due to the first peak in the rUmp pattern, and rUmp's secondary peak also clearly shows modulation on the following slope, the mAss pattern is further visible in the later rise in energy.

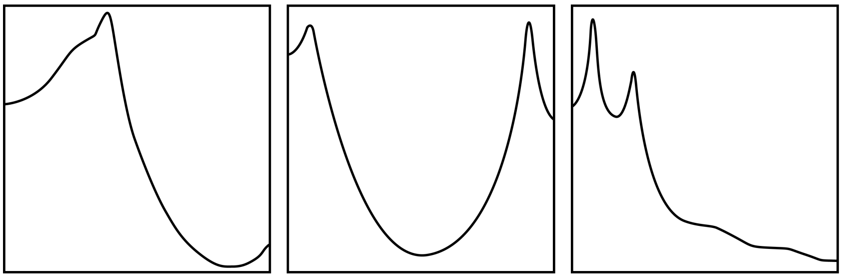

Figure 1: Spectral patterns of |A|, |I|, and |U|: mAss, dIp, and rUmp.

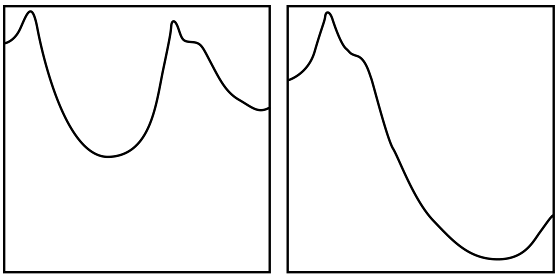

Figure 2: Spectral patterns of the compounds |A, I| and |A, U|.

With such evidence that supports not only an association of the isolated primitives themselves with acoustic properties, but also relate their compounds to the same, salient characteristics, one can then see how basic combinatorics of these elements allow modelling of the basic vowel space, from just the three corner vowels (6) to a larger system with intermediates (7) which follows from simple two-element compounds. When single headedness is introduced, one can additionally distinguish two intermediary forms (8) on either side, depending on which primitive headedness falls. Of course, the SOHC still permits unheaded compounds, further increasing the modellable intermediary vowels between the corners of the vowel space (9). Of course, adding three-element compounds further increases what can be modelled in this space, as does the introduction of elements other than |A|, |I| and |U|.

(6)
i u
a

(7)
i y u
e̞ o̞
a

(8)
i y ɯ u
e o
ə
ɛ ɔ
a

(9)
i y ɨ ɯ u
e o
e̞ ə o̞
ɛ ɔ
a

2.4 THE ELEMENTS

|A|, |I| and |U| are the core resonance elements on which vowel systems in ET are modelled. While such effects as vowel length can be relegated to the surrounding autosegmental model, the representations of consonants and even such effects as nasality and voicing in vowels obviously require the introduction of further elements. Most current work in ET assumes at least six elements, given in table 1 (cf. Backley, 2011). However, there is a good range of variation between precisely which elements are assumed, and what specific properties are ascribed to them. Table 2 gives a number of further elements that have frequently been used in the past and are sometimes still assumed today (see esp. Harris, 1994). One particular trend has clearly been to reduce the number of elements used however, exemplified by work such as Nasukawa (1999, 2000, 2005) who argued for the representation of voice and nasality by one element instead of two, or Pöchtrager (2006) who argued that what is often encoded by the elements |A|, |H|, |ʔ| are not actually elemental but structural properties.

Element Theory also assumes what is often referred to as *consonant–vowel unity* (Backley and Nasukawa, 2010; Backley, 2011). This is the assumption that the exact same primes are active both in the segments that are realised as vowels as in those that are later realised as consonants; whether they are realised

Element	Name	Characteristics
|A|	mAss	Lowness, uvular and pharyngeal place
|I|	dIp	Frontness, palatal and coronal place
|U|	rUmp	Roundness, velar place
|H|	high	High tone, aspiration, frication
|L|	low	Low tone, voicing, nasality
|ʔ|	edge	Stopness

Table 1: The six most commonly assumed elements of ET.

Element	Name	Characteristics
|@|	neutral	Centralisation, identity function
|R|	coronal	Coronal place
|h|	noise	Noise, glottalness, frication
|N|	nasal	Nasality, murmur

Table 2: Some further elements that have been proposed.

as one or the other depends on the structure above the segment. A familiar case in hand are vowel–approximant pairs such as [u, w] and [i, j], which are both very similar phonetically and pattern together phonologically. Take as an examples possible pronunciations of the final segment in the verb /du/ 'to be.1SG' in colloquial Welsh. Look at the three sentences in (10 b–10 a).

(10) (a) [post.mɔn.du.i] 'I'm a postman.'

(b) [du.in.ha.pɪs.i.jaun] 'I'm very happy.'

(c) [dwim.ən.ha.pɪs] 'I'm not happy.'

In each of the three sentences, /du/ is followed by the 1SG pronoun /i/, resulting in the form [du.i] as in (10 a). Attaching a clitic particle as in (10 b) and (10 c) can force resyllabification, so that (10 b) can be pronounced as either [du.in] or [dwin] and (10 c) is always pronounced as [dwim]. From this we can clearly observe that one and the same segment can be realised either as [u] or as [w], depending on how it is syllabified. The phonology requires it to be a vowel to function as a nucleus where no cliticisation takes place. The appropriate generalisation is that the representation of the segment itself does not change at all, as would be the case for instance with a [±syll] feature[2]: in both cases this is simply |U| in the underlying representation.

A good example to illustrate how resonance elements also encode place of articulation comes from Inuktitut[3]. In Inuktitut the 1SG pronoun is attached to a verb as a suffix and can appear either as [tuŋa] or as [juŋa]. The form [juŋa] appears following a vowel, as in [nanuːqujuŋa] 'I'm a polar bear', while [tuŋa] surfaces after consonants, e. g. in [uqalimaːqtuŋa] 'I read'. This pattern can be seen with all the Inuktitut personal pronouns and also with other suffixes. Indeed, this is a case of intervocalic lenition: the stop between two vowels becomes more sonorous, and this is realised through loss (delinking in autosegmental terms) of the edge element, so that |I, ʔ| = [t] becomes |I| = [j]; thus showing a clear relationship between resonance and place of articulation in a stop consonant.

Elements, then, are shared across all segmental representations, but may receive different interpretation depending on the dominating category of the segment itself. In this way, the resonance elements |A|, |I|, |U| receive a place interpretation in consonants, while elements such as |L| which encode nasals or true voicedness (depending on headedness, cf. Nasukawa 2005) in consonants may encode nasalisation in vowels.

2.5 GOVERNMENT PHONOLOGY

The framework from which ET originated is GP, which itself is largely based on AP but aside from assuming elements as the phonological primes instead of features also introduces the notions of licensing and government to account for intersegmental restrictions on phonological structures. Importantly, GP inherited the principal autosegmental skeleton from AP, a typical GP variant of which is illustrated in (11).

2 Depending on assumptions, the features [±tense, ±phar] may also need adjusting in [u] ∼ [w] alternation.

3 This data is in the Uqqurmiut dialect, principally sourced from the Pirurvik Centre's website for learners of Inuktitut, http://www.tusaalanga.ca/.

(11) 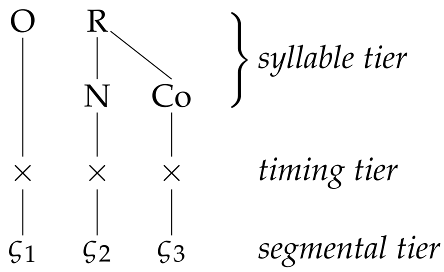

Notably, syllabic structure is relegated to two syllabic tiers above a timing tier (or, the 'x-slots' as they are sometimes called), on which each position assigns a unit of temporal space to the segment attached below it. The syllabic structure is limited to onsets and rhymes, where early on rhymes were allowed to branch into nuclei and codas, but more recent proposals abandon the notion of a coda. There is no level of syllabic structure dominating both the onset and rhyme, and their relations are instead determined by government and by licensing constraints, which for lack of space and immediate relevance I will not discuss further here.

Along with the adoption of the autosegmental skeleton in this framework, subsegmental primes were also assumed to occupy their own tiers, so that the segmental tier may be further divided into an |A|-tier, an |L|-tier, an |ʔ|-tier and so forth, as apparent from the structure of /n/ in (12):

(12) 

Essentially, this then assumes full, language dependent, ordering of the primes below the segment. Additionally, it is usually assumed that two primes can share a single such tier. This is an easy way to account for restrictions on the phonological content of a melodic system. For instance, voiceless or aspirated nasals are rather uncommon cross-linguistically. Taking into account proposals that suggest that both voicing and nasality are represented by the same prime (Nasukawa, 1999, 2000, 2005), call it |L|, where headed |L̲| represents voicing and unheaded |L| nasality, it clearly follows that a system that has voiceless nasals must be an H-system, i. e. a system that represents aspiration via |H| instead of true voicing by |L|, because we could not represent a voicing contrast in nasals with a single such prime. Let us then look at two languages with aspiration systems, English and Welsh. English does not have aspirated nasals, but Welsh does. In Welsh, $|\underline{\text{A}}, \text{ʔ}, \text{L}|$ may thus represent voiced alveolar nasal /n/, while $|\underline{\text{A}}, \text{ʔ}, \text{L}, \text{H}|$ represents an aspirated nasal /n̥/. To account for the impossibility of such representations in English, it is then sufficient to propose that |L| and |H| share a single tier, thus only allowing the three representations for /n,t,d/ in (13), but none that could be interpreted as a voiceless or aspirated nasal.

(13)
$$
\begin{array}{cccc}
\times & \times & \times & {}^*\times \\
| & | & | & | \\
\underline{\text{A}} & \underline{\text{A}} & \underline{\text{A}} & \underline{\text{A}} \\
| & | & | & | \\
\text{ʔ} & \text{ʔ} & \text{ʔ} & \text{ʔ} \\
| & | & & | \\
\text{L} & \text{H} & & \text{L,H} \\
\\
/\text{n}/ & /\text{t}/ & /\text{d}/ & {}^*/\text{n̥}/
\end{array}
$$

2.6 ELEMENT CALCULUS

The compounding operation of the EC in early GP is a binary, *non*-commutative function, i. e. $X \circ Y \neq Y \circ X$. Thus under EC elements had to be combined cyclicly pairwise and the ordering of elements within a segmental representation was important — in contrast to the unordered complement assumed in current versions of ET. Thus, complex representations such as $|X, Y, Z|$ have to be calculated as $X \circ (Y \circ Z)$[4], forcing us to actually specify segmental representations as fully ordered pairs of the type $\langle X, \langle Y, Z \rangle\rangle$, or in short Kuratowski notation as used to represent treelets in much of current syntactic work $\{X, \{X, \{Y, \{Y, Z\}\}\}\}$.

This of course aligns perfectly well with the assumption made in the preceding section that each element occupies its own tier, giving full ordering of elements in the autosegmental skeleton. Compare the two trees in (14) and (15):

4 I. e. assuming leftward compounding, otherwise they could also combine as one of $(X \circ Y) \circ Z$, $(X \circ Z) \circ Y$, $(Y \circ X) \circ Z$, $Z \circ (X \circ Y)$, or any other possible combinations of bracketing and linear order.

(14) 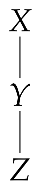

(15) 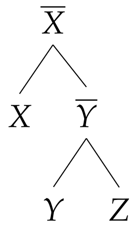

We can see from this how the autosegmental representation in (14) is really just an abbreviated version of (15), which avoids representing the intermediate $\overline{X}$-level. In addition, this gives us a good reason to assume that, where two primes are required to share a single tier, only one of them can appear in any one representation, given that the compounding operation fed by this representation only operates with binary input.

Since the idea of an EC and the level of phonetic representation to which it mapped have however been abandoned, this can no longer be taken as a justification for either the assumption of linearity or that of shared tiers, which critically hinges on that of fixed tiers. And of course no such ordering, apart from a single head, is commonly assumed nowadays, as discussed in section 2.2.

2.7 COMPOSITION AND DECOMPOSITION

In the previous section it has been discussed how EC used to support the idea that elements below the timing tier are ordered on their own tiers, and presumably that they can be affected by phonological processes much like those in AP, linking and delinking. However, such an assumption would be problematic for a number of reasons.

First, if every element has its own tier in a representation and their order is fixed, delinking a segment that is not on the bottommost tier always requires us to also delink all the segments below that prime. Even under the weakest assumption of tier ordering, namely that tier-order can vary from segment to segment, this predicts that the same segment cannot be made to delink two independent primes in different circumstances.

This is however precisely what we see in processes such as initial consonant mutations in the Celtic languages. Here morphosyntactic context can trigger a phonological change on the initial segment of a targeted word, but different paradigms coexist. For instance in Welsh, there are three such mutations: soft mutation, which turns a voiceless plosive into a voiced plosive[5]; nasal mutation, which turns a plosive into a nasal; and aspirate mutation, which turns a plosive into a fricative (Ball and Müller, 1992). At least two of these, soft mutation and aspirate mutation, require 'delinking' of an element from the representation. To turn a voiceless into a voiced plosive, |H| must be delinked[6], and to turn a plosive into a fricative the stop element |ʔ| must be delinked. This is also what both existent ET accounts of Welsh mutations, Buczek (1995) and Cyran (2010), argue. Let us assume a representation of |U, ʔ, H| for /p/ (Cyran, 2010, p. 59). We can then illustrate both processes as in (16) and (17):

(16)

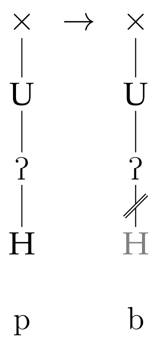

(17)

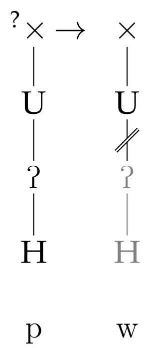

While the change /p/ → [b] in soft mutation is no problem, (17) shows that with the same ordering of tiers we cannot get the change /p/ → [f] which would be expected for aspirate mutation, instead predicting a labial-velar approximant. If we were to reverse the order of the two tiers, we could model aspirate mutation but not soft mutation.

Second, some spreading processes such as vowel harmony, where an element spreads from one vowel to another without affecting intervening consonants, are a problem for the most simple frameworks of AP as they would normally violate the Non-Crossing Constraint, which says that no two association lines in a representation may cross each other. Resolving this requires additional assumptions, for instance that nuclei and onsets project to different parts of the representation so that their lines need not cross to connect them, or proposals that limit constraints on line crossing to certain contexts (cf. Coleman and Local, 1991; Hyman, 2013).

The operation by which segmental representations are altered in ET is thus much more general than the autosegmental notion

5 Soft mutation also targets /m,ɬ,r̥/, but this is not of relevance to the argument made here.

6 we know that Welsh is an |H|-system by virtue of it possessing aspirated nasals, but for further discussion see also Cyran (2010).

of linking and delinking. It is mostly assumed that a phonological process can add or remove single elements to segmental representations without either delinking the remaining representation or establishing any representational association or link between the two segments, though of course other constraints in target selection &c. still apply; especially noteworthy is what is sometimes called the Compositionality Principle (see e.g. Cyran, 2010), which requires that elements to be added to a segmental representation are present in the segment from which the change originates (i. e. elements cannot appear 'out of the blue').

While practically all work in ET now assumes these more loose operations for adding and removing elements from segmental representations, the terms linking, delinking and spreading are still commonly used for this. The departure from actual AP-type linking and delinking is also visible in the notation adopted for graphical illustration in Backley (2011), who shows composition with an arrow, as in (18), and decomposition through a grey box, as in (19). For clarity I will from here on use the terms *composition* and *decomposition* to refer to operations that add and remove elements from a representation.

(18)

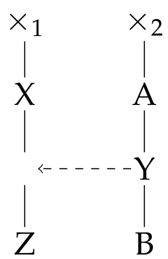

Composition: spreading of Y from $\times_2$ *to* $\times_1$.

(19)

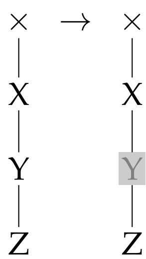

Decomposition of Y, without affecting the Z-tier.

Of course abandoning the notion of ordering, tiers and autosegmental relations in the segmental representation also presents a problem for the account of impossible segmental representation via assignment of the same tier to these elements, and the more general composition and decomposition operations used now are not powerful enough to exclude such combinations from occurring. However, this does not mean that the idea of elemental antagonism is incompatible with this system, and I will introduce an alternative account for elemental antagonism without reference to tiers or ordering in section 3.9.

2.8 ELEMENT GEOMETRY

In FT, the observation that certain features seem to function in groups has led to the proposal of a 'geometric' arrangement of subsegmental primes in place of the simple view that they are simply arranged one by one on tiers below the timing tier (cf. Clements, 1985; Halle, 1995). ET of course must account for the same observations, and the proposition of a subsegmental geometry can be readily adapted to elements. Since the number of primitives is significantly smaller than that in FT however, the geometric arrangement will shrink accordingly. Harris (1994, p. 128), who still uses a larger number of elements than are commonly assumed today, proposes the arrangement in (20) as one that is representative of the consensus at the time[7]:

(20)

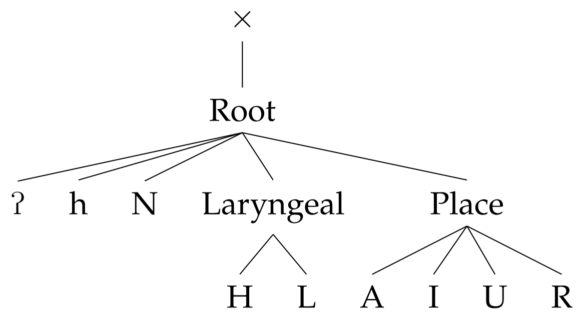

A common example that illustrates the benefit of such 'geometric' arrangements is that of nasal place assimilation. Nasals have been noted as a natural class with a particular tendency to undergo place assimilation, often across all places of articulation. Padgett (2002) gives Kpelle as a prime example of this. The 1SG possessive prefix in Kpelle is a placeless nasal /N/ which always assumes the place of articulation of the first consonant of the root to which it attaches, giving forms such as [mbolu] 'my back', [nduɛ] 'my front', [ŋgɔɔ] 'my foot', and so forth (Padgett, 2002, p. 81).

Instead of requiring separate processes for the assimilation of each of the place elements in these consonants, an element geometry as in (20) can generalise this to leftward spreading of the entire place node, as illustrated for [nduɛ] in (21).

7 See van Oostendorp (2005) for what he sees as the analogous consensus in FT.

(21) 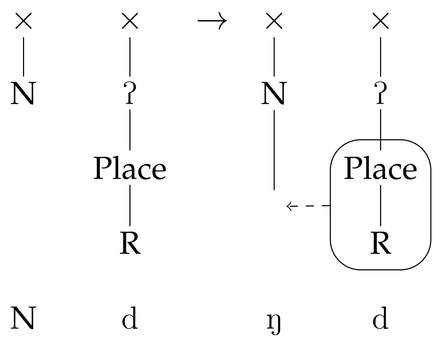

2.9 SUMMARY

In this overview of ET, I have first set out its key properties by stating them as explicit principles. I have then drawn on some aspects of these and historical development within GP to clarify, explicate and narrow down to the particular strand of ET which will form the basis of the remainder of this work. In doing so I have also pointed out some inconsistencies, in particular concerning assumptions surrounding tiers, shared tiers and orderings, which are denied both by existing work in the framework and by the explicit principles of ET, but still seem to sometimes be carried into current proposals through the historic development of the theory.

3 A SET-THEORETIC MODEL OF ET

3.1 INTRODUCTION

In the second chapter, I have given a brief outline of what I see as the current state of ET. In showing how ET developed out of GP and adopts or abandons many previous notions of that framework and earlier ancestors such as AP, it became apparent that this theory of subsegmental structure, while at home in a broadly autosegmental framework, is rather self-contained and does not itself adhere to the same structural principles and relations of other autosegmental systems. In particular, it has been shown how current versions of ET do not appeal to any of the notions of association lines, tier ordering, linking and delinking and how these have been replaced by the SOHC, composition and decomposition.

With his dissertation *Formal Phonology*, Kornai (1995) has demonstrated that AP can be readily captured by a more explicit and rigorous mathematical model. Apart from his interest in bridging the gap between theoretical research and computational application, one of Kornai's main motivations is probably the following statement, which he cites, from Pullum's 1989 comment on the lack of formalisation in large parts of linguistic research:

> Even the best friends of the nonlinear phonology that has driven the relatively formal pre-1977-style segmental phonology into the wilderness [...] will admit that it isn't trying to meet the conditions set out above for formal theories. True, a very significant outpouring of new ideas and new diagrammatic ways of attempting to express them has sprung up over the past decade; but it is quite clear that at the moment no one can say even in rough outline what a phonological representation comprises, using some exactly specified theoretical language. Nor is there much sign of published work that even addresses the issues involved in a serious way. Drifting this way and that in a sea of competing proposals for intuitively evaluated graphic representation does not constitute formal linguistic research, not

> even if interesting hunches about phonology are being tossed around in the process. (Pullum, 1989, p. 138)

Clearly, with the substantial departure from the structures dictated by AP, this criticism is just as applicable to ET now as it was to AP before Kornai (1995).

In this chapter, I will seek to address this issue by giving a formal characterisation of ET as outlined above, ground in basic set theory and logic. For this I will first discuss the extent and delineation of this model. The following sections will give definitions for the different components of this model. Of particular importance here is section 3.5, which gives a mathematically precise definition of what constitutes the core element of ET, a well-formed segmental representation[1]. I then discuss the relations within segments and operations on segments. Finally, I introduce extensions to the presented model which add mechanisms that are able to capture the notions of Element Geometry and elemental antagonism (shared tiers).

3.2 REPRESENTATIONS AND DERIVATIONS

Traditionally, grammars are systems that operate on a set of symbols, often divided into terminal and non-terminal symbols. Via a set of relations and operations (rules, transformations, &c.). The grammar then specifies the way in which these symbols may be combined into representations. For instance given a set of symbols $V = \{a, b\}$, a grammar can be specified that allows either a or b to follow a, but only b to follow b, via a set of rewrite rules such as $O = \{\varnothing \rightarrow a, \varnothing \rightarrow b, a \rightarrow aa, a \rightarrow ab, b \rightarrow bb\}$, which can operate on the rightmost symbol in the representation. From an empty representation $\varnothing$ we can then obtain a or b, from b we can obtain a sequence of any numbers of b's such as $\{b, bb, bbb, bbbb, \ldots\}$. If we begin with a, all the representations featuring only a, i. e. $\{a, aa, aaa, aaaa, \ldots\}$ are derivable. But we can also derive representations of any number of a's followed by any number of b's, e. g. $\{ab, aab, aaab, abb, abbb, aaabb, \ldots\}$, but representations where any b precedes an a, such as $*\{ba, baaa, bba, aba, aaaba, \ldots\}$ are ruled out. The grammar provides a mechanism by which we

1 Note that my use of the terms *segmental representation* and *segment* is largely interchangeable in this chapter.

can *derive* representations from its vocabulary, the set of symbols V.

A system that models the phonology of segments has to work in a somewhat different way however. Clearly, the primary input to the system are not a set of atomic symbols such as features or elements, which are selected and combined to form full segmental representations. The primitives of segmental composition in themselves are largely meaningless and there is no semantics or similar system which would tell us which to select for our representations. Instead, a grammar of segmental phonology has both as its input and its output some kind of already existing, well-formed representation. After all, these are the primary units which have to be learned and stored in the speaker's lexicon to enable the transition from an abstract representation to the linear, pronounceable units of speech that are segments. Thus, the way in which such a grammar operates is by altering representations, not by deriving them[2]. Since representations are principally composed of symbols, it is of course still necessary for the grammar to make reference and enable access to these primitives.

We can specify such a kind of grammar as a set

(22) $$G = \langle V, \Sigma, R, O \rangle,$$

where

A. V is the *vocabulary*, the set of atomic primitives v, e. g. the set of elements in ET;

B. Σ is the set of all well-formed segments ς defined over V;

C. R is the set of relations from Σ to V or $\wp(V)$, such as the *head* relation; and

D. O is the set of operations, mappings from S to S such as *composition* and *decomposition*.

2 Though this may of course only be partially so for phonological models that posit a transition from one vocabulary to another within the same level of representation, e. g. if it is assumed that the lexicon stores [±voice] but the phonological representations translate this into [spread glottis] and [constricted glottis]. For such models many more factors are relevant than for the type of system presumed here, such as what is allowed in the lexicon and in the final representations, where translation happens and through which system, &c. (In fact, most of these models can probably be analysed as systems with different levels of representation with overlapping vocabulary).

The following sections set out how these components can be defined to give a grammar G of this type which models the theory of ET.

3.3 THE VOCABULARY

As discussed in section 2.3, one major factor on which work in ET deviates is the precise set of elements which are assumed. This type of variation is however nearly irrelevant for the basic characteristics of ET, and whether some element v is assumed or not does not affect the overall functioning of the grammar. To reflect this, the elements ought to be referred to in an arbitrary fashion in formally characterising ET, and so the set of all elements V does not need to be defined specifically. It suffices to state,

(23) Let the vocabulary V be the set of all atomic primitives, i. e. elements, in G.

It is in concrete instances of application where we will then want to substantiate, or 'populate', the vocabulary. Following this, for instance for Backley (2011), V is given by

(24) $$V = \{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{H}, \mathrm{L}, ʔ\}.$$

In contrast, under reference to Harris (1994), V is given by

(25) $$V = \{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}, @, \mathrm{N}, \mathrm{H}, \mathrm{L}, ʔ, \mathrm{h}\}.$$

There is a further gain from the arbitrariness of V not discussed so far however: it allows for straight forward implementation of the parametric presence of an element in a system (Cyran, 1996). Cyran, who generally assumes the existence of a noise element |h|, argues that due to parametric variation, some languages such as Munster Irish (Cyran, 1997) and Welsh (Cyran, 2010) do not employ this element in representations, instead using resonant heads to represent this type of noise property (cf. Ritter, 1992). We can now formulate Cyran's h-Parameter as follows:

(26) (a) $\{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{H}, \mathrm{L}, ʔ\} \subseteq V$ for all languages,

(b) $\{\mathrm{h}\} \subset V$ for any language that is an h-system,

(c) there are no other $x \in V$.

Additionally of course working with an arbitrary vocabulary, rather than relying on at least some substantial information about its content, allows for investigations of the effect that such matters as vocabulary size have on other properties of G, such as its Strong Generative Capacity (GenC); and this is precisely what is exploited for the measure of Generative Power (GenP) in chapter 4.

3.4 SEGMENTS

In section 3.2 it was argued that the primary operands in ET are the segments ς, not atomic symbols as in most derivational grammars. Segments are of course still composed of the primitives $v \in V$, and so we need to define how these are to be placed formally in the segment. This now must reflect what was discussed about the composition of segments in ET in section 2: they are generally unordered, with one element optionally singled out as a head (recall the Single Optional Headedness Condition). Headedness in most theories of linguistics generally asserts asymmetry over the part of the representation that is headed by it, and this is also the case in ET. This asymmetry is represented in much of syntactic theory as a structural property, and it is the head which assigns the category of a phrase. This is commonly represented by a tree as in (27):

(27)
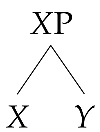

Where X is the *head*, XP the *phrase of category* X, and Y the complement of X. In mathematical terms, such asymmetry is simply expressed as an ordered pair, so that we can say $\mathrm{XP} = \langle X, Y \rangle$. This is in fact what is commonly done in the formal description of phrase structure, in the more low-level Kuratowski notation (Kuratowski, 1921) for ordered pairs

$$\langle a, b \rangle \doteq \{a, \{a, b\}\}. \tag{28}$$

For instance Chomsky (1995) defines the tree adjoining operation *merge* as

$$\text{merge}(\alpha, \beta) = \{\alpha, \{\alpha, \beta\}\} \tag{29}$$

which not only suffices for the first selection of Y by X, but if XP were later selected by Z, we could further substitute Z

for α and $\{X, \{X, Y\}\}$ for β to receive $\{Z, \{Z, \{X, \{X, Y\}\}\}\}$, illustrated in (30)[3]:

(30)

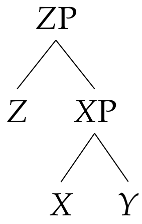

If we want to adopt this notation for the representation of segments in ET we are faced with a challenge: there may be more than two elements in a representation but only one head, and the head may be optional. The latter could be easily addressed by saying that, for instance $\emptyset \notin V$ and then the head may simply be the empty set, but this does not get us much further. Instead, let us begin with the fuller, more standard Kuratowski notation for ordered pairs

(31) $$\langle a, b \rangle \doteq \{\{a\}, \{a, b\}\}.$$

This notation has the advantage that no atomic elements are direct members of the outermost set, instead, we find two sets which in themselves contain the substantial content of the ordered pair. In (31), $\{a\}$ is representative of the head position, and $\{a, b\}$ with the exclusion of what is in the head (i. e. $\{a, b\} \setminus \{a\}$) is representative of the complemental position. To resolve the problem of a single head and several 'complements', we can now extend this notation by allowing any number of further elements in the set that now contains only a and b, i. e.

(32) $$|\underline{a}, b_1, \ldots, b_n| \doteq \{\{a\}, \{a, b_1, \ldots, b_n\}\}.$$

The examples in (33 a) to (33 c) illustrate how some representations would then have to be rewritten as such a set:

(33) (a) $|\underline{\mathrm{A}}, \mathrm{I}| = \{\{\mathrm{A}\}, \{\mathrm{A}, \mathrm{I}\}\}$,

(b) $|\underline{\mathrm{A}}, \mathrm{H}, \mathrm{ʔ}| = \{\{\mathrm{A}\}, \{\mathrm{A}, \mathrm{H}, \mathrm{ʔ}\}\}$,

(c) $|\mathrm{I}, \mathrm{L}, \mathrm{ʔ}, \underline{\mathrm{N}}| = \{\{\mathrm{N}\}, \{\mathrm{N}, \mathrm{I}, \mathrm{L}, \mathrm{ʔ}\}\}$.

Similarly, we can allow the number of b's to be nought to take care of the case where there is only one element which is also the head (e. g. $|\underline{\mathrm{A}}|$ or $|\underline{\mathrm{I}}|$). In this case, the two subsets are identical and collapse, since by definition in set theory $\{X, X\} = \{X\}$,

3 See e. g. Bury (2003) and Bury and Uchida (2008) for some more substantial work on these 'constituent structure sets'.

i. e. a representation where the only element is the head yields a set $\{\{a\},\{a\}\} = \{\{a\}\}$. To make headedness optional we can permit that a simply be omitted, which will leave us with the empty set $\{\} = \emptyset$ and a set $\{b_1, \ldots, b_n\}$, i. e. we receive a set $\{\emptyset, \{b_1, \ldots, b_n\}\}$. These possibilities now cover the full range of permitted segmental representations, as is illustrated by the further examples in (34 a–34 c):

(34) (a) $|\underline{\mathrm{U}}| = \{\{\mathrm{U}\}\}$,

(b) $|\mathrm{I}| = \{\emptyset, \{\mathrm{I}\}\}$,

(c) $|\mathrm{I}, \mathrm{U}, \mathrm{N}| = \{\emptyset, \{\mathrm{I}, \mathrm{U}, \mathrm{N}\}\}$.

To define this format for segmental representations more concisely, let us refer to the two sets in the extended notation in 32 as the head position H and the complement position C[4]. The complement position C can then essentially be any subset of V, i. e.

$$C \subseteq V,$$

which includes the empty set since by definition the empty set is a subset of any set. The head position must be a set with exactly one member or no members at all (i. e. the empty set), and since a in $\{\{a\},\{a, b_1, \ldots, b_n\}\}$ is also in C, must be a subset of C, i. e.

$$H \subseteq C \text{ and } |H| \leq 1.$$

By combining these conditions, we can give the general set theoretic notation for a segmental representation as

(35) $$\varsigma \doteq \{H, C\} \text{ where } C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1.$$

3.5 THE SET OF ALL SEGMENTS

Segmental representations are the central operational component of ET, and in order to make meaningful generalisations it was argued for the vocabulary that populates it that this should for all intents and purposes be an arbitrary set of elements V. Similarly, to be able to make generalisations about segmental representations, we must define from these a set that contains all the well-formed segmental representations in ET, based on that arbitrary vocabulary. Given that we have already generalised from V to a notation to represent all the segments in ET

4 The issue of overlap between H and C is addressed later in section 3.6, where a differentiation between *complement* and *dependent* is introduced.

in (35), this task is rather simple. We can substitute the conditions in (35) into the definition of a generalised set Σ of all the sets $\{H, C\}$ that conform to them. This gives us the set Σ of all well-formed segments:

(36) $$\Sigma \doteq \{\{H, C\} : C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1\}.$$

We can now generalise the notion of well-formedness in light of this. Any set x is a well-formed segmental representation if and only if it is a member of Σ, or

(37) $$(\forall x)(x \text{ is well-formed } \leftrightarrow x \in \Sigma).$$

Thus, all the sets given in (33 a–34 c) are well-formed (assuming the right V), since they are all members of Σ. However, the three sets

(38) (a) $\{\mathrm{A}, \mathrm{I}\}$,

(b) $\{\{\mathrm{A}, \mathrm{U}\}, \{\mathrm{A}, ʔ\}\}$, and

(c) $\{\mathrm{N}, \{\mathrm{A}\}\}$

are all examples of sets which are not well-formed for any V. This is because the definition of Σ does not include any sets of this form. A segment can then potentially fail to be well formed either by its form (i. e. the structure of sets in it) or by its content (i. e. by including items which are not members of V). Indeed, containing a set which is not a subset of V implies that a set is not in Σ:

(39) $$(\forall x)((\forall y \in x)(y \not\subseteq V \rightarrow x \notin \Sigma)).$$

This is especially relevant for proposals such as Cyran (1996), where V is affected by parametric variation.

In addition to (39), there are three other theorems of importance about membership in Σ. First, ET assumes that there is a completely empty representation, || (usually realised as [ə]), i. e. there should exist in Σ a segment ς such that this ς only contains the empty set:

(40) $$(\exists \varsigma \in \Sigma)(\varsigma = \{\varnothing\}).$$

PROOF: Let V be an arbitrary set and Σ be the set $\{\{H, C\} : C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1\}$. $\varnothing \subseteq V$ since by definition $\varnothing$ is a subset of all sets, thus $C \subseteq V$ is true for $C = \varnothing$. Similarly, since

$\emptyset \subseteq \emptyset$ by definition, $H \subseteq C$ is also true for $H = \emptyset$. $|\emptyset| = 0$ and $0 \leq 1$, thus $|H| \leq 1$ is true for $H = \emptyset$. Since $H = \emptyset, C = \emptyset$ satisfies all conditions, $\{\emptyset, \emptyset\} = \{\emptyset\}$ is a member of Σ. □

Second, from the IIP it follows that, for each element, there must be in principle a well-formed representation only comprised of that one element. This is called a *simplex* representation. However, stating and proofing this theorem is a lot easier once we have introduced the head and complement relations in section 3.6 and I will return to it there.

Third, it is argued that the purest independent interpretation of an element is visible in the simplex representation where it is also the head, i. e. representations such as $|\underline{\text{A}}|$ and $|\underline{\text{N}}|$. Let us call these types of representations *singletons*, because their set-representations collapse into a singleton set, e. g. $|\underline{\text{A}}|$ becomes $\{\{\text{A}\}, \{\text{A}\}\} = \{\{\text{A}\}\}$ (A set $\{x\}$ with only one member is called a singleton). While reference to this is somewhat obscured by the common practice to only indicate headedness when it is important for differentiation in work on ET[5], there should be a well-formed singleton representation for every member of V, i. e.

(41) $$(\forall v \in V)((\exists \varsigma \in \Sigma)(\varsigma = \{\{v\}\})).$$

PROOF: Let V be an arbitrary set, $v \in V$, and Σ given by $\{\{H, C\} : C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1\}$. Assume $C = \{v\}$, then $C \subseteq V$ is true since all members of C are necessarily members of V. Similarly, assume $H = \{v\}$, then $H = C$ which implies $H \subseteq C$ is true. Since $|\{v\}| = 1$, $|H| \leq 1$ is also true and thus for any $v \in V$, $\{\{v\}, \{v\}\} = \{\{v\}\}$ is a member of Σ. □

3.6 HEADS, COMPLEMENTS AND DEPENDENTS

Now that the format for segmental representations has been fixed and we have the means to refer to the segments of ET via Σ, we need to establish relations within the segments. In section 2, the notion of *head* and *dependent* elements were introduced, however in 3.4 I have drawn a picture of heads as partial order over a set, in analogy with representations of syntactic constituents. In actuality, headedness in segments can be

5 A practice that clearly adds to the vagueness of definition of the theory and should therefore probably be avoided, or at least based on the category of 'where it is obvious' rather than 'where it makes a difference'. All too often the latter case only becomes apparent in later work.

thought of as both structural and relational. Consider the following statement of clarification from Harris (1994):

> Terms such as *onset* and *nucleus* refer to *categories* of syllabic structure. *Head*, on the other hand, is not a categorial term but rather refers to a phonological function or *relation*, specifically one that is contracted between positions. (Harris, 1994, p. 149)

This clearly expresses the idea that the head occupies a different position its complement or dependents. This becomes even clearer if we illustrate the scheme from (35) as a tree, as in (42). A slightly more useful representation is tree (43) however.

(42)

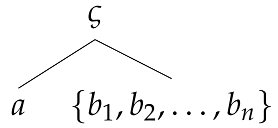

(43)

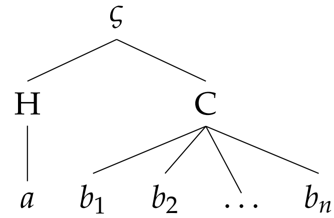

Here the positional properties are clearly apparent, and different relations result structurally. In familiar terms from syntactic discourse, ς dominates both a and all b. The first node dominated by ς would then be referred to as the head (in actuality, ς would be expected to form a category aP, a being the head), and the second node the complement (where there is an extra step indicated, with the dominating node not indicated, in 43). While the segmental representations used in this model clearly indicate a similar structure to syntactic treelets, I will advocate a somewhat different definition of these relations.

Although only the terms head and dependent are commonly used, work in ET actually makes reference to segmental content in three ways: (a) it refers to the head of a segment, (b) it refers to the dependents, i. e. all the elements but the head, and (c) it refers to all the elements of the segment regardless of whether they are head or dependent. I will use the term *complement* to refer to precisely this last frame of reference. All three of these relations are mappings from a segment $\varsigma \in \Sigma$ to a subset of V, and in the notation $\{\{a\}, \{a, b_1, \ldots, b_n\}\}$ one can intuitively pick out the head and complement (they are the first and second set, respectively).

In set theoretic terms, reference is also simple. The head relation indicates what both subsets have in common, i. e. it is exactly the intersection of the subsets of ς. The head relation can thus be defined as

(44)
$$\begin{aligned} \mathrm{H} : \Sigma &\mapsto \wp(V) \\ \mathrm{H}(\varsigma) &= \bigcap \varsigma. \end{aligned}$$

Similarly, the complement is equivalent to the union of the subsets of ς. Since $\{a, a\} = \{a\}$, and the head is always a subset of the complement, this always turns out identical to larger of the two subsets, the complement:

(45)
$$\begin{aligned} \mathrm{C} : \Sigma &\mapsto \wp(V) \\ \mathrm{C}(\varsigma) &= \bigcup \varsigma. \end{aligned}$$

The set of dependents is indicated by the relation that maps from the segment to all the elements in the complement except the head. We can thus use the two relations for head and complement, and indicate the dependency relation as that of complement minus head, i. e.

(46)
$$\begin{aligned} \mathrm{D} : \Sigma &\mapsto \wp(V) \\ \mathrm{D}(\varsigma) &= \mathrm{C}(\varsigma) \setminus \mathrm{H}(\varsigma). \end{aligned}$$

With these three relations we can also return to the theorem of the existence of a simplex element for all $v \in V$ discussed in section 3.5. A simplex element is one that only employs one element in its representation. There are actually two types of segments to which this applies, those that are one-element, self-headed segments (i. e. singletons, cf. theorem 41) and those which have only one element but that element is not also the head. The characteristic properties of this type of segment is thus that the complement contains exactly one element, the complement is $\{v\}$. The singleton is then a special case of the simplex representation (we have another proposition, namely that every singleton representation is also a simplex representation). The proposition then is that such a representation exists for all elements in V, that is

(47)
$$(\forall v \in V)((\exists \varsigma \in \Sigma)(\mathrm{C}(\varsigma) = \{v\})).$$

PROOF: Let v be a member of an arbitrary set V and ς a member of the set Σ, which itself is given by $\{\{H, C\} : C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1\}$. For any $v \in V$, $\{v\} \subseteq V$ necessarily, and thus assuming $C = \{v\}$ satisfies $C \subseteq V$. For both $H = \emptyset$ and $H = \{v\}$ the other two conditions are satisfied (cf. theorems 40 and 41). By the definition of $\mathrm{C}(\varsigma) = \bigcup \varsigma$, i. e. either $\emptyset \cup \{v\}$ or $\{v\} \cup \{v\}$, for both of which $\mathrm{C}(\varsigma) = \{v\}$. $\square$

3.7 OPERATIONS: COMPOSE AND DECOMPOSE

While we now have the means to accurately define segmental representations and covered the relations that hold within segments, of central importance are of course the means with which such segmental representations can be altered. There are two operations in ET, *compose*, which adds elemental content to a representation, and *decompose*, which removes elemental content from representations. Since they alter, rather than 'derive' representations, their essential function is to map one segment ς_1 to another ς_2, i. e. they are $\Sigma \mapsto \Sigma$, so much is clear from the literature at least. In addition, reference has to be made to the material to be in– or excluded however, and the means by which this is done is something that the literature has been entirely quiet on. There are several options: First, we could let the operations make use of an additional well-formed representation, which can be united/intersected with the targeted segment to map to a different segment. These could be defined as

(48)
$$\begin{aligned} \text{comp} : \langle \Sigma, \Sigma \rangle &\mapsto \Sigma \\ \text{comp}(\varsigma_1, \varsigma_2) &= \{H(\varsigma_1) \cup H(\varsigma_2), C(\varsigma_1) \cup C(\varsigma_2)\}, \text{ and} \end{aligned}$$

(49)
$$\begin{aligned} \text{decomp} : \langle \Sigma, \Sigma \rangle &\mapsto \Sigma \\ \text{decomp}(\varsigma_1, \varsigma_2) &= \{H(\varsigma_1) \cap H(\varsigma_2), C(\varsigma_1) \cap C(\varsigma_2)\}. \end{aligned}$$

While this is very simple formally, there are several problems with this approach. For one, it would mean that a suprasegmental grammatical component that effects such operations to take place would have to have the full means to build appropriate well-formed representations in order to feed an appropriate second argument to the operations. This is both somewhat problematic for the few of compartmentalisation and contrary to the point of providing a means to validate representations and define operations at the subsegmental level at all. There is also an empirical problem, arising principally from the requirements of the SOHC: While it would be easy enough to remove several elements from a representation's complement at once, any instance of decomp() could only ever remove one single element from the head position at once. There are however phonological processes that seem remove any head at all, such as head-alignment (a type of harmonic process found for instance in Korean vowel harmony, see Lee 1996), or the advanced stages of lenition which eventually lead to an empty representation. This would require that the grammar either look ahead at the head of the target segment and then remove it, which is easy enough if somewhat awkward intuitively by decomposing the singleton of the target segment's head, i. e. via $\text{decomp}(\varsigma_1, \{\text{H}(\varsigma_1), \text{H}(\varsigma_1)\})$, or that it iterate through all elements in V to form their singletons and remove them, which does not require reference to the content of the target segment by another component.

Second, they may add or remove only a single element at a time, i. e. all operations are $\langle \Sigma, V \rangle \mapsto \Sigma$. This is the case that requires the least assumptions and support, but it will require two variants of either operation, one for targeting heads and one for targeting the complement. This would leave us with four operations,

(50)
$$\begin{aligned} \text{comp} : \langle \Sigma, V \rangle &\mapsto \Sigma \\ \text{comp}(\varsigma, v) &= \{\text{H}(\varsigma), \text{C}(\varsigma) \cup \{v\}\}, \end{aligned}$$

(51)
$$\begin{aligned} \text{decomp} : \langle \Sigma, V \rangle &\mapsto \Sigma \\ \text{decomp}(\varsigma, v) &= \{\text{H}(\varsigma), \text{C}(\varsigma) \cap \{v\}\}, \end{aligned}$$

(52)
$$\begin{aligned} \text{hcomp} : \langle \Sigma, V \rangle &\mapsto \Sigma \\ \text{hcomp}(\varsigma, v) &= \{\text{H}(\varsigma) \cup \{v\}, \text{C}(\varsigma)\}, \text{ and} \end{aligned}$$

(53)
$$\begin{aligned} \text{hdecomp} : \langle \Sigma, V \rangle &\mapsto \Sigma \\ \text{hdecomp}(\varsigma, v) &= \{\text{H}(\varsigma) \cap \{v\}, \text{C}(\varsigma)\}. \end{aligned}$$

Of course these face the same problem for processes removing all heads from a representation, but they do not strike an unjustified imbalance between the ability to target multiple elements at once for complements but not heads. A further advantage is that through their openness of the second argument this is less restrictive on the requirements for phonological processes. This prevents us from precluding without justification the existence of processes that may be prevented by their way of operation.

Imagine for instance a phonological process that is triggered by representations with headed $|\underline{\text{A}}|$ which causes adjacent segments which contain $|\text{A}|$ to promote $|\text{A}|$ to headhood, as proposed to account for Korean vowel harmony by Lee (1996). We can imagine two cases of such a process. First, one in which already existing heads, e. g. in $|\underline{\mathrm{I}}, \mathrm{A}|$ are overwritten to $|\mathrm{I}, \underline{\mathrm{A}}|$, as is the case for Korean where [e] $= |\underline{\mathrm{I}}, \mathrm{A}|$ changes to [ɛ] $= |\mathrm{I}, \underline{\mathrm{A}}|$ under the process of A-head alignment he proposes. Second, one in which this scenario blocks the process from applying (i. e. where [e] in the same situation would have blocked A-head alignment). With the previous proposal the second case is unmodellable, while in the latter it is easy enough to say that the result of $\mathrm{hcomp}(\{\{\mathrm{I}\}, \{\mathrm{I}, \mathrm{A}\}\})$, which is $\{\{\mathrm{I}, \mathrm{A}\}\}$, is not actually a well-formed segment (it is not a member of Σ) and thus the term $\mathrm{hcomp}(\{\{\mathrm{I}\}, \{\mathrm{I}, \mathrm{A}\}\})$ is undefined since hcomp is defined as a mapping to Σ. This prevents the process from applying, unless the head is decomposed first (which is what would be done to model the first case).

Third, we can make the second variant more convenient for reference by retaining the split between operations that target the head and complement separately, but allowing multiple elements to be composed or decomposed at once. This is achieved by making the second argument any subset of V, as in

(54)
$$\begin{aligned} \mathrm{comp} : \langle \Sigma, \wp(V) \rangle &\mapsto \Sigma \\ \mathrm{comp}(\varsigma, \bar{v}) &= \{\mathrm{H}(\varsigma), \mathrm{C}(\varsigma) \cup \bar{v}\}, \end{aligned}$$

(55)
$$\begin{aligned} \mathrm{decomp} : \langle \Sigma, \wp(V) \rangle &\mapsto \Sigma \\ \mathrm{decomp}(\varsigma, \bar{v}) &= \{\mathrm{H}(\varsigma), \mathrm{C}(\varsigma) \cap \bar{v}\}, \end{aligned}$$

(56)
$$\begin{aligned} \mathrm{hcomp} : \langle \Sigma, \wp(V) \rangle &\mapsto \Sigma \\ \mathrm{hcomp}(\varsigma, \bar{v}) &= \{\mathrm{H}(\varsigma) \cup \bar{v}, \mathrm{C}(\varsigma)\}, \text{ and} \end{aligned}$$

(57)
$$\begin{aligned} \mathrm{hdecomp} : \langle \Sigma, \wp(V) \rangle &\mapsto \Sigma \\ \mathrm{hdecomp}(\varsigma, \bar{v}) &= \{\mathrm{H}(\varsigma) \cap \bar{v}, \mathrm{C}(\varsigma)\}. \end{aligned}$$

All of these could of course be mapped on the second proposition made above by applying them one to one to every $v \in \bar{v}$. These are also very convenient in terms of referring to the kind of head-decomposition discussed for option one: $\mathrm{hdecomp}(\varsigma, V)$ will decompose any head. Since options two and three are then the best justified in terms of not going beyond what the literature seems to implicitly suggest in restrictions on composition, these are clearly to be preferred. The last option introduced is moreover notationally more convenient, and I therefore advocate this as a definition for the set of operations in G. However this is clearly an aspect of ET that requires further clarification in the literature and through empirical studies on the limitations of composition and decomposition of heads.

3.8 ELEMENT GEOMETRY

The content discussed so far gives a complete outline of ET's core. However, in section 2 I have also discussed two aspects of ET so far left unaccounted for: element geometry and elemental antagonism. Both of these can easily be added to the above outline of G, but let us first turn to element geometry.

Essentially element geometry is an arrangement of the elements in V, which allows to refer to groups of these elements. Further, none of these groups may overlap, i. e. an element |X| cannot be headed by two different nodes at the same time. These properties are mirrored in the set-theoretic notion of a partition. A partition P of a set X is a set which divides all the members of X into subsets which are collectively exhaustive (i. e. $\bigcup P = X$) and where there exist no two subsets in P which share a member (i. e. $(\nexists(A, B) \in P)(A \cap B \neq \emptyset)$, or negated $(\forall(A, B) \in P)(A \cap B = \emptyset)$). If we are satisfied that such a partition is sufficient to replace element geometry, then we can give a set Γ of all the possible partitions ('geometries') of V as

(58) $$\Gamma = \left\{\gamma : \emptyset \notin \gamma \wedge \bigcup \gamma = V \wedge (\forall (\alpha, \beta) \in \gamma)\,(\alpha \cap \beta = \emptyset)\right\}$$

where γ is a specific partition of V. There is a caveat to this approach however: it is a group based approach, not a hierarchical approach, and thus there are no subdivisions possible, i. e. we cannot represent the idea that two or more nodes themselves are dominated by another node. Given geometries of only two hierarchical levels as in example (20), repeated in (59) below, this is not problematic however.

(59)

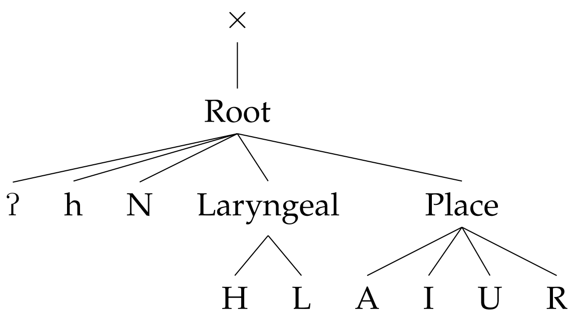

Given the vocabulary $V = \{ʔ, \mathrm{h}, \mathrm{N}, \mathrm{H}, \mathrm{L}, \mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}\}$, this geometry could be replaced by a partition $\gamma = \{\{ʔ, \mathrm{h}, \mathrm{N}\}, \{\mathrm{H}, \mathrm{L}\}, \{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}\}\}$. We can call the set $\{\mathrm{H}, \mathrm{L}\}$ *Lar* and the set $\{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}\}$ *Pla*. Now delinking of the root node in (59) can instead be accounted for by $\mathrm{decomp}(\varsigma, V)$, delinking of the laryngeal node by $\mathrm{decomp}(\varsigma, Lar)$, and so forth. A more interesting case is the linking of one node, say the place node, from a segment ς_1 to a segment ς_2. To illustrate this, let us revisit the example of Kpelle nasal place assimilation from section 3.8. Here a segment $\varsigma_1 = \{\emptyset, \{\mathrm{N}\}\}$ only specified as a nasal, assimilates in place to a segment ς_2, in the case of [nduɛ] a voiced alveolar stop, i. e. $\varsigma_2 = \{\emptyset, \{ʔ, \mathrm{R}\}\}$. To spread any and all place elements from ς_2 to ς_1 without making explicit reference to ς_2's content, when then can use the intersection of $\mathrm{C}(\varsigma_1)$ and *Pla*, which is then composed to ς_1, or in generalised form $\bar{\varsigma}_1 = \mathrm{comp}(\varsigma_1, \mathrm{C}(\varsigma_2) \cap Pla)$. Specific for the [nduɛ] example we can then substitute in the appropriate values and resolve this as follows:

$$\begin{aligned}\bar{\varsigma}_1 &= \mathrm{comp}(\{\emptyset, \{\mathrm{N}\}\}, \mathrm{C}(\{\emptyset, \{ʔ, \mathrm{R}\}\}) \cap \{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}\}) \\ &= \mathrm{comp}(\{\emptyset, \{\mathrm{N}\}\}, \{ʔ, \mathrm{R}\} \cap \{\mathrm{A}, \mathrm{I}, \mathrm{U}, \mathrm{R}\}) \\ &= \mathrm{comp}(\{\emptyset, \{\mathrm{N}\}\}, \{\mathrm{R}\}) \\ &= \{\emptyset, \{\mathrm{N}, \mathrm{R}\}\}. \qquad (= [\mathrm{n}])\end{aligned}$$

Of course decomposition and composition can here be combined to first remove any place elements and then replace them with those of a different segment. For instance let $\varsigma_1 = \{\emptyset, \{\mathrm{A}, \mathrm{R}, \mathrm{L}, ʔ\}\}$ and $\varsigma_2 = \{\emptyset, \{\mathrm{U}, \mathrm{H}\}\}$ and the goal is to replace all place elements in ς_2 with those from ς_1. This then involves decomposition of *Pla* from ς_2, and we have $\bar{\varsigma}_2 = \mathrm{decomp}(\varsigma_2, Pla) = \{\emptyset, \{\mathrm{H}\}\}$. This is followed by composition of all the elements in $\mathrm{C}(\varsigma_1)$ which are also in *Pla*, i. e. the intersection $\mathrm{C}(\varsigma_1) \cap Pla = \{\mathrm{A}, \mathrm{R}\}$. Thus, $\bar{\bar{\varsigma}}_2 = \mathrm{comp}(\bar{\varsigma}_2, \mathrm{C}(\varsigma_1) \cap Pla) = \{\emptyset, \{\mathrm{A}, \mathrm{R}, \mathrm{L}, ʔ\}\}$ maps to the desired segment which has the same place elements as ς_1 but retained all other elements.

While we can of course make up any number of such partitions $\gamma \in \Gamma$, the basic idea of element geometry is that the geometry is universal and not language dependent. Even if that is not assumed, it would not make sense to have more than one geometry for any specific language, and thus an extension to G is best made by adding a specific $\gamma \in \Gamma$, i. e. $G^\Gamma = \langle G, \gamma \rangle$, or more fully,

(60) $$G^\Gamma = \langle V, \Sigma, R, O, \gamma \rangle.$$

Given the common disagreement on the layout of both feature and element geometry, a question of interest is of course still how many such partitions there can be, i. e. what is the cardinality of Γ for any one V (cf. also Kornai, 1993)?

3.9 ELEMENT ANTAGONISM

While not much of the current work in ET makes use of element geometry (for instance Backley 2011 does not even discuss it), which is possibly related to the stark reduction in the number of elements assumed, the notion that two elements can share a single tier, termed *element antagonism*, is still common and of some importance in current work. Antagonistic elements are essentially pairs of elements which may not occur in a representation together; while apparently motivated in opposing acoustic characteristics, as opposed to geometric arrangements, these antagonistic pairs are language dependent (Backley, 2011). Thus, we need to extend G with an arbitrary set Δ of excluded element combinations. Δ can simply be a set of sets which are to be excluded from occurring as complements, since heads are limited to maximally one element, e. g.

(61) $$\Delta = \{\{\mathrm{H}, \mathrm{L}\}, \{\mathrm{I}, \mathrm{U}\}, \dots\}.$$

However, simply extending G with this set is not enough. Because it functions somewhat similarly to the vocabulary V (it is somewhat of an antagonist to V), we must also modify the definition of what a well-formed segment is to the exclusion of all $\delta \in \Delta$. Essentially, we must introduce a condition on Σ that there must not be any $\delta \in \Delta$ which is a subset of any C, i. e.

(62) $$(\forall \varsigma \in \Sigma)((\nexists \delta \in \Delta)(\delta \subseteq \mathrm{C}(\varsigma))).$$

We then have to revise the definition of Σ as follows:

$$\bar{\Sigma} \doteq \{\{H, C\} : C \subseteq V \wedge H \subseteq C \wedge |H| \leq 1 \wedge (\forall \delta \in \Delta)(\delta \nsubseteq C)\}.$$

This would mean that, to include element antagonism, we would define G as

(63) $$G^{\Delta} = \langle V, \bar{\Sigma}, R, O, \Delta \rangle$$

with the modified set $\bar{\Sigma}$. If, for any reason it is not desirable to replace Σ, another possibility is to include $\bar{\Sigma}$ as a further, reduced set of well-formed segmental representations. In this case the definition of $\bar{\Sigma}$ can be abbreviated to

(64) $$\bar{\Sigma} \doteq \{\varsigma : \varsigma \in \Sigma \wedge (\forall \delta \in \Delta)(\delta \nsubseteq \mathrm{C}(\varsigma))\}$$

and G extended to $G^{\Delta} = \langle V, \Sigma, R, O, \Delta, \bar{\Sigma} \rangle$.

While it appears from the literature that antagonistic relations are limited to pairs of elements, no explicit statement is to be found anywhere to that effect and clearly in AP it would be no problem to assume more than two primes share a tier. Nonetheless, an important limitation on possible Δ's is that they may only contain elements that are actually part of the vocabulary. The set $\mathbf{D}$ of possible Δ can then be given as

(65) $$\mathbf{D} = \left\{ \Delta : \left(\forall x \in \bigcup \Delta \right) (x \in V) \right\}.$$

If we want to be more restrictive and actually limit δ's to being pairs of elements, we can give $\mathbf{D}$ as

(66) $$\mathbf{D}^2 = \left\{ \Delta : \left(\forall x \in \bigcup \Delta \right) (x \in V) \wedge (\forall \delta \in \Delta) (|\delta| = 2) \right\}.$$

3.10 SUMMARY

In this chapter I have shown how the ET model of subsegmental phonology can be captured in formally precise terms using basic constructs from set theory and first order predicate calculus. At the centre of the proposed formalism was the set Σ which, for a given vocabulary V, defines the entire set of well-formed representations permitted by ET. Using this it was proved that some of the key theorems about segmental representations in ET, such as the existence of empty representations and simplexes of all elements, hold of the formalisation presented here. In giving a formal account of the relations of head and dependent, it was advanced that although not discussed explicitly in the literature, ET also makes use of a relation which maps to the entirety of the elemental content of the segment, which I have introduced with the term *complement* here and which has shown to be useful in the following definitions of relations and operations. It was pointed out that something that has been overlooked in the informal, graphical discussion of ET before is how the composition and decomposition operations are to be implemented, but I have advocated the view that these come in two flavours each, one to target the head of a segment and one the complement. Finally, it was shown how the basic core model of ET can be extended to include an alternative for element geometry, namely partitions of V, and the type of element antagonism seen in Backley (2011) to replace the notion that certain elements share tiers.

4 THE GENERATIVE CAPACITY OF ELEMENT THEORY

4.1 INTRODUCTION

The ET approach to segmental representation is principally characterised by its privative, independently interpretable primes. This is opposed to the bivalent features of theories in the SPE tradition of FT and Underspecification Theory (UT) (Archangeli, 1988; Steriade, 1995).

A frequent argument made by advocates of ET is that it has a generative capacity more closely reflective of the phonological segments and patterns attested across languages, while feature models faced a problem of overgeneration (e.g. Backley, 2009, 2011; Chen, 2010). A principal factor behind this argument is that ET employs a much smaller set of primitives, with usually around six or seven elements, while most feature-systems posit upward of 20 features[1]. However, this argument has not been formally substantiated in the literature to-date; though discussions of the generative capacities of FT and UT exist (e.g. Reiss, 2012).

In this chapter, I will thus look at the Strong Generative Capacity (SGC) and Generative Power (GenP) of ET and compare this explicitly to the two feature-based models FT and UT. Following on from the work presented in (Reiss, 2012) and my work on ET in the previous chapter, I will adopt a set-theoretic approach to the combinatorics of subsegmental primitives in the three theories and discuss this in the context of cross-linguistic phonological variation, Universal Grammar (UG) and the phonetics–phonology interface.

4.2 GENERATIVE CAPACITY AND GENERATIVE POWER

The Strong Generative Capacity (SGC) of a grammar G is defined as the set of strings that its rules allow us to generate from its vocabulary V (Bach, 2003).

1 SPE makes use of at least 23 features (Chomsky and Halle, 1968), Kornai discusses the possibility of up to 36 features (Kornai, 2008).

Assuming that theories of phonological representation can be seen as (at least partial) grammars of the type outlined in section 3.2, SGC refers to the set Σ of all well-formed representations. However, as a direct consequence of the same combinatorial rules and operations contained in that grammar, one may also usefully look at sets other than Σ. In this context, Reiss (2012) specifically investigates the two questions of possible inventories and definable linear phonological rules of the type $x \rightarrow y/context$.

Clearly, the set of possible inventories I predicted by a theory of segmental representation is of the utmost importance in accounting for cross-linguistic phonemic variation, since this essentially answers the question of how many possible languages such a grammar predicts, all else being equal (Reiss, 2012). However, in looking at theories of segmental representation specifically, such matters as definable rules are of little relevance since they are attributable mainly to system-external factors; i. e. whichever form these modifications take, they are operand on the level of the full segmental skeleton (cf. Scheer, 2013). An important factor that contributes to the formulation of rules over the segmental tier however are the natural classes N which are definable from the inventory of primitives V, as this is what most if not all phonological rules will make reference to and what essentially allows for the types of phonological generalisations such rules make.

I thus argue that we should focus on the three following properties of theories of segmental representation here:

1. The set Σ of all the possible segments
2. The set I of all the possible inventories
3. The set N of all the defined natural classes

However, since we are actually discussing the cardinality (viz. size) of these three sets, it may additionally be useful to make a distinction between the specific generative capacity given the size of the specific vocabulary assumed in that grammar, and a more generic version where the vocabulary sizes are equated. While these terms are often used interchangeably elsewhere, I will adopt the two terms Generative Capacity (GenC) for the specific output of a grammar G given its own vocabulary V_G and the term Generative Power (GenP) for the output of different grammars $G_1, G_2, \ldots, G_n$ with a vocabulary V of equal cardinality.

4.3 VARIATION AND GENERAL CONSIDERATIONS

It has been amply pointed out that many issues other than accuracy of input–output need be considered in the design of grammatical models, especially UG and principle and parameter frameworks (e.g. Chomsky, 2002, 2007; Hale and Reiss, 2008; Reiss, 2012). One particular concern in what has been referred to as the 'bottom-up approach' to UG is the consideration that it is preferable to attribute as little as possible to the biological design of the language faculty. The reasoning is along the lines of Chomsky's statement that 'the less attributed to genetic information for determining the development of an organism, the more feasible the study of its evolution' (Chomsky, 2007) and we may consequently wish to keep the number of innate phonological primitives (i. e. V) minimal.

Yet, with these primitives we have to cover not only the phonological systems attested across the world's languages, but equally consider speaker-to-speaker variation at the I-language level, and what we may refer to as the attestability–computability divide in the sense of Hale and Reiss (2008). That is, we must account not only for attested or even attestable I-languages, but for all those I-language which are computable but may be unattested for any number of reasons, not least because it would be a logical fallacy to expect that every I-language would be borne out solely because it is both computable and attestable. It is this divide that may be the source of the apparent disagreement on whether a large generative capacity is desirable or not. As Reiss (2012) points out, it appears as though discussion of syntactic parameters favours a large margin over attested languages (e.g. Kayne, 2000; Newmeyer, 2004), while phonological overgeneration appears to be disapproved of (e.g. Giegerich, 1992; Kornai, 2008; Backley, 2011). Reiss himself finds it surprising that one should disapprove of a large generative capacity in phonology, given the array of cross-linguistic variation and that present models would allow us to easily model more than 10 billion languages and several tens of thousands of segments with just a few features (Reiss, 2012, p. 173).

If we follow this, we are then left with two counterbalancing guidelines for designing our model of a grammar: minimise the vocabulary and maximise generative capacity. However, recall the arguments for looking at Σ, I and N separately made in section 4.2: if we acknowledge that these sets make different contributions to a grammar, then it may well be that we cannot apply the same reasoning to them. For instance the size of attested phoneme inventories seems to be limited, and we may not want to predict an extensive number of natural classes and rules that are not necessary to account for phonological patterns — this would easily lead us to an 'anything goes' position, contrary to the aim of the UG enterprise: to make generalisations about what is and what is not computable by the language faculty. In the following sections, I will discuss this further and give my reasoning for why we should account for some of these sets with a large margin and for others with only a small margin over what is attested.

4.4 SEGMENTS AND INVENTORIES OF SEGMENTS

Individual languages' phonemic inventories have been argued to be anywhere between 11 (Rotokas, Pirahã) and 141 (!Xũ) (UPSID; Crystal, 2010) and this has been taken as one of the motivations for restricting the generative capacity of phonological theories (cf. Giegerich, 1992; Kornai, 2008). Conversely, it may be pointed out that variation in-between phonemic inventories is abundant, which in turn should lead to the desirability of a large generative capacity (Reiss, 2012). I argue that we have to differentiate between two factors here, maximal size of a single inventory and the possibility of variation between inventories of such limited size. The latter issue principally touches upon the phonetic resolution of the assumed primes.

Considering first the issue of maximal size of single inventories, without paying regard to inter-inventory variation. We can of course note that these appear to be limited. The question is whether it is desirable to attribute this limitation, at least in part, to the design of the innate part of the grammar. From the assumption that the principle set of primitives V is innate and that UG dictates how these primitives may combine, it directly follows that the combinatorial possibilities allowed over V are a limit to the size of any single phonemic inventory.

If we subscribe to the notion that a theory that assumes a smaller set of innate primes is preferable (all else being equal), maximal single inventory size is limited directly by the generative power of the grammar. Since the set of all possible inventories is the powerset of the the SGC (viz. the possible segments), by default we predict a set of possible inventories which is not evenly distributed in space. In fact, since the distribution of cardinality of the subsets of the powerset of any set is normally distributed, there will be many available inventories of the median size and very few at either extremes.

From this at least we can conclude that it is a direct consequence of the assumptions we make about UG that phonemic inventories of very small or very large size should be rare (since there are increasingly few available) and the vast majority of inventories should fall somewhere in-between the two extremes. In effect, if languages' phonemic inventories were assigned by chance (which of course is not the case), we would expect their sizes to be normally distributed.

Figure 3 illustrates the distribution of phonemic inventory sizes in the UCLA Phonological Segment Inventory Database (UPSID). Each inventory is mapped by a cross (×), while the line behind is a normalised average[2]. Note that this almost reflects a normal distribution but is skewed toward the lower end of the available breadth of inventory sizes, and there must thus be some factors external to UG which disfavour larger inventories — which following Hale and Reiss's (Hale and Reiss, 2008) arguments for a substance free approach to phonology, we should not concern ourselves with here. More importantly however, there are still some inventories to be found in these disfavoured intervals, so that these external factors may be absolute only to a certain degree[3]. Since UG always limits inventory size and the space containing all possible inventories is normally distributed, it appears reasonable to aim for a generative power which, given a certain size of V, limits the size of any single inventory with a small margin over what is maximally attested but allows for a large enough margin to account for a basic degree of variation. This should not require modelling the internal distribution within that set through active design considerations of the grammar itself.

From this we come to the second consideration, variation and the phonetic resolution of phonological primitives. If we accept that inventory size is limited, both by UG and external factors, and through this relegated to the space of possible inventories which does not in fact allow for the most variation between inventories (recall the skewedness of real phonemic inventory

2 $g : x \to \frac{f(x-1)+f(x-2)+f(x)+f(x+1)+f(x+2)}{5}$, where $f(x)$ gives the number of languages with a given inventory size and $x \in \mathbb{N}$.

3 For instance a system with just one or two phonemes may simply not provide enough possibilities to encode information given certain other properties of the world (including the language faculty).

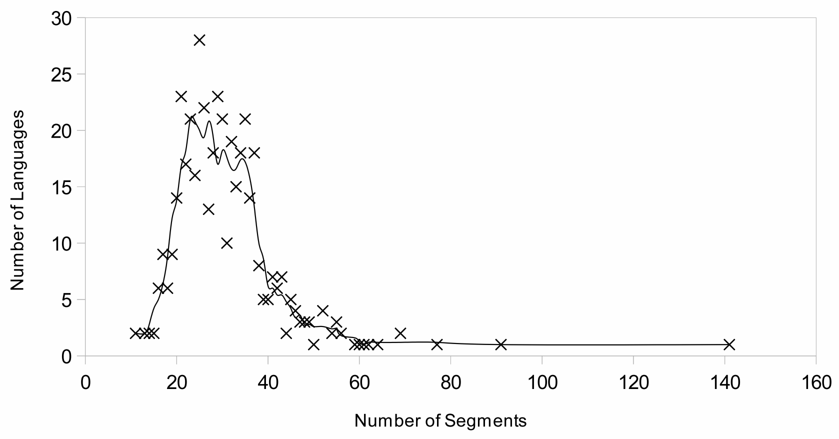

Figure 3: Distribution of phonemic inventory sizes in UPSID.

sizes in figure 3), then how can we sufficiently account for all the detailed surface variation we see cross-linguistically?

This question is principally related to what may be referred to as the phonetic resolution of segmental primitives, i. e. the richness in detail encoded by the primitives of grammar at the interface to phonetic interpretation. The less detail is directly encoded in these primes and the more is ascribed to phonetic interpretation, the more systematic variation is possible on the surface. We may ask how many different phonetic outputs can be assigned to an identical representation in the grammar. For instance 'voiced' plosives in some languages are fully voiced, in others they are partially voiced, unaspirated, tense, &c., yet they encode the same two- or three-way contrast across languages. Then, is it really necessary to encode the precise mechanism with which this happens in the set of primitives of the grammar; do we need features for fully-voiced, partially-voiced, aspirated, tense/slack, and so forth or is it enough to be able to encode a three-way contrast and account for the different output in terms of variation of the phonetic interpretation of the same primitive? One possible assumption is that a binary feature [±voice] is interpreted as voiceless/aspirated in one language, but tense/slack in another, and partially voiced/voiceless in yet another language. Conversely, the phonological patterning and behaviour of voicing contrasts in different languages may provide evidence that there are more than two underlying categorial distinctions (cf. e. g. Halle and Stevens, 1971; Honeybone, 2005). Both types of assumption are in fact quite common in phonological work involving voicing contrasts. Consequentially, we may ask how many possible phonetic realisations can a single phonological representation have? A theory that allows enough variation of this kind does not need a very large set of possible inventories, but many languages can have highly similar or even identical sets of segmental representations, yet show fine-grained variation in the phonetic 'shapes' of their phonemic inventories.

To summarise, following my argumentation, the size of the set of possible segments Σ must minimally be able to account for the largest attested phonemic inventory, with a small surplus. Beyond this, there is no need for a large size of Σ. In order to keep our grammar simpler and less detail rich, it is thus desirable to keep the size of Σ small. Conversely, variation between the phonemic inventories of languages is contributed to both by the size of Σ and the phonetic resolution of the grammar's primitives; yet an argument may be made that it is preferable to have a large set of possible inventories I to allow for enough variation in the underlying segmental representations at a lower size of the specific inventory ι — low phonetic resolution can then further contribute to this basic level of variation[4].

4.5 NATURAL CLASSES

A largely separate issue is that of the natural classes predicted by a grammar. Natural classes are the means by which the grammar can select specific subsets of an inventory ι. As such, they are what determines the domain and co-domain of phonological processes (or rules) which can be applied to a representation. By definition, natural classes are principally intersections between the content of different segments ς, i.e. given an inventory ι with the members $\varsigma_1 = \{a, b, c\}$, $\varsigma_2 = \{b, c, d\}$, $\varsigma_3 = \{a, d\}$, we can say that ς_1 and ς_2 form a set of natural classes N with the members $A = \{x : a \in x\}$, $B = \{x : b \in x\}$, $C = \{x : c \in x\}$ and $D = \{x : d \in x\}$, but also of their combinations, e.g. $BC = \{x : c \in x \wedge b \in x\}$; effectively, a natural class exists for every possible combination of primitives, i.e.

(67) $$(\forall v \in \wp(V))((\exists \nu \in N)(\nu = \{\varsigma : \varsigma \in \Sigma \wedge v \subseteq v\})).$$

4 An additional point that needs to be considered is of course the markedness of the predicted inventories. The assumption made for the sake of comparability here is that all three theories are able to generate a large enough and broadly similar set of unmarked inventories.

Similarly to Σ, I argue for a set N of natural classes that is more restrictive and closely matches the system-internal patterning of primes found in any given language. It is undesirable to predict natural classes where there is no evidence that such a class exists.

4.6 POSSIBLE SEGMENTS

In FT, each segment is composed of sets of features paired with a value $\beta \in B$. Following Reiss (2012), if we say that $V_{\mathrm{FT}} = \{F_1, F_2, \ldots, F_n\}$ for n features, then a segment ς_{FT} can be given as a set of simple feature value pairs. For $B = \{+, -\}$, we get

(68) $$\varsigma_{\mathrm{FT}} = \{\langle F_1, \beta_1\rangle, \langle F_2, \beta_2\rangle, \ldots, \langle F_n, \beta_n\rangle\}.$$

UT essentially extends this mapping by a further possible value to $B = \{+, -, \diamond\}$, where '$\diamond$' means 'unspecified' and G_{UT} does not allow mapping any $\diamond \mapsto \{+, -\}$.

The combinatoric possibilities for FT and UT can then be given as

(69) $$|\Sigma_{\mathrm{FT}}| = 2^{|V_{\mathrm{FT}}|}, \text{ and}$$

(70) $$|\Sigma_{\mathrm{UT}}| = 3^{|V_{\mathrm{UT}}|}.$$

$\|V\|$	$\|\Sigma_{FT}\|$	$\|\Sigma_{UT}\|$	$\|\Sigma_{ET}\|$
0	1	1	1
1	2	3	3
2	4	9	8
⋮			
6	64	729	256
7	128	2,187	576
8	256	6,561	1,280
9	512	19,683	2,816
⋮			
20	1,048,576	3,486,784,401	11,534,336
21	2,097,152	10,460,353,203	24,117,248
22	4,194,304	31,381,059,609	50,331,648
23	8,388,608	94,143,178,827	104,857,600
24	16,777,216	282,429,536,481	218,103,808

Table 3: Some values of $|\Sigma|$ for different sizes of V.

The set Σ_{ET} of all the possible segments in ET has already been defined in (36) in section 3.5; we only need to find the cardinality of the set, $|\Sigma_{ET}|$. This is given by[5]

(71) $$|\Sigma_{ET}| = \left(1 + \frac{|V_{ET}|}{2}\right) \times 2^{|V_{ET}|}.$$

Where $2^{|V|}$ are the combinatoric possibilities covering all complements in Σ (i. e. $\wp(V)$); single optional headedness then adds $\frac{|V|}{2}2^{|V|}$ to the cardinality of Σ.

Sizes of Σ for all three are given for some vocabulary sizes in table 3. We see that SGC for most instantiations of ET is somewhere between 256 and 1,280 segments. In comparison, FT and UT with at least 20 features generate upward of 1 million to 3 million segments. Notably, an instance of ET with this many elements would generate far more segments than either of the feature theories.

ET has a greater GenP than UT and UT a greater GenP than FT. This can be seen especially in the plots in figures 4 and 5. Plot

5 This coincides with sequence A001792 in the Online Encyclopædia of Integer Sequences, http://oeis.org/A001792.

4 also shows clearly how the exponentiality of their functions leads to an explosion of the size of Σ well below 20 vocabulary items for all three theories.

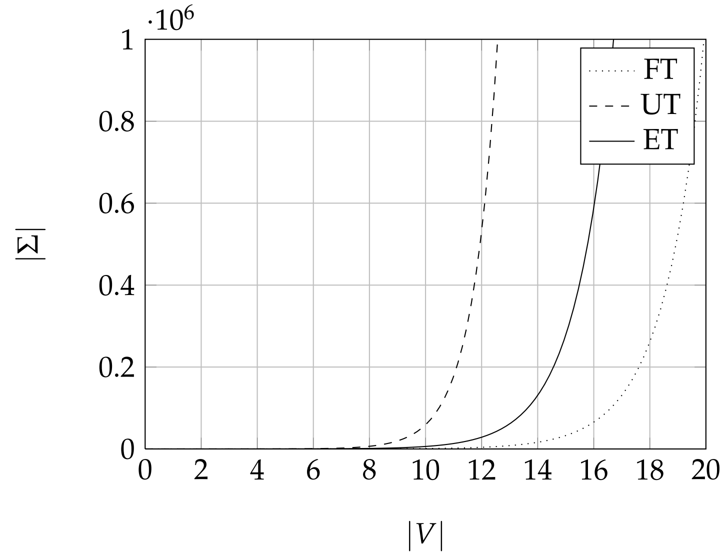

Figure 4: Plot of $|\Sigma|$ for vocabulary sizes from 0 to 20, y-axis in millions.

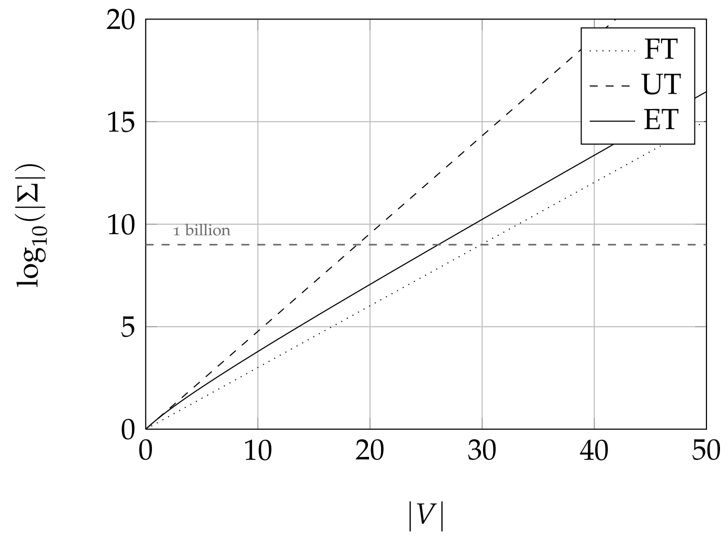

Figure 5: Plot of $\log_{10}(|\Sigma|)$ for vocabulary sizes from 0 to 50.

4.7 POSSIBLE INVENTORIES

The set of all the inventories I derivable from a given set Σ consists of all the possible combinations of the segments in Σ, i.e. $I = \wp(\Sigma)$. Since $|\wp(X)| = 2^{|X|}$, we can give $|I|$ for each of the three theories as

(72) $$|I_G| = 2^{|\Sigma_G|}.$$

As can be seen from the plot in figure 6, this function grows extremely fast for any of the three theories. The data in table 4 show how incredibly large $|I|$ really becomes for only a few primitives.

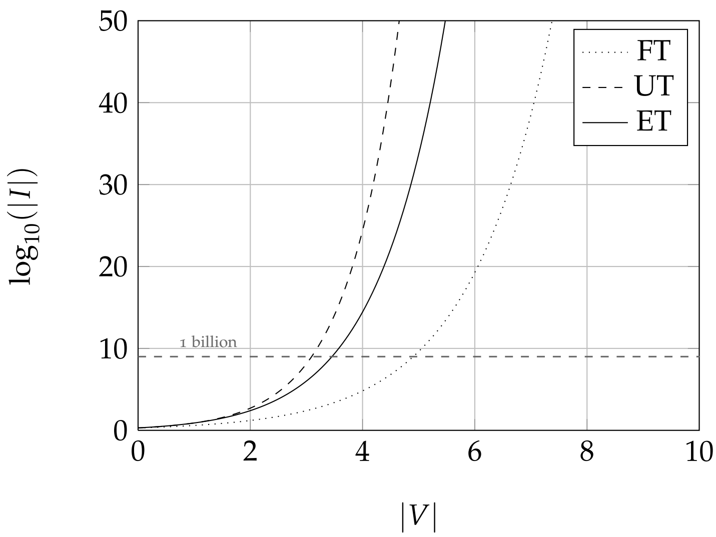

Figure 6: Plot of $\log_{10}(|I|)$ for vocabulary sizes from 0 to 10.

$\|V\|$	$\|I_{\mathrm{FT}}\|$	$\|I_{\mathrm{UT}}\|$	$\|I_{\mathrm{ET}}\|$
0	2	2	2
1	4	8	8
2	16	512	256
3	256	1.342×10^{8}	1.049×10^{6}
4	65,536	2.418×10^{24}	2.815×10^{14}
5	4.295×10^{9}	1.414×10^{73}	5.192×10^{33}
6	1.845×10^{19}	2.824×10^{219}	1.158×10^{77}
7	3.403×10^{38}	$\vdots$	2.47×10^{173}
8	1.158×10^{77}		$\vdots$
9	1.34×10^{154}		

Table 4: Some values of $|I|$ for different sizes of V.

4.8 POSSIBLE NATURAL CLASSES

The set of natural classes defined by FT/UT and ET is an area where the privative–binary distinction matters. With privative primes, only the existent primes and their combinations form natural classes. In binary systems, natural classes exist for both the $+$ and $-$-valued primes as well as partial matrices.

All possible combinations in ET are covered by powerset, i. e.

$$N_{\mathrm{ET}} = \wp(V_{\mathrm{ET}}), \text{ and} \tag{73}$$

$$|N_{\mathrm{ET}}| = 2^{|V_{\mathrm{ET}}|}. \tag{74}$$

This is the same as $|\Sigma_{\mathrm{FT}}|$.

For FT, in addition to full matrices, partial matrices also define natural classes. This adds over the $2^{|V_{\mathrm{FT}}|}$ system by allowing exactly what UT does: leaving v unspecified as $\diamond$, thus

$$N_{\mathrm{FT}} \equiv \Sigma_{\mathrm{UT}}. \tag{75}$$

N_{UT} cannot add any further onto the capacity defined by Σ_{UT}, since this already includes all possible partial matrices.

Consequently, ET with 7 primes predicts 128 natural classes; but FT/UT with 20 features predict over 3 billion natural classes (cf. table 3).

4.9 DISCUSSION

The results in themselves illustrate that it is important to make a distinction between GenP and SGC. While it is true to say that the SGC of FT and especially UT is larger than that of ET, ET is more powerful than both UT and FT if a similar number of primitives are assumed. This also shows that the decisive factor is not the form of representations or a binary–privative opposition, but vocabulary size.

In terms of fit to the criteria laid out for Σ, I, and N, $|\Sigma_{\mathrm{ET}}|$ provides the best fit since it has a narrow margin over the maximal attested size of any single inventory. The rapid growth of $|I|$ showed that ET, FT and UT are all able to account sufficiently for variation, regardless of vocabulary size. Indeed, they all provide a margin of several billion possible unique inventories for every single person alive today.

A notable difference exists in the number of natural classes: ET predicts hundreds, whereas FT and UT predict billions. If

ET accounts sufficiently for attested phonological patterns, this then poses the question why so many natural classes should be necessary. Further of course, at this magnitude, one may ask whether generalisations over such a large set are meaningful at all.

To summarise, the overall most important factor appears to lie in the number and kind of primes assumed rather than a feature–element opposition. A small enough UT system with a low phonetic resolution could easily satisfy the same criteria, but as is, the SGC of ET appears to indeed provide the best fit to general observations about segmental systems.

5 CONCLUSION

The main aim of this dissertation has been to address the lack of a formally precise definition of the model of segmental representation known as Element Theory. For this it was first necessary to narrow down current research in the framework to what can be understood as the broader consensus on ET in the recent past and to point out where the main disagreements and unclarities lie: much debate and change in the theory is centered around the elements themselves, while unclarity exists regarding the status of factors such as tier-ordering and the definition of the composition and decomposition operations.

The aim was then achieved by detailing a set of formal definitions, ground in set theory, which together are able to model ET as described before. A central proposition of this account has been that segmental representations can be modeled by a simple extension of the notation for ordered pairs, $\{\{h\}, \{h, d_1, \ldots, d_2\}\}$, where h represents the head and d the remaining dependent elements. It has been shown how this can be generalised into a precise definition of well-formedness for segmental representations and proofs have been presented that show the central properties of ET, discussed under the headings of the SOHC, ERP and IIP, do hold for these constructs. In formally defining the head and dependent relations, it has been shown how a complement relation must also exist, which has been shown to be useful in the remaining work. Discussion of the composition and decomposition operations has revealed several possible ways of defining them. While no definite conclusion was made as to the correct definition, it has been discussed how these differ in the predictions they would make and the more conservative definition was adopted based on this. Finally, it was shown how the notions of element geometry and shared tiers can be replaced by introducing partitions of the vocabulary, without reference to tier ordering.

Based on this formal account of ET, the secondary aim was to evaluate a key claim made by ET advocates: that the generative capacity of ET provided a better fit to empirical data than feature theories. It has been argued that this claim must be dismantled into three separate factors: that of the number of segments predicted by a theory, the number of predicted inventories and the number of natural classes made available. Based on the data about the variation of segmental inventories and considerations about the design aims for the UG aspects of theories of segmental representation, the argument has been advanced that overgeneration of individual segmental inventories and natural classes should be limited, but that it is beneficial to predict many possible inventories. By evaluating these three aspects as sets, it has then been demonstrated that the claim that ET overgenerates the least is true, but that this is solely attributable to ET's use of the fewest primes in practice, while, possibly somewhat surprisingly, ET is otherwise significantly more powerful than feature theories given an identical inventory of primitives.

A major limitation of this work is that many propositions contained in the formal definition of ET, as well as the cardinality of Σ, have not been formally proved and this may usefully be done in future research to further substantiate the claim that ET, or at least the account presented here, is internally consistent. While the discussion of generative capacity certainly demonstrated the usefulness of a formal account, another major benefit, the formal verifiability of phonological analyses as mathematical propositions, has not been demonstrated beyond few examples of composition and decomposition and further work could usefully demonstrate how this can lead to further insights and greater accuracy in phonological research. Beyond this, a major area of future research could aim to explore how the sets of segmental representations can be computationally mapped into an acoustic signal — something that would make a substantial contribution to demonstrating that an intermediate level of phonetic representation is indeed not required.

BIBLIOGRAPHY

Archangeli, D. (1988). Aspects of underspecification theory. *Phonology*, 5:183–207. (Cited on page 37.)

Bach, E. & Miller, P. (2003). Generative capacity. In Frawley, W., editor, *The International Encyclopedia of Linguistics*, volume 3, pages 20–21. Oxford University Press, Oxford, 2 edition. (Cited on page 37.)

Backley, P. (2009). Element theory and the structure of English vowels. Ms., Tohoku Gakuin University. (Cited on page 37.)

Backley, P. (2011). *An Introduction to Element Theory*. Edinburgh University Press, Edinburgh. (Cited on pages 3, 5, 6, 7, 14, 20, 33, 35, 37, and 39.)

Backley, P. and Nasukawa, K. (2010). Consonant-vowel unity in Element Theory. *Phonological Studies*, 12:21–28. (Cited on page 7.)

Ball, M. and Müller, N. (1992). *Mutation in Welsh*. Routledge, London. (Cited on page 13.)

Buczek, A. (1995). Lenition and fortition: Consonant mutation in Welsh. In Gussmann, E., editor, *Licensing in Syntax and Phonology*, volume 1 of *PASE Studies and Monographs*, pages 187–216. Lublin, Folium. (Cited on page 13.)

Bury, D. (2003). *Phrase structure and derived heads*. PhD thesis, University College London. (Cited on page 22.)

Bury, D. and Uchida, H. (2008). Constituent structure sets. *UCL Working Papers in Linguistics*, 20:87–111. (Cited on page 22.)

Chen, Y.-L. (2010). How powerful are elements? An evaluation of the adequacy of element theory in phonological representations. MSc dissertation, University of Edinburgh. (Cited on page 37.)

Chomsky, N. (1995). *The Minimalist Program*. MIT Press, Cambridge, MA. (Cited on page 21.)

Chomsky, N. (2002). *On Nature and Language*. Cambridge University Press. (Cited on page 39.)

Chomsky, N. (2007). Approaching UG from below. In Sauerland, U. and Gärtner, H.-M., editors, *Interfaces + recursion = language? Chomsky's Minimalism and the View from Syntax-semantics*, volume 89, pages 1–24. Mouton de Gruyter, Berlin. (Cited on page 39.)

Chomsky, N. and Halle, M. (1968). *The Sound Pattern of English*. Harper & Row, New York, NY. (Cited on pages 1, 3, and 37.)

Clements, G. N. (1985). The geometry of phonological features. *Phonology Yearbook*, 2:225–252. (Cited on page 15.)

Coleman, J. S. (1990a). Charm theory defines strange vowel sets. *Journal of Linguistics*, 26(1):165–174. (Cited on page 5.)

Coleman, J. S. (1990b). Vowel sets: a reply to Kaye. *Journal of Linguistics*, 26(1):183–187. (Cited on page 5.)

Coleman, J. S. and Local, J. K. (1991). The "no crossing constraint" in autosegmental phonology. *Linguistics and Philosophy*, 14:295–338. (Cited on page 13.)

Crystal, D. (2010). *The Cambridge Encyclopedia of Language*. Cambridge University Press, Cambridge, 3 edition. (Cited on page 40.)

Cyran, E. (1996). The parametric occurence of elements in phonological systems. In Kardela, H. and Szymanek, B., editors, *A Festschrift for Edmund Gussmann from His Friends and Colleagues*, page 75098. Redakcja Wydawnictw KUL, Lublin. (Cited on pages 20 and 24.)

Cyran, E. (1997). *Resonance Elements in Phonology: A Study in Munster Irish*. Folium, Lublin. (Cited on page 20.)

Cyran, E. (2010). *Complexity Scales and Licensing in Phonology*. De Gruyter Mouton, Berlin. (Cited on pages 13, 14, and 20.)

Giegerich, H. (1992). *English Phonology: An Introduction*. Cambridge University Press, Cambridge. (Cited on pages 39 and 40.)

Goldsmith, J. (1976). *Autosegmental Phonology*. PhD thesis, MIT. (Cited on pages 1 and 3.)

Hale, M. and Reiss, C. (2008). *The Phonological Enterprise*. Oxford University Press, Oxford. (Cited on pages 39 and 41.)

Halle, M. (1995). Feature geometry and feature spreading. *Linguistic Inquiry*, 26(1):1–46. (Cited on page 15.)

Halle, M. and Stevens, K. (1971). A note on laryngeal features. Quarterly Progress Report 101. Research Laboratory of Electronics, MIT. (Cited on page 42.)

Halle, M. and Vergnaud, J.-R. (1982). On the framework of autosegmental phonology. In van der Hulst, H. and Smith, N., editors, *The structure of phonological representations*, volume 1, pages 65–82. Foris, Dordrecht. (Cited on page 1.)

Harris, J. (1994). *English Sound Structure*. Blackwell, Oxford. (Cited on pages 3, 5, 7, 15, 20, and 26.)

Harris, J. (1996). Phonological representations are redundancy-free and fully interpretable. In Durand, J. and Laks, B., editors, *Current trends in phonology: models and methods,*, volume 1, pages 305–332. European Studies Research Institute, Manchester. (Cited on page 5.)

Harris, J. (2007). Representation. In de Lacy, P., editor, *The Cambridge Handbook of Phonology*, pages 119–137. Cambridge University Press, Cambridge. (Cited on page 1.)

Harris, J. and Lindsey, G. (1993). There is no level of phonetic representation. *UCL Working Papers in Linguistics*, 5:354–373. (Cited on pages 1, 3, 5, and 6.)

Harris, J. and Lindsey, G. (1995a). The elements of phonological representation. In Durand, J. and Katamba, F., editors, *Frontiers of phonology: atoms, structures, derivations*, pages 34–79. Longman, Harlow, Essex. (Cited on page 3.)

Harris, J. and Lindsey, G. (1995b). Segmental decomposition and the signal. In Dressler, W. U., Prinzhorn, M., and Rennison, J. R., editors, *Phonologica1992: proceedings of the 7th International Phonology Meeting*, pages 97–106. Rosenberg & Sellier, Turin. (Cited on page 5.)

Harris, J. and Urua, E.-A. (2001). Lenition degrades information: consonant allophony in Ibibio. *Speech, Hearing and Language: Work in Progress*, 13:72–105. (Cited on page 5.)

Honeybone, P. (2005). Diachronic evidence in segmental phonology: the case of obstruent laryngeal specifications. In van Oostendorp, M. and van de Weijer, J., editors, *The Internal*

Organization of Phonological Segments, pages 319–354. Mouton de Gruyter, Berlin. (Cited on page 42.)

Houston, K. (2009). *How to Think Like a Mathematician*. Cambridge, Cambridge. (Cited on page 2.)

Hyman, L. (2013). How autosegmental is phonology? *UC Berkeley Phonology Lab Annual Report*, pages 32–59. (Cited on page 13.)

Jakobson, R., Fant, C. G. M., and Halle, M. (1952). *Preliminaries to Speech Analysis: The Distinctive Features and their Correlates*. MIT Press, Cambridge, MA. (Cited on page 1.)

Jakobson, R. and Halle, M. (1956). *Fundamentals of Language*. Mouton & Co., The Hague. (Cited on page 1.)

Kaye, J. (1990). The strange vowel sets of charm theory: the question from top to bottom. *Journal of Linguistics*, 26(1):175–181. (Cited on page 5.)

Kaye, J. (1992). Do you believe in magic? the story of s+c sequences. *SOAS Working Papers in Linguistics and Phonetics*, 2:293–313. (Cited on page 3.)

Kaye, J. (2000). A users' guide to government phonology. Ms., University of Ulster. (Cited on page 3.)

Kaye, J. and Harris, J. (1990). Segmental complexity and phonological government. *Phonology*, 7:255–300. (Cited on page 3.)

Kaye, J., Lowenstamm, J., and Vergnaud, J.-R. (1985). The internal structure of phonological elements: a theory of charm and government. *Phonology*, 2:305–328. (Cited on pages xiv, 1, 3, and 5.)

Kaye, J., Lowenstamm, J., and Vergnaud, J.-R. (1990). Constituent structure and government in phonology. *Phonology Yearbook*, 7:193–231. (Cited on page 3.)

Kayne, R. (2000). *Parameters and Universals*. Oxford University Press, Oxford. (Cited on page 39.)

Kornai, A. (1993). The generative power of feature geometry. *Annals of Mathematics and Artificial Intelligence*, 8:37–46. (Cited on page 33.)

Kornai, A. (1995). *Formal Phonology*. Outstanding Dissertations in Linguistics. Garland Publishing, New York, NY. (Cited on pages 2, 17, and 18.)

Kornai, A. (2008). *Mathematical Linguistics*. Advanced Information and Knowledge Processing. Springer, Berlin. (Cited on pages 37, 39, and 40.)

Kuratowski, C. (1921). Sur la notion de l'ordre dans la théorie des ensembles. *Fundamenta Mathematicæ*, 2(1):161–171. (Cited on page 21.)

Lee, D.-Y. (1996). Vowel harmony in Korean. In *Proceedings of 3rd Government Phonology Workshop*. University of Vienna. (Cited on page 29.)

Lindsey, G. and Harris, J. (1990). Phonetic interpretation in generative grammar. *UCL Working Papers in Linguistics*, 2:355–369. (Cited on page 5.)

Nasukawa, K. (1999). Prenasalisation and melodic complexity. *UCL Working Papers in Linguistics*, 11:53–70. (Cited on pages 7 and 10.)

Nasukawa, K. (2000). Nasal harmony as prosody-driven agreement. *UCL Working Papers in Linguistics*, 12:257–285. (Cited on pages 7 and 10.)

Nasukawa, K. (2005). *A Unified Approach to Nasality and Voicing*. Mouton de Gruyter, Berlin. (Cited on pages 7, 9, and 10.)

Newmeyer, F. (2004). Typological evidence and universal grammar. *Studies in Language*, 28:527–548. (Cited on page 39.)

Padgett, J. (2002). The phonology of kpelle welmers. *Language*, 78(1):81–110. (Cited on page 15.)

Partee, B. H., ter Meulen, A., and Wall, R. E. (1990). *Mathematical Methods in Linguistics*. Kluwer Academic Publishers, Dordrecht. (Cited on page 2.)

Pöchtrager, M. (2006). *The Structure of Length*. PhD thesis, University of Vienna. (Cited on page 7.)

Pullum, G. K. (1989). Formal linguistics meets the boojum. *Natural Language and Linguistic Theory*, 7:137–143. (Cited on pages 17 and 18.)

Reiss, C. (2012). Towards a bottom-up approach to phonological typology. In Di Sciullo, A. M., editor, *Towards a Biolinguistic Understanding of Grammar: Essays on Interfaces*. John Benjamins, Amsterdam. (Cited on pages vii, 2, 37, 38, 39, 40, and 44.)

Ritter, N. (1992). Headedness as a means of encoding stricture. In Booij, G. and van der Weijer, J., editors, *Phonology in Progress – Progress in Phonology*, pages 333–365. Holland Academic Graphics, The Hague. (Cited on page 20.)

Scheer, T. (2013). Melody-free syntax. Paper presented at the Workshop on complexity at UQAM, Montréal. (Cited on page 38.)

Steriade, D. (1995). Underspecification and markedness. In Goldsmith, J. A., editor, *The Handbook of Phonological Theory*, pages 114–174. Blackwell, Oxford. (Cited on page 37.)

van Oostendorp, M. (2005). Feature geometry. Ms., University of Leiden. (Cited on page 15.)