

Morphology and Compositional Interpretation

MSc Thesis (*Afstudeerscriptie*)

written by

Nathan Malik

(born 30th December 1997 in Stamford, Connecticut, USA)

under the supervision of **Dr. Maria Aloni** and **Dr. Thomas Roberts**, and
submitted to the Examinations Board in partial fulfillment of the requirements
for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense: **Members of the Thesis Committee:**
3rd February 2023

Dr. Balder ten Cate (Chair)
Dr. Maria Aloni (Co-supervisor)
Dr. Thomas Roberts (Co-supervisor)
Dr. Paul Dekker
Dr. Roland Pfau



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

Every definition of the Principle of Compositionality presupposes some set of atoms that provide the primary inputs to a compositional interpretation. When it comes to the study of the compositionality of natural language, such atoms are often equated with words. The diversity of word-structure, or morphology, attested cross-linguistically is vast and highly complex. This naturally leads to the question of how best to understand theoretically the relevance of morphological structure to a compositional semantics; namely, which parts of natural language are the most basically meaningful elements, the compositional atoms that may be stored in the lexicon and serve as the initial inputs to a process of interpretation? Any morphological theory which maintains that proper parts of morphologically complex words can be isolated and are individually meaningful is subject to a number of notable challenges, once it is related to broader grammatical issues involved in a compositional natural language semantics. The alternative position, which takes morphologically complex words in their entirety to be syntactically atomic and semantically basic, invites several questions of its own. Upon the adoption of a holistic conception of the lexicon, that systematically organises the forms and meanings of whole, internally complex words as nodes in a structured network of lexical and grammatical relations, those questions may begin to be addressed.

Contents

1	Introduction	3
1.1	A Conception of the Grammar	5
1.2	The Relevance of Morphology	8
1.2.1	Inflectional Morphology	10
1.2.2	Word-Formation and Derivational Morphology	12
2	Morphological Theory	13
2.1	A Taxonomy of Morphological Theory	14
2.2	Evaluation	19
2.2.1	Morphemic-Incremental	21
2.2.2	Morphemic-Realisational	26
2.2.3	Inferential-Incremental	29
2.2.4	Inferential-Realisational	33
3	Word-Based Morphology	39
3.1	In Defense of “the Word”	40
3.1.1	Special Sound	41
3.1.2	Special Meaning	45
3.1.3	Special Structure	47
3.2	Words and Paradigms	51
3.2.1	Learnability	54
3.2.2	Processing	58
3.3	Semantic WP - Morphological Holism	60
4	Morphological Case Studies	65
4.1	Number	66
4.2	Agreement	69
4.2.1	Anaphoric Pronominal Reference	69
4.2.2	Adjective and Determiner Gender Agreement	76
4.2.3	Null Subjects	80
4.3	Recalcitrant Case	84
4.4	Typology and Diachrony	89
5	Conclusion	94

Chapter 1

Introduction

The Principle of Compositionality provides a fundamental basis for the construction of formal theories of interpretation. There are several different specific characterisations resulting from different interpretations of its basic formulation, and each may be subject to a certain set of challenges (Partee, 1984). When it comes to its role in the study of natural language, however, such specificity may distract from the task at hand. As Dowty (2007) puts it, instead of settling on a particular definition and aiming to determine *whether* (some fragment of) natural language is compositional or not, the question is better viewed as asking *how* natural language is compositional. In light of this, the loose formulation below is enough to serve the present purposes:

- **Compositionality:** the meaning of a complex expression is derived from the meaning of its parts and the way they are syntactically combined

While the questions of what kinds of things natural language meanings are and which syntactic analyses best yield them have been probed significantly throughout the literature, another question remains implicit: just what are the most basic, individually meaningful parts of natural language? Indeed, many studies of natural language semantics call these minimal compositional elements words, but there is currently no consensus concerning the ontological status of words (should they even exist as such).

Morphology is a subdiscipline of linguistics that studies words and their structure, and therefore helps in answering part of this question. Accounts of word structure vary widely, but the theoretical interest is in a large part concerned with how this structure determines the role of words within sentences and throughout the lexicon. Different views on the matter hark back to antiquity, with the Graeco-Roman tradition taking the latter role to be of greater importance, where words are best seen as wholes relating to those of similar form and meaning, while the Sanskrit tradition gives priority to the former, where the (proper) parts of words individually contribute to forming sentences (Blevins, Ackerman & Malouf, 2018). Conversely to the semantic analyses mentioned

above, morphological theories rarely seem to foreground questions of meaning when it comes to the analysis of word structure.

In compositional theories of interpretation that treat natural language fragments, there is always a set of atomic expressions, together with their basic meanings, often referred to as the lexicon. This, in combination with a theory of syntax, provides the compositional architecture according to which the meanings of complex expressions are derived. It is with respect to these atoms and the lexicon that morphology becomes relevant to compositional semantics, and vice versa. Many languages possess rich morphology, where the internal structure of words becomes very complex, and these complex words may individually express meanings that take several words to express in other, morphologically simpler languages. This phenomenon is referred to as ‘synthesis’ (Sapir, 1921), such as in the Latin *dormivero*, meaning “I will have slept”, where the individually meaningful (to some extent) words in the by contrast ‘analytic’ language English are instead ‘synthesised’ into one word in Latin. The question arises of whether whole words themselves should be seen as syntactically atomic, meaning the same as what may be expressed by several words in another languages, or whether they are better taken to be decomposable, with their parts in fact constituting the basic parts whose meanings correspond to what is expressed by simplex full words in other languages. Either option entails significant consequences for compositionality, as the atomic expressions to which meanings are associated differ, and the basic meanings that are associated also vary accordingly.

A proper investigation of the effects of either choice is thus warranted. This demands the articulation of a theory of grammar onto which the morphological phenomena can be mapped, which takes up the rest of this chapter. The following chapter provides both an outline of the space of morphological theories from which to choose, and an evaluation of them in terms of the conception of the grammar and the resulting criteria of assessment spelled out below.

A theory that does not decompose morphologically complex whole words settled upon at this juncture. A different semantic role for word-internal structure must therefore be proposed, where the consequences and an implementation of this provide the subject of the remainder of the thesis.

Chapter 3 defends the word-based conception of morphology, and elaborates on one such theory from which the overall discussion takes inspiration, treating its philosophical relevance in turn. Lastly, in chapter 4, a formal semantic implementation informed by this theory is attempted for a range of morphological phenomena. This is supplemented by a brief consideration of the relevance of the adopted position to the issues of diachrony and typology.

The second and fourth chapter form the main body of the thesis in terms of the novel theoretical positions developed regarding the relation between morphological theory and a compositional semantics; the former constitutes the more negative part of the exposition, being somewhat exegetical in nature, while the latter provides the more positive side, contributing novel data and analyses to the discussion. The third chapter may be seen as somewhat ancillary, providing a discussion of certain, perhaps more peripheral threads of interest.

In the present investigation, the data considered are the ways in which basic meanings appear to be encoded in the lexicon cross-linguistically, although it is ultimately the theoretical accounts of this range of morphological phenomena which are up for discussion. The employment of a particular theory of the grammar that derives from the literature on Direct Compositionality (Jacobson, 2002) along with a set of theoretical criteria closely related to such a conception should, it is hoped, help to achieve the goal of maintaining parsimony, clarity and concreteness. The evaluation of morphological theories in chapter 2 is undertaken with a similar aim at heart.

1.1 A Conception of the Grammar

The theory of the grammar articulated in Dowty (2007) provides a starting point from which to assess the status of morphology with respect to compositionality, and brings with it a number of helpful criteria. Three levels of the grammar are distinguished: meaning, *tectogrammatcs*, and *phenogrammatcs*. The first of these levels amounts to a model-theoretic interpretation with the relevant units and operations as is typical in the Montagovian tradition. The distinction between the latter two levels originates in Curry (1961) and reflects what Montague (1970) captured in the opposition between syntactic rules and syntactic operations: the former is akin to abstract, combinatoric phrase-structure rules that take as inputs and return as output syntactic categories, often according to a system of categorial grammar; and the latter relates to the actual, concrete linguistic manifestations upon application of the categorial rules, such as through word order and prosody. Morphology, and in particular inflectional morphology (see section 1.2.1) tends also to be included as a part of the phenogrammatcs, but this is an assumption that is critically assessed here. Crucially, theories tend to associate a unique syntactic operation to each syntactic rule.

Following the tradition of Direct Compositionality (Jacobson, 2002), it is desirable for the interface between the syntax and semantics to be maximally transparent, partly because this is simply parsimonious and not stipulative, and partly because the arguments for a more complex architecture involving intervening interfaces lack weight (see Jacobson, *ibid.* for details). Directness concerns the (lack of) postulation of any intermediate levels to the grammar beyond those specified above, in order to impose theoretically the desired transparency, although in its original formulation it relates only to the interface between the tectogrammatcs and the semantics. This results in the syntax and semantics working *in tandem*, given the homomorphic characterisation of compositionality: as the syntax builds constructions (or proves them well-formed), the semantics interprets them immediately and in lockstep. An important correlate of this is the Hypothesis of Local Interpretation: “each linguistic constituent has a well-formed and complete denotation that does not depend on any linguistic element external to that expression” (Barker, 2002).

Locality is implicit in a slightly different but closely related sense too, in terms of the immediacy with which the interface operates. The strength of

directness may be thus articulated, as there remain options as to how complex and rich the syntax may be. Minimally, the syntax need only ‘see’ the syntactic categories of the immediate constituents it works with, not requiring any additional structural information. Applications of the relevant rules may then be enacted through a simple concatenation operation on the phonologically monolithic strings that embody the constituents, resulting in appropriate word order. Working in tandem with the syntax, the semantic operations corresponding to the syntactic rules, typically functional applications, are carried out per well-formed constituent, yielding their interpretations in a stepwise manner. A strongly directly compositional grammar as a whole therefore takes as input(s) and returns as output triples of the form $\langle \textit{phonology}, \textit{syntactic category}, \textit{meaning} \rangle$, often referred to as *signs* (Jacobson, 2002).

Three relevant criteria additional to that of directness fall out of the above, although (i) and (ii) are reimagined somewhat below to best serve the present purposes:

- (i) **Compositional Transparency** (Dowty, 2007): the degree to which the compositional semantic interpretation of natural language is readily apparent (obvious, simple, easy to compute) from its syntactic structure.
- (ii) **Syntactic Economy** (ibid.): the degree to which the syntactic structures of natural language are no more complicated than they need to be to produce compositionally the semantic interpretation that they have.
- (iii) **Locality of Interpretation** (Jacobson, 2002): every constituent has a meaning.

Another criterion that is of occasional interest here is the semantic analogue to (ii):

- (iv) **Structural Semantic Economy** (Dowty, 2007): the degree to which the meanings and operations on meanings used during compositional interpretation to build up complex meanings out of simpler meanings are no more complicated than they need to be to derive, in a stepwise fashion, all the complete sentence meanings that natural languages in fact express.

Locality in the sense of immediacy outlined above is equivalent to a *context-free* level of complexity of the grammar. This means that at any point (node) in a derivation (represented by a tree), it is only the information of immediate daughter nodes themselves that are relevant to the computation of the value of the parent node, and not that of any further descendants. In the semantic case, this means that the value of some semantic operation applying to the meanings of the daughter nodes alone may not depend on the construction of the formulae that represent them in the relevant logical translation language, but only on the meanings (the model-theoretic structures) themselves.

Potentially of greater relevance here is how this manifests syntactically, where context-freeness essentially amounts to a constraint on the syntactic operations that they may only concatenate strings, at least when the grammar

operates on the above signs that build in minimal representation. Importantly, however, certain recalcitrant cases appear to require a weakening of the directness in some manner. Jacobson (2002) notes that there are two options as to how to account for such extra complexity into the overall architecture of the grammar. The first would be to drop the requirement of uniqueness of context-free syntactic operations indexed to the combinatory rules. The second would be to allow more structure into the first two elements of the triples that constitute grammatical signs: instead of phonological strings, we have phonological representations, and instead of syntactic categories alone, the syntactic rules involve structural information, represented by trees.

It is difficult to arbitrate between these options. The first results in a dissociation of the tecto- and phenogrammatcs. One of the initial motivations for the uniqueness of the indexing of syntactic operations to syntactic rules was likely some kind of extension of transparency throughout all the levels of the grammatical system, beyond the syntax-semantics interface *qua* tectogrammatcs, and on to the phenogrammatcs; the hierarchical structure and rules enacting the computation of the meanings of complex expressions from their parts are determined in conjunction with and simultaneously to the syntactic operations that produce the observed linguistic forms at each step of the derivation. If this is lost, the application of the same syntactic rule would be manifested by a simple, concatenative operation on the constituents in some situation, but by a different, non-context-free operation in another, with the choice of operation having to be determined by some relevant stipulation. The process of interpretation in the grammar as a whole would thus be rendered more opaque, with the undesirable result that the linguistic expressions themselves bear a less direct relation to their meanings.

The infixation operation of *WRAP* (Dowty, *ibid.*) is a clear example of taking the second option. It seems to be demanded by the phenomenon of ‘rightward extraposition’, such as in the following sentence:

- (1) [A woman] just came into the room [who we met at the station earlier]

The relative clause and the main clause’s subject, which it modifies, ultimately form a single constituent in the structure of the whole sentence, but one that is discontinuous due to the rest of the main clause intervening. The *WRAP* operation allows for compositional transparency to be maintained by interpreting this in the same way as if the relative clause were not extraposed. This is at the expense of syntactic economy, however, with infixation being a non-concatenative syntactic operation.

While the discontinuity needed for *WRAP* indeed adds extra structure, it is minimal, as conceivably only a break in a string is needed to determine the infixation point, when it comes to the phenogrammatcs. Moreover, concatenation is simply a special case of infixation, where one of the two strings between which constituent is infixed is the empty string. As for the tectogrammatcs, however, the break in the string corresponds to the point deeper in the tree where the discontinuous constituent is formed from the relevant parts (here,

from the relative clause and the subject of the sentence). If the combinatory syntactic rules work on trees instead of bare categories, then context-freeness can be recovered. Again, for *WRAP*, this concession may be minimal, as it appears only to require the depth of one branch to determine the infixation point. Allowing for these added representations, every syntactic rule can be uniquely associated with some infixation operation in a context-free manner.

It appears, therefore, that in order to maintain the transparency throughout the grammatical system, it is necessary to build more structure into the elements operated on by the grammar and lose some simplicity of representations, and thus syntactic economy. While this is minimal for the case of *WRAP*, it may be that more recalcitrant cases require significantly more representation to a theoretically undesirable extent. The following section demonstrates how morphology appears to provide such phenomena.

1.2 The Relevance of Morphology

The issue of morphology can be stated clearly with respect to the above trade-off between (i) and (ii). If complex word-forms¹ can be decomposed, as is proposed in certain morphological theories, and their constituent parts are thus elements in the compositional process of interpretation, then the range of non-concatenative structure demonstrated cross-linguistically calls for an account that respects both (ii) and (i) as much as is possible.

It is mooted by Dowty (*ibid.*) and seemingly tacitly assumed elsewhere that (inflectional) morphology in general constitutes but a phenogrammatical reflex of sorts in natural language, with little to no relevance to the levels of the grammar on which the process of interpretation actually depends. This may be a result of the fact that the languages with which formal theories of natural language interpretation have tended to be concerned —often English —are morphologically impoverished. Dowty (*ibid.*) concedes, however, that when languages have rich morphology and correspondingly freer word order, we ought to expect them nonetheless to be beholden to the same level of compositional transparency as morphologically simpler languages (never mind that the other criteria ought still to apply too). Dowty (*ibid.*) conjectures that this is likely a problem best suited for psycholinguistic inquiry, where it becomes a question of how parsing of various actual linguistic structures can reveal the underlying tectogrammatical structure on which the interpretation in fact depends.

That puts off a proper grammatical investigation with respect to a theory of interpretation of such phenomena as they occur cross-linguistically. If inflectional morphology is observably directly relevant at the syntax-semantics interface, then its contribution cannot be isolated at phenogrammatical level. Such relevance is conceivably self-evident in certain cases and may be argued for in less obvious cases (see chapter 4). Something of a necessary interdependence

¹Throughout this thesis, the term *word-form* is used to refer specifically to the linguistic expression, or form, constituting a word, considered apart from its meaning, while *word* itself is neutral with respect to the form-meaning distinction.

between the syntax-semantics interface in the abstract and morphology as it is phenogrammatically manifested appears, therefore, to be demanded.

Such interrelation also follows from the uniqueness requirement for the association of syntactic rules to context-free operations, extending the transparency at the syntax-semantics interface to the phenogrammatcs, thereby engendering some necessary degree of interdependence. Indeed, (i) can be extended such that it is not only abstract syntactic structure that is relevant to determining the apparentness of the compositional semantic interpretation, but also the way in which the linguistic forms themselves manifest this structure. Namely, the phenogrammatically realised morphological markings may thus ultimately be of compositional import. The resultant question is of just what this import is, and it will be important in providing such an account to adhere to the above criteria.

This conclusion is in consonance with Jacobson’s (2004) suggestion that “we maintain that it is the actual (pronounced) expressions of a language which are compositionally interpreted”, albeit in a slightly different sense. Instead of maintaining this through a *vertical* transparency, where no more levels of the grammar are postulated than required, a more *horizontal* transparency is proposed here, which can be viewed as being somewhat analogous to strength of directness. In this sense, transparency concerns the richness of representations, where the obfuscation of the role of the foremost elements (read “actual (pronounced) expressions”) involved in the process of interpretation through a dependence on arguably unnecessary and stipulative abstract theoretical entities ought to be kept to a minimum.

(ii) and (iv) are thus important in providing theoretical constraints on what representations at the different levels may be, and what we take the operands of the grammar to be constituted by. This is especially true when (ii) is seen to incorporate the phenogrammatcs, in much the same way as for (i), where this is implicit in the relation between context-freeness and concatenative syntactic operations. The novelty here is a result of considering the phenogrammatcal level on the same footing as the other levels of the grammar, in lieu of the goal of minimising the complexity of the grammar as a whole. To reemphasise the above discussion of the interdependence between the phenogrammatcs and the (tectogrammatcal) syntax-semantics interface, an equivalent extension is proposed for (i). We might thus demand at every mention of syntactic structure in our above criteria, that the ideally context-free operations concretely manifesting this structure be accounted for too.

When it comes to morphology, this boils down to the question of what items are deemed to be basically meaningful, as elements of the lexicon. Can sub-word units —if morphologically complex whole words are indeed syntactically decomposable —be associated with meanings in a stable manner that does not require building in an undesirable amount of representation for the grammar to work with? While (i), (ii) and (iv) are posed separately from one another, they may be seen to coincide to a given degree in the lexicon, as its items are both syntactically atomic and semantically basic, meaning that the compositional machinery depends on these elements to provide the primitive representations

with which it works. This is why both a theory of morphology and a compositional theory of interpretation, taken together, are required to probe the question at hand.

1.2.1 Inflectional Morphology

Inflectional morphology concerns the expression of morphosyntactic properties (singular, plural...; present, past, future...; masculine, feminine,...; nominative, accusative,...; and so on) as instances of inflectional categories (respectively: number; tense; grammatical gender; case; and so on) by the relevant word-forms through their internal structure. While the meaningfulness of certain of these morphosyntactic properties, and as a result the words, too or their parts that express them, may be up for debate (see chapter 4), some inflectional morphology has seemingly undeniable semantic import.

A straight-forward example that applies even to English would be that of plural nouns, where the morphological encoding of this morphosyntactic property, most commonly the suffix *-s*, appears to carry the information, in some manner, that the noun attached to it (at least possibly - see section 4.1) denotes a non-atomic entity.² Even in English, however, cases that challenge the minimality of our grammatical representations exist. Irregular, non-concatenative plurals such as *mouse-mice* are not too uncommon. If these are decomposed into syntactically atomic parts that are basically meaningful, in the same way as is assumed for the suffixal cases, an increase in the complexity of syntactic operations would be required. It may be that the different medial vowel sounds are what carry the singular and plural meanings respectively, while the surrounding, discontinuous consonants carry the lexical meaning. Unlike before, however, the source of this discontinuity is purely phonological. This is arguably undesirable, as the non-concatenative operation is thus syntactically and semantically unmotivated.

Suppletion, where two inflectional forms of the same word are phonologically unrelated, such as *person-people*, or *go-went*, obfuscates the issue of morphological decomposition further still. It appears impossible to maintain that the proper tectogrammatically relevant syntactic parts are equivalent to their linguistic manifestations. Which parts carry the lexical and morphosyntactic number meanings respectively? And moreover, when they are composed, how does the relevant syntactic operation return the appropriate phonological forms? One way to solve this would be to add another level to the grammar, in-between the phenogrammatics and tectogrammatics, mediating between the two, but this of course flies in the face of directness, especially as it is taken here to extend

²Even in this apparently simple case, not all is as it seems. The suffix takes a number of different forms in particular phonological environments. Compare the variation between *bet-s*, where it is the phoneme /s/ that is appended to a voiceless consonant, *bed-s*, where /z/ is beside a voiced consonant, and *princess-es*, where // is epenthesised before /s/ due to a clash of similar consonants. This is not to mention irregular cases such as *knife-knives*, where the preceding consonant is modulated as well. Postulating the suffix as a standalone meaningful element clearly requires some phonological abstraction to account for its various incarnations in a uniform manner.

to and include the former. Abandoning that option, so much additional representation would then be required to maintain compositional transparency and in turn drastically diminish syntactic economy. While an account that includes suppletion may seem like generalising to the worst case, especially given how rare a phenomenon it is, if we wish for our theory of the grammar to be as complete as possible, then to stipulate that suppletion has no place in it and is merely an exception could be deemed unsatisfactory. Suppletion is of course never systematic, however, as is to be expected taking pressures such as learnability into account, and this perhaps lends credence to the attribution of the phenomenon as exceptional.

Other forms of non-concatenative morphology may be highly systematic, though. Reduplication, where some part of or the whole of a word-form is repeated in constituting a new inflectional form, is often seen to be systematic, for example in forming plurals in certain Austronesian languages such as Malay: *rumah* means “house”, while *rumah rumah* means “houses”. Subtraction, also known as disfixation, is when part of a word-form is deleted to constitute an inflectionally related form. While not systematic, though not an entirely isolated case throughout the language, a clear example comes from the French *os*, meaning both “bone” and “bones”, where the final consonant is pronounced for the singular, but is silent for the plural. Root-and-pattern morphology provides further examples of recalcitrant non-concatenative phenomena, where, in languages such as Arabic, for example, lexical semantic information appears to be encoded in a ‘templatic’ consonantal root, such as *k-t-b*, while the pattern of infixes and affixes to this template determines fully specified inflectional and derivational forms; in general, the abstract root appears to have the loose meaning of having to do with writing, but there are indeed inflectional distinctions, such as between *kitāb* meaning “book” and *kutub* meaning “books”. In all of these cases, the questions of which parts to isolate, and how they are individually meaningful become highly pertinent.

The purpose of this extremely brief exposition of a rudimentary fraction of the range of non-concatenative morphological phenomena displayed across natural languages is not to invite a grammatical analysis for each particular case; rather, it serves to raise the issue of the relevance of inflectional morphology to our conception of the grammar. In any attempt to construct a formal analysis of the above examples, concessions have to be made with respect to the guiding criteria that inform what the inputs and outputs of the grammar should be. The nature of these concessions depends primarily on the theory of (inflectional) morphology that is adhered to, as it provides the interface with the rest of the grammar, namely through an articulation of what units are deemed to constitute the lexicon and thus be syntactically atomic and basically meaningful as the fundamental inputs to the grammar. An outline of the range of morphological theories is provided in the following chapter, with respect to which an assessment can be carried out of how they fit together with the conception of the grammar provided above and the extent to which they fulfil the criteria.

Another question that arises is what the semantic import of different inflectional categories may be, if it can be maintained that they are indeed meaningful

in the first instance. The choice of certain theories of morphology and their resultant grammatical status may suggest that semantically murky inflectional categories such as case and grammatical gender, as well as the related phenomenon of morphosyntactic agreement, should indeed have some role to play in the process of interpretation. This is treated in chapter 4.

1.2.2 Word-Formation and Derivational Morphology

A traditional split in the study of morphology exists between inflectional morphology that pertains to morphosyntactic properties of related word-forms, as outlined above, and ‘derivational’ morphology, or more generally ‘word-formation’, that is concerned with the the formation of new words that may have their own inflections, widely referred to in the literature as ‘lexemes’. A lexeme is a theoretical abstraction from the more concrete entities of inflectionally related word-forms. For example, the abstract lexeme **write** manifests as the inflectionally related word-forms *write*, *writes*, *wrote*, *writing* and *written*. The word-forms are related according to the relevant inflectional categories, such as person, number, tense and aspect. Derivational morphology is thought, however, to constitute new lexemes. For example, **writer** is one possible derivation, which has its own inflectional forms, such as in the plural *writers*. **Rewrite** is another derivation, again with its own inflections. Word-formation is a generalisation of this phenomenon, where more than one lexeme may be involved in the creation of a new one, such as in the formations **ghostwriter** and **screenwriter**, again each with their own inflections.

Derivational morphology and word-formation sometimes result in a change in syntactic category, but they always lead to considerably less semantically transparent effects than inflectional morphology appears to. The meaning of a derived form may change drastically from referring to an entity instead of an event, a property instead of an entity, or something more subtle, and potentially (formally) semantically inarticulable such as for diminutives, or certain prefixes such as “anti-”, “meta-”, and so on. The present goal is not to develop an analysis of just how to interpret derivational morphology; rather, it is only to highlight that if the compositional status of word-internal structure is under investigation, then the full range of possibilities of the meaningful contributions of that structure and the conceivable parts ought to be borne in mind, where derivational morphology presents a series of significantly challenging and obscure possibilities.

Chapter 2

Morphological Theory

To conduct a survey of several different specific morphological theories would require too much time and space for the present purposes. Instead, the taxonomy of morphological theories of inflection provided in Stump (2001) provides the scope of theoretical possibilities that can be assessed with respect to a compositional theory of interpretation in line with the above conception of the grammar and the accompanying criteria. Accordingly, the association of morphosyntactic properties to morphological units is taken throughout the following discussion to be equivalent to providing the corresponding semantic information (for example, plural number bearing plural meaning), such that the interpretation procedure builds this information in accordingly.

The important questions are as follows: firstly, what are the syntactically atomic parts assumed to make up the lexicon in a given morphological theory?; secondly, what kinds of meanings are associated with these elements stored in the lexicon?; and lastly, what kind of operations are required on (the meanings of) these atomic parts to derive (the meanings of) more complex expressions? Throughout, more general theoretical criteria of parsimony, concreteness and transparency also play a role when arbitrating between the respective consequences engendered by different theories.

In brief, the following discussion concludes with the judgement that any theory which decomposes morphologically complex words in some way is likely not to fulfill some or all of the above criteria satisfactorily, despite the advantages that they may offer in certain other respects. By contrast, a theory which asserts that sub-word units of whatever kind are not of direct compositional semantic import suffers no such problems (as in section 2.2.4). This is because fully realised word-forms plausibly provide the most stable and coherent morphological unit with which lexical and morphosyntactic meaning may be identified. Without any word-internal compositional interpretation, it is assumed that the relations amongst these forms, holistically throughout the lexicon, serve to organise and distinguish their meanings systematically. The questions that arise upon the adoption of such a theory, as well as the outline of a formal semantic implementation, are the subject of the subsequent chapters.

2.1 A Taxonomy of Morphological Theory

The taxonomy presented in Stump (2001) involves two dimensions of theoretical variation along each of which two alternatives are distinguished. The first axis captures the difference between *lexical* and *inferential* theories. What is at stake here is which forms are deemed basic and thus stored in the lexicon. This means that the associated morphosyntactic properties and meanings are of secondary importance, with them instead being the concern of the second axis. To avoid possible confusion down the line, the term *lexical* will be replaced by the term *morphemic* here. In morphemic theories, it is minimal sub-word units that are stored in the lexicon individually, whatever their meaning. Such sub-word units are typically referred to as *morphemes* (hence the name of the category here), which may encompass any supposed minimal syntactic unit (that may or may not be phonologically realised, and/or basically meaningful).

Such theories are necessarily ‘syntactocentric’, as morphemes are inserted in the relevant (word-internal) structure that determines the role they play in constituting the larger expressions that they make up. This structure, and how the morphemes are inserted into it, is motivated by independent principles of syntactic structure, instead of being dependent on any morphology-specific explanations. The relevant questions that arise for such theories are which sub-word units, or morphemes, comprise the lexicon, and how does syntax manifest the construction of larger expressions, namely words, in which the morphemes appear.

Inferential theories are by contrast ‘morphocentric’. Relations between word-forms are the primary concern of such theories, and these relations are characterised according to some kind rules of inference that yield the appropriate shape of the word-forms. The nature of the rules of inference determine what is ultimately taken to be stored. Importantly, however, the parts of words that express morphosyntactic properties, such as *-s* in English expressing plural (as well as certain other morphosyntactic properties, or sets thereof), are not stored individually. These parts of words that appear to express these properties are still theoretically relevant, and tend to be termed *exponents* instead of morphemes. Exponents represent relations amongst like word-forms, and thereby inform the inferential process that yields the word-form.

Two kinds of morphological rules of inference may be distinguished. Firstly, there are *rules of exponence*, which are applied in a linear manner, manifesting the exponents of the morphosyntactic properties onto the lexical ‘root’ or ‘stems’ of a given lexeme, ultimately to yield the fully realised word-form. An example of the application of a rule of exponence would be to deduce the plural form *dogs* from the stem *dog-*, by way of the rule of exponence that yields the plural as output through providing the exponent *-s*. This *-s* holds no independent theoretical status, however, and it is not stored in the lexicon nor is it in any way individually meaningful. It represents but an abstraction from the many applications of this (kind of) rule, which takes as input stems such as *dog-* and *cat-* and returns the whole word-forms *dogs* and *cats*, and in so doing associates the plural number information to the morphosyntactically bare, or underspeci-

fied lexical root. Application of this specific rule of exponence fails, however, in cases such as *mice*, where a different rule of exponence would be required, the same that yields *lice*, modifying the medial vowel in a specific manner. While the rules that yield *dogs* and *mice* associate the same morphosyntactic property, it is the fact that they relate different patterns of word-forms (phenogrammatically) that takes theoretical precedence, and as such they represent distinct morphological inferences, and thus two separate rules. It is thus the range of relations amongst word-forms that determine scope of rules of exponence, reflecting the morphocentric character of the inferential component in this guise.

As a sidenote, the difference between a root and stem is subtle. A stem may be seen as a word-form without some or all of its exponents, but with a distinct phonological realisation. Roots abstract over all the possible stems of a given lexeme, and are therefore somewhat analogous to a lexeme, although they play a different role theoretically. To provide an example, in Spanish, the verb *jugar*, meaning “to play”, has two stems: *jueg-* used in the present indicative for all the singular forms, as well as the third person plural, such as *juego* meaning “I play”, while *jug-* is used for the remaining plural forms, such as *jugamos* meaning “we play”. If exponents express the morphosyntactic properties, and the (basic) stems, the lexical meaning (although the stem distinction arguably carries some morphosyntactic information as well), the root may be seen in some sense as an abstraction over specific forms that express lexical meaning. In that sense, it coincides somewhat with a lexeme, but the senses differ in that a lexeme is an abstraction over sets of fully realised word-forms, whereas a root is abstracted over stems. Certain theories —particularly less morphocentric ones where stems are no longer so relevant —simply see roots as the bearers of lexical meaning

Importantly, no theoretical precedence is necessarily given to concatenative morphology by rules of exponence, as the linear steps in the process determining new stems or forms are typically taken to be phonologically constrained only by the (systematic) relations of word-forms to one another across the morphological system. It is thus relations amongst word-forms, as manifested by stems and exponents, that provide the space of possible rules, and these rules in turn define the morphological component. These relations are ‘syntagmatic’, however, in that they are determined by the linear arrangements of morphological units, internally to word-forms through identifiable exponents, not that this necessarily privileges concatenative structure.

Rules of exponence are distinguished from *paradigmatic* rules, which relate word-forms externally, according to their place in the paradigmatic structure of the lexicon as a whole. Each full word-form occupies a cell in a tabular structure known as a paradigm, that is in essence a matrix for each lexeme yielded by the possible combinations of morphosyntactic properties given the inflectional categories that are morphologically encoded in the language. In the English noun system, given that only one number is marked, each nominal lexeme’s paradigm consists of two cells, as displayed in Table 2.1.

By contrast, nouns in Danish, for example, inflect for both number and definiteness, with one distinction for both, yielding $2 \times 2 = 4$ paradigm cells for Danish nouns. Below in Table 2.2, we see the paradigm for the Danish *hånd*,

SINGULAR	PLURAL
dog	dogs

Table 2.1: English noun paradigm

meaning “hand”. In some cases, two cells may be expressed by the same form, which is known as ‘syncretism’, such as for Danish *år*, which means both “year” and “years”, or simply the English *sheep*, which is both the singular and plural form.

	SINGULAR	PLURAL
INDEFINITE	hånd	hænder
DEFINITE	hånden	hænderne

Table 2.2: Danish noun paradigm

Proportional analogy across related paradigm cells is the inferential method, or rule taken to yield (new) word-forms. For example, if the plural and singular cells of the lexeme **dog** is known, and one wishes to determine the plural cell of the lexeme **cat**, given knowledge of the singular form, then by analogy with respect to both the morphosyntactic properties and the related pattern(s) of forms expressing it, the likely form *cats* may thus be deduced. More specifically, the following proportional analogy between singular and plural forms, namely *dog : dogs :: cat : ?* may be solved to yield the correct plural form.

The patterns of relations amongst fully realised word-forms are therefore of theoretical precedence, relating, for example, *mouse* with *louse*, due to the shared plural forms *mice* and *lice*, but these are unrelated to *house* and *grouse*, and essentially all other nouns, as the pattern does not extend to them. Conversely, *mice* and *lice* are related due to the shared form of their singulars. Switching in *mouse* for *cat* in the above proportional analogy would yield the incorrect plural form **mouses*, so such a paradigmatic rule would ultimately be ruled out. However, the fact that the incorrect plural form is not part of the language means that the rule is also not considered part of the language - rather, it is a possible inference a learner with insufficient linguistic knowledge could make.

Lexemes are only of secondary theoretical status with respect to determining the relations between paradigm cells, as it is instead the actual word-forms themselves that constitute the nodes in the pattern of relations, a network that is structured by the given language’s inflectional system.¹ The morphological system of the lexicon is in turn constituted by these networks of relations over paradigms. The application of proportional analogies is what is taken to be the way in which paradigmatic rules can be realised linguistically, where paradigm

¹It was claimed in section 1.2.2 that derivational morphology relates lexemes. If we wish to background lexemes theoretically, and reinforce their abstract and derivative status by comparison to fully realised word-forms, then we may just as well state that derivational morphology relates sets of word-forms, or (sub-)paradigms themselves.

cells themselves are involved as opposed to stems and exponents, as is the case for rules of exponence.

The terminology can be misleading, as paradigmatic rules represent more so patterns of relations across paradigms than any kind of rule *per se*, beyond licensing proportional analogies. Certain theories that lend a more prominent role to lexemes may take it that paradigmatic rules are in fact rules more properly conceived, represented by so-called ‘paradigm functions’. Paradigm functions take as input a pair, one element of which being the lexeme and the other, a set of morphosyntactic properties, and return the word-form that is realised given the coincidence of those properties for the lexeme, which may be inferred both from the rest of the lexeme’s paradigm, and from other inflectional forms of other lexemes that share some or all of the set of morphosyntactic properties specified.

In the following, such a conception of paradigmatic rules is not assumed; the alternative view that they represent simply the patterns of relations is preferred. This is because the proportional analogies licensed across word-forms that represent the active side of inferential component may be seen to yield new word-forms without recourse to reified lexemes and morphosyntactic property sets. These are better seen as theoretically useful abstractions from the inherent relations amongst both the forms and the lexical and grammatical meanings of the actual, most concrete linguistic elements, namely the word-forms themselves. See section 3.2 discussion of some of the theoretical arguments against paradigm functions from the perspective of the word-based morphological theory adopted in light of the evaluation in the following section.

The choice as to whether rules of exponence or paradigmatic rules are preferred is for a theoretician to make, with each having its own set of qualities that could make it suitable for different purposes. Indeed, some theories may wish to include both. It is observed below that theories to which rules of exponence are better suited tend to see morphosyntactic properties as manifesting through parts of word-forms, whereas those to which paradigmatic rules are better suited see morphosyntactic properties as being associable to whole word-forms. Such a distinction brings us on to the second axis of the taxonomy, but let us first recapitulate what is implicated in first axis.

It is most useful to reflect on the differences between what is stored, how so, and what is derived in morphemic and inferential theories. In morphemic theories, morphemes are stored, such as root morphemes typically carrying lexical meaning, *e.g.* *dog-*, and inflectional morphemes that bear the meanings of morphosyntactic properties they express, such as the suffix *-s* for plural, or a phonologically empty suffix for the singular. These may be concatenated simply upon the application of a combinatoric syntactic rule that takes inputs of the appropriate categories (say, noun root and number morpheme) and derives the plural form *dogs*, with the corresponding lexical and morphosyntactic aspects of its resultant composed meaning. When it comes non-concatenative morphology, morphemes may either be considered as being abstract, with their ultimate linguistic manifestations handled by some other component of the grammar separate to the syntax, or the syntactic operations on morphemes may build

in richer phonological representations. This is to be considered further in the evaluation below.

By contrast, for inferential theories involving rules of exponence, a basic stem, or set of basic stems bearing lexical meaning may be stored similarly, but instead of any inflectional morphemes bearing morphosyntactic properties, a rule of exponence, *e.g.* $\Rightarrow_{\text{PL};-s}$ would be stored, which could be applied to the stem *dog-* to derive the plural form *dogs*, *e.g.* *dog-* $\Rightarrow_{\text{PL};-s}$ *dogs*. Importantly, the application of such a rule is conceived of as an inference to the plural form from the basic stem, where the lexical meaning is consistent throughout. Therefore, this rule of exponence yielding the inference to this regular plural form would be distinguished from those applied in irregular formations. As for the plural meaning, it is likely the case that it is only part of the full plural form, following the application of the rule that associates it to the lexical root.

A question therefore arises as to whether there is an independent, concrete element of the language with plural meaning, as for a plural morpheme, or whether it only appears given a derivation; the rule conceivably only has the potential to associate the morphosyntactic property upon its application, without necessarily bearing the meaning itself. This is different from an inflectional morpheme which bears the morphosyntactic property itself. While both may ultimately be functions semantically that associate plural meaning to the lexical item similarly, the theoretical senses are distinct, as the rule itself never materialises linguistically, only its effect. It is nonetheless deemed a part of a speaker's linguistic knowledge, and thereby stored appropriately, so it is reified theoretically at least. The benefits and pitfalls of these alternatives are explored further at the relevant points in the evaluation below.

Again, this is to be distinguished from an inferential theory involving paradigmatic rules, where there are no roots, morphemes, or rules stored individually, but rather only whole, potentially morphologically complex word-forms, such as the singular *dog* and plural *dogs*, along with their respective meanings. These and other singular and plural forms are related by the paradigmatic structure of the lexicon, organised according to the inflectional categories encoded (number alone for English nouns). This provides the patterns according to which inferences can be made as to the forms of any as yet unencountered, and therefore empty paradigm cells. Therefore, there is decidedly no individual linguistic element bearing singular or plural meaning; such meanings are only a part of the meanings of full word-forms along with their lexical meanings.

The second axis concerns the units with which morphosyntactic properties are basically related. *Incremental* theories are thus distinguished from *realisational* theories, where, in the former, it is the sub-word units themselves, be they concrete or abstract morphemes, or only stems of increasing complexity, that are associated with the morphosyntactic content, while for the latter, the association of morphosyntactic properties to fully phenogramatically realised word-forms logically precedes their expression through any of the relevant parts.

Importantly for the present purposes, in incremental theories, complex combinations of morphemes, or stems acquire their morphosyntactic properties cumulatively, step-wise from the combination of the individual morphemes, or

through rules of exponence. *Dogs* is a plural form only given some process that has attached the morphosyntactic property PLURAL to it by way of some sub-word unit (through a concrete or abstract morpheme carrying number information, or through a rule of exponence applied to the stem *dog*).

For realisational theories, the association of the morphosyntactic properties to the word-form itself depends of course also on whether its form is conceived of morphemically or inferentially. In either case, however, the grammar is taken to govern the sets of morphosyntactic properties it realises, while the determination of the form itself is what is relevant morphologically with respect to the given set of morphemes, or the given lexeme. The grammar specifies that there is a morphosyntactic property PLURAL to be manifested linguistically in some way, and the morphology specifies that the forms *dogs*, *cats*, *mice*, *people* each realise this property, either with respect to the given lexeme or root morpheme.

2.2 Evaluation

The dimensions of theoretical variation as presented above correspond to the difference between a theory being in the one sense ‘morphotactically’ either morpheme- or word-based, and in the other sense morphosyntactically morpheme- or word-based (Blevins, Ackerman & Malouf, 2018). The former regards the theoretical status assigned to units of form, while the latter concerns which units have grammatical import and how morphosyntactic properties are thus associated to words. This coincides somewhat with the distinction between phenogrammatcs and tectogrammatcs (or the syntax-semantics interface in general), particularly when the matter of the (compositional) semantic import of sub-word units is raised.²

It is noted that it is desirable for a theory be *consistently* either morpheme-based, or word-based, in both senses. This is consonant with the point belaboured in section 1.2 that the phenogrammatcs and the rest of the grammar are interdependent. Seeing as morphemic and incremental theories tend to be morphotactically and morphosyntactically morpheme-based respectively, while inferential and realisational theories appear more word-based in the respective dimensions, when observing the possible combinations of the above kinds of theories, it appears that the Morphemic-Realisational, and Inferential-Incremental alternatives are disfavoured from the off. The former remains influential, however, and the latter serves to demonstrate what a morphocentric theory that considers complex word-forms to be decomposable involves. They receive a proper treatment, therefore, following the initial evaluation of Morphemic-Incremental theories. This order of discussion serves to set up the theoretical problems that derive from an analysis that is entirely morpheme-based, so that it can be explored whether and how these problems may be treated.

²Interestingly, and perhaps somewhat anachronistically, Curry (1961) originally articulated the distinction between phenogrammatcs and tectogrammatcs by analogy to morphophonemics as opposed to morphology.

The above example of *mouse-mice* serves below as a starting point for discussion, due to the way in which its simple and neat non-concatenative morphological structure encodes the simple morphosyntactic distinction between singular and plural. Other examples of morphological phenomena are also referenced to highlight specific issues. The question of how the number information is incorporated with the lexical meanings of the whole words *mouse* and *mice* is what is in question. Each kind of theory has something different to say on the matter, although this may be limited to a discussion of the relation between morphology and syntax alone, perhaps only with implicit consideration of the semantics. Making explicit the relevance of the theoretical discussion to a compositional theory of interpretation is the addition of the present investigation. What appears to be required specifically for each theory remains subject to evaluation with respect to the criteria outlined in the first chapter.

It will first be necessary, therefore, to give an interpretation to the whole words so that the semantic targets of the process by which these meanings are arrived at is known from the off. There are several different accounts of the semantics of number distinctions, but one is chosen here for expository purposes, as the ultimate interpretations of the whole words is not of primary interest. Rather, once the interpretations are settled on, whatever they may be, an analysis of how these meanings and the forms expressing them may be arrived at in each case given the variation of (morphological-)theoretical assumptions can proceed. The details of the semantics adopted are laid out in the following section.

Prior to getting underway with the evaluation, it will be useful to recapitulate the space of theoretical alternatives in the taxonomy, in the order that they are treated, by briefly stating in the abstract how each kind of theory conceives of the items stored in the lexicon and associates morphosyntactic properties (or composes meanings). Following Stump (*ibid.*), let us talk of a word-form w , a set of morphosyntactic properties, σ , and where relevant, a lexeme L

- **Morphemic-Incremental:** w is associated with σ in the syntax, where the syntactic structure yielding w is built in lockstep with the association of individual morphemes bearing the properties in σ to w 's lexical root, whose form is altered at each step in the derivation.

Ex.: *Dogs* is constructed from two daughter nodes, DOG-, and -S, bearing lexical and plural meaning respectively and being composed at parent node. Both morphemes are morphotactically concrete and concatenated at parent node.

- **Morphemic-Realisational:** w is also associated with σ in the syntax, where there is a single node comprising the structure that associates the properties in σ to w 's (abstract) lexical root. At this node, w is formed through a process that selects the appropriate morphemic forms given the set of properties and the lexical content.

Ex.: *Dogs* comprises the structure with two branches, one with an abstract lexical root morpheme **dog** and the other with the abstract plural

morpheme [PL]. The whole abstract structure is sent to a phonological read-out that provides the forms *dog* and *-s* to yield whole word-form. The structure is interpreted separately.

- **Inferential-Incremental:** the association of σ to w occurs morphologically, where w acquires the properties in σ given an appropriate series of steps that add to a basic stem of w in an inferential process that ultimately yields w

Ex.: *Dogs* is ‘inferred’ from the linear application of a plural exponent rule to the stem *dog-* that yields the output of the full form with the appropriate exponent on it, and associates plural meaning to it.

- **Inferential-Realisational:** the association of σ to w is a feature of the morphological system where w inhabits a cell in a paradigm (or specifically in the paradigm of the lexeme L) that is the intersection of the set of properties in σ . If w is not simply stipulated in advance, it is inferred from $\langle L, \sigma \rangle$, for example, by proportional analogy from other appropriately related forms.

Ex.: *Dogs* is listed in the lexicon already. If it is not known, it may be inferred by proportional analogy from its singular *dog* and other like singular-plural pairs, such as *cat-cats*, *ferret-ferrets*, and so on.

It should be borne in mind throughout the following discussion, that, while arguments concerning the kinds of morphological theories treated here may be found in the relevant bodies of literature, the field of morphology does not, in general, foreground issues relating to (compositional) semantics. By the same token, the range of theoretical perspectives on morphological structure is also not typically considered in the literature on the compositionality of natural language. This means that the mapping of these kinds of morphological theories onto a putative compositional process of interpretation explored in the following 4 subsections is a novel addition of the present investigation.

2.2.1 Morphemic-Incremental

Morphemic-Incremental theories see words as decomposable into minimal meaningful constituent parts —morphemes stored in the lexicon with their basic meanings —and then aim to derive the form and meaning of larger expressions, namely morphologically complex words, through a (syntactic) process of composition equivalent to that operating at the sentence level. Such a theory is parsimonious, because it postulates one set of combinatory syntactic rules that work internally and externally to words. One example would be Lieber (1992), which derives from the school of generative grammar. Indeed, such theories involve an extension of the application of syntactic principles from sentential to word-internal structure, where the latter may undergo a compositional interpretation in much the same manner. It is presupposed here that the syntax-semantics

interface is the same as that postulated for the rest of the grammar: direct, compositionally transparent and syntactically (and semantically) economical. This differs of course from the typical generative theories that postulate an intervening ‘logical form’ (LF) between the syntactic structure and the semantics.

Roots (or root-morphemes) are thus composed with inflectional morphemes to form word-like syntactic structures in the same way that (morphologically simple) words form sentences, and all are interpreted in tandem with their construction, which allows for directness to be broadly maintained. Tectogrammatically, this demands that roots and inflectional morphemes be of the appropriate syntactic categories so that they combine only with one another, and not with other words, until a full word is formed. This may only be a small concession theoretically, particularly given the recursive manner in which syntactic categories can be derived in systems of Categorical Grammar.

This added constraint on the combination of syntactic parts, in that inflectional morphemes only combine with other sub-word units, is captured through the resultant syntactic categories in a straightforward manner. While the greater number of syntactic categories arguably leads to a reduction in syntactic economy, this is only quantitatively so. The fact that this constraint emerges from the theoretical architecture does away with any need to reserve a special theoretical place for ‘words’ (see section 3.1 for arguments against the concept of the ‘word’ and defences against these arguments). A ‘word’ is thus an emergent concept, a consequence of syntactic selectional restrictions, reflecting the syntactocentric nature of such theories. This amounts to greater qualitative parsimony, and may be seen as a benefit of such theories.

When it comes to non-concatenative morphology, however, many problems alluded to in section 1.2 manifest very clearly in such a theory. Focussing on the example of *mouse* and *mice*, we must isolate constituent parts that can be listed with their meanings and syntactic categories, so that they can thus be appropriately composed. The first option would be to store the singular and derive the plural compositionally. This forces an order of precedence on morphosyntactic features of the same inflectional category, which is by no means necessary. Singulars may appear simpler, but this leads to no *a priori* reason why their meanings should be stored and plural meanings derived. Indeed, as in the case of subtractive plurals from French presented in section 1.2.1, the plural form appears at least phonologically simpler.

There are reasons to presume greater simplicity of one morphosyntactic property ahead of another in the same inflectional category (see Farkas & de Swart, 2010, and the discussion of number in section 4.1), but it is typical in any case for such theories to abstract a root and an individual inflectional morpheme for at least the more marked morphosyntactic property, if not for each. The imposition of an order of precedence on morphosyntactic properties, as well as inflectional categories more generally, may be motivated in certain other cases (see Harley & Ritter, 2002), but this is not pressing in the present expository discussion of simple two-way number distinctions.

In the present discussion, we will use the following interpretations for singulars and plurals. We take it that nouns in general denote sets of (atomic)

entities, and, following Link's (1983) seminal analysis, a (complete free join semi-)lattice is generated from the relevant sets of atoms, thereby including in the domain the non-atomic elements, or plural individuals. And let us take it that singular nouns denote the atoms, and plurals, the non-atoms. This is massively simplified, as we would like to include atoms in the plurals denotations too for 'inclusive' readings such as in "did you see mice?", where a singular mouse could be one such reference of the plural form. Section 4.1 treats the issues concerning the different possible plural denotations comprehensively, but for now, let us deal exclusively with 'exclusive' readings for simplicity's sake.

Let us therefore explore a theory where a lexical root is abstracted from both *mouse* and *mice* that is composed with a singular and plural inflectional morpheme in the relevant case. We thus postulate a root morpheme, made of discontinuous consonants, *m-s*, that is thus listed in the lexicon with the associated interpretation as the entire lattice of sets of mice. The respective intervening diphthongs may also be listed in the lexicon with their singular and plural meanings, namely functions that takes as input a lattice and returns the subset of atoms or non-atoms respectively.

This now means that every number exponent with a different phonological form has its own listing. There are thus several different possible morphemes, all with the same meaning. A singular morpheme may be silent as is most common in English, represented by a 'zero-morpheme' (one that is phonologically null); it may be represented as a medial infixal vowel, such as for *m-ou-se* and *l-ou-se*, and also *m-a-n*; or something more unwieldy may be required, such as in the suppletive *person*. There are analogous requirements for the plural morphemes.

Returning to the examples from section 1.2.1, the notion of silent, but nonetheless meaningful lexical entries is perhaps of concern. Take the case of *sheep-sheep*. How are the singular and plural meanings distinguished by the process of interpretation despite them being marked by identical phonologically empty morphemes? Conversely, in the case of the suppletive *person-people*, perhaps we take just the shared *p-* to be the root that refers to the lattice of sets of people. For a case such as "go-went", where there is no phonological relation whatsoever between the two forms, then regardless of their precise interpretation, some silent syntactic unit would thus be required to bear the common lexical meaning and thereby feature as an element in the compositional interpretation. What to do in the case of fully reduplicative plurals such as in Malay? Which of the two identical reduplicands has plural meaning associated to it? Compositional transparency is at stake when the phenogrammatrics are factored in if so many phonologically null or indeterminate units provide the inputs on which the grammar exacts an interpretation.

Many morphemes requires specific, new syntactic operations to give the appropriate phenogrammatical manifestation when they come to be part of a larger expression. Some may simply require suffixation, others infixation, while in other languages, there may be combinations of, such as in the Danish *hånd* and *hænder* meaning "hand" and "hands", as in table 2 - never mind the case of root-and-pattern morphology of Arabic, as was mentioned in section 1.2.1. If the syntactic rule effecting the combination of these plural morphemes with

the relevant roots is not to be disassociated from the phenogrammatcs entirely —equivalent to the uniqueness of indexing of the rules with operations being maintained —then the amount of phonological representation required to deal with the range of non-concatenative phenomena is conceivably extremely cumbersome, and thus drastically increases the complexity of the resultant syntactic operations.

Another interesting and difficult phenomenon is that of ‘deponent’ verbs, which are passive in form but active in meaning. An example is the Latin *obliviscor*, meaning ‘I forget’. The suffix *-or* is typically marks the first person singular present passive, as can be seen in *am-or* meaning ‘I am loved’. Deponents may be considered an instance of ‘defective’ morphology, where morphological markings are somehow misleading. Moreover, in the case of Latin deponent verbs outlined here, there are also related ‘morphological gaps’, where words that appear to be derivable through a process of composition do not in fact exist in the language - there is no form related to *obliviscor* with passive meaning.

If *-or* is to carry the (first person singular present) passive information across by far the majority of verbs, why does it have the opposite, active meaning for *obliviscor*, as well as for other deponent verbs? Is it that there are in fact two stored morphemes, identical in form, but directly contrastive in meaning (there are only two voices in Latin: active and passive)? This solution appears highly counter-intuitive, and may imply difficulties for language learning and processing, especially given that deponents constitute a non-negligible class of verbs. There is nothing to say which morpheme with its specific voice meaning ought to be selected, other than the verb it attaches to, and this could arguably force some non-locality in the process of interpretation. Moreover, how could such an analysis explain why there is no form expressing passive meaning of these deponent verbs? These meanings surely exist semantically, but are never linguistically manifested such that they could be arrived at through an appropriate process of compositional interpretation. There is conceivably no explanation of this phenomenon that can be derived solely from the internal parts of such words.

A final point to mention when it comes to inflection is that the interpretations of highly ‘fusional’ morphemes, which encode strictly more than one morphosyntactic property into a single morphological unit, are difficult to attain when they are viewed as an atomic syntactic part. This is partly because their morphology suggests they should not be syntactically decomposed any further, and partly because the interpretation of the greater expressions they may feature in appears to be more stable. In Kalaallisut, also known as Greenlandic, a single verbal suffix serves to give both subject and object agreement. For example, *-vaa* is the inflection for a third person singular subject and third person plural object; *asavaa* means ‘he/she/it loves them’. Importantly, pronouns are dropped widely in Greenlandic, so the suffix is thus the only means of determining the referents for the verbal arguments.

In morphologically simple languages, when many-place verbal predicates take their arguments, this is typically done in consort with a syntactic analysis

that determines which argument is taken first and which is taken second, and so on, given the curried argument places of the verbal predicate. In the case of Greenlandic, the syntactic analysis that would have to be postulated is not linguistically present, as the fusional morphemes encode both simultaneously, and cannot thus be decomposed any further (at least not phenogrammatically). A syntactic analysis providing the supposed ordering of the application of the verbal predicate to its arguments would therefore require a potentially undesirable leap in abstraction, dissociating the phenogrammatics from the compositional interpretation. Perhaps some compositional analysis could be given to resolve the situation, where the fusional suffixes are a syntactic atom with a basic meaning, but this seems as though it would be difficult, and is not explored any further here. The point remains that the situation for such a language appears wildly different from that of morphologically simpler languages that encode the same information across several compositional elements, as is reflected in the fusion exhibited by the inflectional morphemes.

This speaks to a broader problem, exemplified by fusional morphology, which is that the order of composition of several morphosyntactic properties that are represented by a single morpheme may be not at all obvious. Other such examples may derive from the postulation of one or more zero-morphemes that feature as part of the same word-form. A somewhat contrived example from English, taken from Stump (2018), is that of the English past tense of the verb *cut*. The absence of any overt past tense marking, as well as any number marking, such as in the past tense forms *was-were*, could suggest that both the number and tense inflections are phonologically empty. What, then, is the order of composition of these morphosyntactic properties with their resultant meanings? The same question may be asked for *was* and *were* given their fusional appearance, as they do indeed encode number distinctions overtly (at least singular vs. non-singular, since “were” is used for the second person singular as well). This may seem a somewhat contrived example, but similar phenomena occur in other languages more concretely and demand a response in terms of an incremental theory of inflection.

Abstract syntactic analyses could be proposed for the above cases, for which there may be little direct linguistic evidence, in order to keep the interpretation of such morphologically complex languages in line with that of morphologically simpler languages. This avoids a proper account of the interaction between phenogrammatical variation and compositional interpretation, however, and employs a methodology that involves a certain theoretical and typological preference towards languages with greater ‘iconicity’, or one-to-one form-to-meaning correspondences. This preference may be borne out in predictions relating to learnability and processing. The one-to-many (fusional) and many-to-one (fissional³) form-to-meaning (or -to-morphosyntactic property) relations that underline a great deal of morphological variation are heavily implicated in questions of syntactic economy and compositional transparency, and therefore

³Morphological fission is perhaps most clearly exemplified in the root-and-pattern phenomena as in Arabic, but in general concerns cases where one morphosyntactic property may be expressed by two or more discontinuous morphological units.

demand the inclusion of phenogrammatical phenomena into the discussion in a way that privileges no particular kind of morphology.

We may corner the problems that arise given the morpheme-based character of such a theory in each of the morphemic and incremental dimensions that correspond to the morphotactic versus morphosyntactic distinction. As regards the former, it is in many cases unclear which sub-word units to isolate phenogrammatically and store with their meanings. To decompose morphologically complex words into proper parts requires additional assumptions about what the lexicon may consist in, and increased complexity of the syntactic operations that recompose morphemes into larger constructions. This entails diminished syntactic economy where parts are not concatenated into larger expressions, not to mention greater opaqueness in terms of what kinds of linguistic objects may provide the input to the interpretation, such as for phonologically empty and discontinuous units.

As for the morphosyntactically morpheme-based incremental component, it appears that the converse association of morphosyntactic properties, in terms of their meanings, to particular morphemes themselves may seem suspect if these morphemes are not monolithic phonological strings. Even when they are, however, there may be clashes of meanings that cannot be resolved locally, as in the case of Latin deponent verbs. Moreover, the fusion of semantic information into one morpheme, as in Greenlandic, challenges tectogrammatical assumptions concerning the manner in which verbs take their arguments in a compositionally transparent manner. Such problems persist as the atomic parts with their associated morphosyntactic properties and corresponding meanings are composed into larger expressions, since the order in which morphemes are added may not always be evident from word-internal structure. This in turn adds another layer of compositional opacity.

While the discussion thus far has centred around inflectional morphology, for derivational morphology, the picture may become murkier still. It is unclear in many cases what the interpretation of derivational morphemes may be, as they may be highly polysemous in an unsystematic way, or convey meanings that are difficult to capture model-theoretically. One example of polysemy comes from the set of related derivational forms *import*, *export*, *import*, *deport*, *purport*, *comport*, and potentially others. There may be some semantic relation between the meaning of the derivational base, or root (*-port*), potentially meaning something like “carry”, and the full derivational forms with their prefixes, but it is rather (compositionally) opaque in some of these cases. Moreover, the exact interpretation of some derivational affixes may be difficult to derive, such as in the English suffixes *-ness* and *-tion*, for diminutive and augmentative formations in other languages.

2.2.2 Morphemic-Realisational

A Morphemic-Realisational theory ideally associates morphosyntactic properties directly to fully realised word-forms, while maintaining that sub-word units, or morphemes, are stored in the lexicon. Given the above conception of the

lexicon as the store of atomic parts with their basic meanings, the at face value inconsistent character of such a theory necessitates some abstraction for sub-word units. One clear example of this is in Distributed Morphology (DM - Halle & Marantz, 1993), which serves to guide and inform the following discussion. The question of what is stored in the lexicon revolves around the notion of ‘listedness’, where any special status for words in particular as being ideal signs, with some special phonology, syntactic status and meaning is denied (again, see section 3.1); rather, any (minimal) part that exhibits such an ideal relationship may be listed (in the lexicon).

In such a theory, syntactic structures are built using purely abstract morphemes, with a split assumed between those representing lexical roots and those representing morphosyntactic properties. This could again simply be a result of the syntactic categories of the morphemes in a categorial grammar, although in the typical DM framework, this distinction is simply stipulated, because it allows for the realisation status of the theory to be upheld in spite of its syntactocentric nature. Having derived a syntactic structure, it is then up to the listings in the lexicon associating (bundles of) features with phonological representations - ideally strings - to provide the phenogrammatical realisations that are appropriate at the given node, concatenating them (or performing some other operation) accordingly. The syntactic structure alone, apart from the phenogrammatical manifestation, is taken to be the input to the semantic interpretation.⁴

‘Vocabulary Insertion’ is the process that governs when phonological forms are inserted into the tree. This was originally proposed to take place exclusively at terminal nodes. This would of course render the theory rather incremental, and requires extra ‘readjustment rules’ to deal with non-concatenative morphology (to turn, say, *mouses* into *mice* at the parent node of the root and inflection), as well as zero-morphemes to represent for phonologically null exponents. Instead, a more faithful realisation approach taken in later guises of DM, that incorporates in ‘Late Insertion’, where the appropriate listed items may be inserted at non-terminal nodes, such as in Radkevich (2010).

For example, *mice* may be listed with the bundle of abstract morphemes that comprises the lexical root and plural morpheme, and then inserted at the superordinate node to the lexical root and plural morpheme, while *dog-* and *-s* may be stored individually with the respective abstract root and inflectional morphemes, and then concatenated in accordance with the structure that the abstract morphemes read out as them feature in. Such a theory thereby copes with non-concatenative morphology better than its incremental counterpart, as well as fusional morphology, so we will only consider DM-variants with late insertion. The listed items that are inserted at non-terminal nodes are typically associated only with unstructured bundles of features, as opposed to the whole syntactic structures they come to express phenogrammatically.

There are other processes and rules to deal with form-to-meaning mismatches, however, such as syncretism, where one form may have several mean-

⁴DM is not directly compositional, but this is besides the point for now, as we could imagine that morphemes are in fact interpreted in tandem with the construction of larger expressions from them, without the need for an intervening LF.

ings (*are* in English covers the second person singular and all plural forms). Here, *are* would be listed only as the present copula, without any number or person feature associated, while *am* and *is* may be listed with bundles that include both singular morpheme and the respective person morphemes. Importantly, all listed items compete for insertion, but only those with the best fit —those which include as many of the morphemes present in the derivation as possible —are ultimately inserted.

While it could appear for words with simple concatenative structure that their parts are associated with morphosyntactic properties more directly than the full word-form, this is not the case. Because an entire abstract syntactic structure is taken to be built first before being sent to the phonological read-out process (and also sent separately to the interpretation process), we can claim the whole word-form is associated with morphosyntactic properties, ahead of its parts. Namely, the most superordinate node in that abstract structure, which gathers all the relevant morphosyntactic properties beneath it, is what provides the input to the read-out (and interpretation). Then, dependent on that structure, the word-form is read out. The theoretical position is maintained that the whole word-form is only associable with the entire structure, and its realisational status is thereby maintained.

The first thing to note is that the conception of the lexicon adopted here that provides the basic operands of the grammar at all levels (phenogrammatical, tectogrammatical, and semantic), is no longer available. The tectogrammatical works on abstract morphemes alone, and it is they that bear meaning separately from the actually pronounced, concrete linguistic elements. During a derivation, a process of competition yields the actual linguistic manifestations, where the basic phenogrammatical stuff is in fact only associated to unstructured bundles of abstract morphemes, and not to their meanings, or the structures that determine them as a result. This means that the grammar no longer works on signs as representations where form, syntax and meaning coincide. Instead, there is one interface between the phenogrammatical and a reduced syntax, namely bundles of abstract morphemes where the tectogrammatical structure is collapsed, while there is another, separate interface between the syntactic structure and the semantics.

This is rather opaque, but a necessary result of the overarching theoretical design. Any transparency between the two interfaces is mostly accidental, being systematic conceivably only for isolating languages with no morphology whatsoever, or for agglutinating languages whose morphology consists only of one-to-one form-to-meaning relations, namely for iconically ideal languages. Such a typological privilege is, again, theoretically undesirable, as it arguably suggests that languages that do not conform to this ideal are in some way aberrant. The widely attested aberrance demands an explanation diachronically, as such languages would surely be disfavoured given plausible constraints on learnability, and perhaps also synchronically, as there may be concomitant predictions when it comes to processing and production. Besides this, the claim that it is not the pronounced elements of language that are ultimately interpreted, but rather that they serve only to encode some abstract intermediate syntactic rep-

resentation instead provides the input to an interpretation, is philosophically suspect.

The likely trade-off between compositional transparency and syntactic economy would suggest that such a compositionally opaque theory ought in fact to be highly syntactically economical. Indeed, when it comes to the phenogrammatics, this appears to be the case, as ‘Late Insertion’ allows for non-concatenative morphology to be simply a consequence of listedness, and not the result of any unwieldy syntactic operation. Despite this, however, there is no equivalent gain when it comes to tectogrammatical syntactic economy, as the combinatory rules need to be able to operate on the most minimal abstract morphemes that can be postulated. An unexplored alternative could store trees with the listed items inserted at non-terminal nodes, rescuing some compositional transparency, but this entails other negative consequences for the complexity of the grammar, both in terms of the richness of representations in general, and the syntactic rules that would operate on trees and not categories.

In sum, DM solves some of the problems encountered in a Morphemic-Incremental theory, particularly for variants that incorporate late insertion. In such a way, non-concatenative morphology can be dealt with simply, where it does not need to be decomposed phenogrammatically. A new problem is developed, however, through the separation of two interfaces, one between the phenogrammatics and syntax, and the other between the latter and the semantics.

This increased opacity could be resolved somewhat if entire trees are stored with the vocabulary items inserted at non-terminal nodes, but the option taken in DM is to store them as unstructured bundles of features. This is to feed the process of competition where the vocabulary item stored with the greatest number of features implicated at the relevant node is inserted, resulting in the most appropriate realisation - it is not immediately obvious how entire structures could feed some analogous process of competition in any straightforward manner. The tectogrammatical structure is therefore being relevant to the second interface with the semantics. Moreover, while that possible solution may increase compositional transparency, the lexicon as the store of the most basic inputs to the grammar becomes much more complex than is typically tolerated, if whole trees upon which an interpretation depends can be stored with basic phenogrammatical units.

2.2.3 Inferential-Incremental

In consideration of the kind of morphological rules of inference that an incremental theory may reasonably employ, it appears that rules of exponence are demanded instead of paradigmatic rules. It is not clear how paradigmatic relations amongst full-word forms associate morphosyntactic properties incrementally, with the notion of their coincidence being crucial to defining the cells in the first instance. Perhaps one could suggest that the rows and columns, such as in Table 2, represent some abstract kind of morpheme, as they indeed represent morphosyntactic properties. Given that the cells at which the full word-forms

are found are simply where these properties coincide (again abstractly with respect to the lexeme), there would be no conceivable reason to impose some order on the association of these morphosyntactic properties to any steps along the way to deriving the full word form. Rules of exponence, by contrast, associate morphosyntactic properties step-wise to increasingly complex stems until a whole word-form is arrived at with its resultant meaning.

As mentioned above, inferential theories include a morphological component as part of the grammar, which is distinct from the (sentential) syntax. Here, this morphological component is embodied by rules of exponence, applications of which derive more complex stems and ultimately whole word-forms in a linear step-wise fashion. The outline provided here roughly follows the framework of Articulated Morphology (Steele, 1995).

We may conceive of rules of exponence as a morphology-specific tectogrammatical analogue, where they associate to the stems that serve as their input some additional (often single-element) set of morphosyntactic properties with the corresponding meaning(s) per instance of application, while simultaneously specifying the forms of the input and output phenogrammatically. Importantly, this association is achieved through the application of the rule and not by way of some individual linguistic element bearing the morphosyntactic property.

Only when all the rules of exponence necessary to yield a full word-form have been applied is an input to a distinct sentential syntax provided. This in a sense mirrors the distinct combinatorial syntactic profiles of sub-word units and full word-forms in morphemic theories, except that this distinction is necessarily built into the architecture of the grammar here, reflecting the morphocentric character of such a theory, as opposed to being epiphenomenal to tectogrammatical selectional restrictions.

Now to provide a loose implementation for the derivations of *mouse* and *mice*, with the same target interpretations as before. First, a choice must be made: either, the singular is the basic meaning, and the application of the plural rule of exponence manifests the plural meaning; or, there is some abstract, basic meaning, and there is a singular rule and a plural rule. For analogous reasons to those in section 2.2.1, the former option is disregarded in favour of the latter.

We thus assume there is some basic stem, *m-s*, with the medial vowel underdetermined, which may be taken to denote the entire lattice of sets of mice. Again, this meaning is not *prima facie* part of the language, given that the form taken to express this meaning is never linguistically manifested itself, yet it would have to be stored to provide an input to the rule. This is perhaps undesirable, as the theory is thus rendered less concrete. In any case, the singular and plural rules of exponence applied to this root yield the forms *mouse* and *mice*, where the rules themselves are stored that specify the appropriate medial vowels phenogrammatically, and semantically as the function that takes the whole lattice and returns the atoms and non-atoms for the singular and plural respectively.

It is now possible to discuss how the morphological component of the grammar fits in with the rest of the grammar in relation to our criteria. The first problem to note derives from the linear nature of rules of exponence. While

there is a basic meaning, for example the root, stored with its meaning, the interpretive procedure yielding the more refined grammatical meanings of the inflected forms is associated entirely with the applications of the relevant rules, and not with any compositional parts or elements *per se*. The plural inferring the word-form *mice* from the root has the stored meaning of a function that takes as input the lattice of all sets of mice, and returns just the non-atomic elements. When the rule is applied morphologically to realise the plural exponent, the associated functional application is also effected. By contrast, in morphemic theories, the same function was instead stored with the actually manifested plural morpheme, while its application occurred in conjunction with the relevant combinatory syntactic rule that combined the individually meaningful sub-word units.

It is not directly important what the semantic operation associated with the rule is, but rather that morphological-tectogrammatical rules of exponence now build in significantly more semantic information in interpreting words than do the bare, abstract combinatory rules of Categorical Grammar that are typically interpreted simply as functional application. This may lead to predictions concerning morphological typology, where morphologically rich languages have an extra set of rules of an entirely different and more complex kind, both structurally and in terms of their semantic content. Whether this is borne out or not in reality is not probed here, but it is a strong prediction that may be difficult to test and generally undesirable. While it is perhaps desirable not to have concrete linguistic elements associated with morphosyntactic properties alone, due to the phenogrammatical issues that arise, to lump them instead into a series of rules that modify stems carrying lexical information does little better, as attributing such grammatical content to an abstract theoretical device is conceivably yet more stipulative. A wedge is thus forced between the word-internal, morphological tectogrammatcs and the sentential counterpart. The latter maintains compositional transparency, as is typical, where the nodes at which branches meet due to their well-formedness having been proved in a system of Categorical Grammar simply effect functional application in the semantics. This transparency is because the abstract rules of Categorical Grammar and functional applications as denoted by the language of lambda calculus work in consort, reflecting the Curry-Howard correspondence (Dowty, *ibid.*). Such a process is rendered significantly more opaque word-internally, however, with rules of exponence themselves providing the semantic information of, for example, the morphosyntactic property PLURAL, which means they have to be stored individually, alongside the lexical items, to allow for the computation of their effects when composed with roots and stems. Again, it may be undesirable to assume the storage of such contentful rules, since they are not concrete linguistic elements, but rather theoretical devices. Besides this, to propose two kinds of tectogrammatcs is of course unparsimonious, although it is demanded by the ‘morphocentric’ character of inferential theories, at least when coupled with an incremental component that requires rules of exponence.

The picture becomes arguably more opaque upon consideration of the phenogrammatcs, as a rule of exponence interpreted to yield the meanings

of plural forms, for example, may have several different manifestations. This problem affects all incremental theories, in fact. Sometimes it adds only a suffix, sometimes both the vowel change and a suffix, as in the Danish *hånd-hænder*, and sometimes there is no change at all, as in *sheep*, amongst other possibilities, such as suppletive plurals.

It may be claimed that one rule of exponence, such as for pluralisation is manifested through several different phonological operations, but this is undesirable as the abandonment of uniqueness has the result of disassociating the phenogramatics from the rest of the grammar. Moreover, the morphocentricity of rules of exponence as relating to the morphophonological exponents of morphosyntactic properties means that they should incorporate the phenogramatics to the greatest extent possible.

The preferable alternative would then be to claim that there are many pluralisation rules of exponence, one for each possible phenogrammatical manifestation of plural number. This has the unwanted result, however, that an important generalisation is lost in the interface with the semantics, where all of these rules have the same effect in the process of interpretation. In fact, what is apparently the proposal of Articulated Morphology in its original formulation is to give one rule for every plural form in the language - one that forms dogs from the relevant root, another that forms *cats*, and so on. This relinquishes another opportunity for generalisation in terms of the phonology. This is because, in its original formulation, it works in terms of pairs of phonological material and sets of morphosyntactic feature, so to avoid abstraction in terms of the former, it is necessary to include each specific phonological input and output for every rule.

Another point of interest is the resultant place held for morphosyntactic properties; while in morphemic theories, they are represented by the morphemes either individually or in clusters for instances of fusion, here they are represented by the sets of rules that associate the same property (or set of properties for fusion). The theoretical emphasis is therefore weighted towards the phenogrammatical side of morphosyntactically related forms, as they are a property of the relevant set of rules of exponence, namely the effect on the meaning of their input. This is unlike morphemic theories, where morphosyntactic properties are reified through the morphemes representing them. Seeing as morphosyntactic properties are more so theoretical entities, this possibility for them to be emergent in Inferential-Incremental theories is arguably a desirable result.

While part of the goal of inferential theories of morphology is to better account for the way in which morphological structure differs from sentential structure, combining this with an incremental process of interpretation of the inflectional exponents leads to a familiar set of issues. Rules of exponence limit both syntactic economy and compositional transparency through their linear nature, demanding that the (morphological-)tectogrammatical rules involve much more concrete semantic import than would typically be tolerated at the sentential level.

The multiplicity of their phenogrammatical manifestations is another potentially undesirable result of the morphocentric character of such theories, although it does allow for a less reified theoretical status for morphosyntactic

properties than in morphemic theories. The trade-off for this is to store rules in the lexicon, which is arguably more stipulative. Something else to mention briefly is that it also remains unclear how such a theory would deal with derivational morphology, as the semantic operations associated to the rules of exponence that infer new derivational forms would have to be even more complex and contentful than for inflections, if they could be captured model-theoretically in the first place.

Returning to the point of the morphotactic and morphosyntactic inconsistency of such a theory, rules of exponence may at first appear to offer something of a happy medium between the morphotactically non-morpheme-based inferential component and the morphosyntactically morpheme-based incremental aspect. In actuality, however, they suffer from many of the same critiques as can be targeted at morphemic theories, except that the source of the problems no longer lies with the ways in which basically meaningful minimal parts are put together to yield words and their meanings, but rather with how the rules themselves change their inputs and the inputs' meanings in a linear manner. While they avoid certain of the issues encountered in morphemic theories, they in turn cause novel problems by alienating morphology from the broader grammatical system through the postulation of a new kind of theoretical entity that is to be stored along with lexical roots, namely the rules of exponence.

2.2.4 Inferential-Realisational

All of the above kinds of theory have in some way imported certain notions from a compositional theory of interpretation of sentence-level structure to word-internal structure. This is because such theories take a broadly 'constructive' view of morphology (Blevins, 2006), where the attempt is made first to isolate minimal meaningful elements somehow before then deriving the whole word-forms in which they occur. When this is coupled with a compositional theory of interpretation, problems arise often due to a clash between which minimal parts are isolated, what meanings are to be associated to them, and how the parts and meanings may be combined in a transparent manner to yield the desired full word-forms with the meanings they are supposed to bear from the beginning of the analysis.

The alternative to this is an 'abstractive' view of morphology (Blevins, *ibid.*), which proposes that word-internal structure is only relevant insofar as it reveals the external relations of full word-forms themselves, in their (*gestalt*) entirety to one another. Such relations may be taken here to be constituted by both the shape of the word-forms and their associated meanings. Indeed, such a view denies a role for compositional interpretation within words. For the time being, the notions of a theory being Inferential-Realisational, and being abstractive may be conflated, although the subtle distinctions are to be treated in section 3.2, where the theory of Word & Paradigm Morphology (Blevins, *ibid.*) is outlined in greater detail, which is followed in section 3.3 by a discussion of its potential philosophical relevance and in chapter 4 by a formal semantic implementation for a range of morphological phenomena. The remainder of this section further

outlines the basic premises of such a theory and evaluates it with respect to the criteria provided in section 1.1.

To provide a picture of an abstractive, word-based theory, (sets of) morphosyntactic properties are related to whole word-forms, and the inferential rules used to postulate new word-forms are captured in terms of the relations across and within inflectional paradigms (*e.g.* tables 1 and 2), namely through the analogical paradigmatic rules discussed in section 2.1. Unlike for rules of exponence, which are stored in the lexicon, the paradigmatic rules emerge from the organisation of the lexicon. There is thus a relation between *mouse* and *louse*, because of their shared plural forms *mice* and *lice*. They form a very small inflectional family, of size 2, unlike the enormous inflectional family whose members are the paradigmatically related sets of forms *house-houses*, *grouse-grouses*, *dog-dogs*, *cat-cats*, and so on.

These patterns of related forms can be exploited to derive new ones, where they are not already given, but this process is not considered to be part of the language itself; it is instead taken to be part of the acquisition process of the morphological system, which is idealised to include all forms stored with their meanings in advance. The greater the size of the inflectional family, the more likely the pattern exhibited by its word-forms relating to a given morphosyntactic property is to be extended by analogy to new forms when associating that property to the relevant lexeme. See section 3.2 and subsections for more in-depth account of these characteristics of the theory. This is the case both synchronically, through overgeneralisation in the acquisition process, and diachronically, through the well-attested spread of often regular or frequent morphological patterns by analogy (Fertig, 2013).

A morphosyntactic property may be exhibited in a number of different ways phenogrammatically, across different inflectional families, but no particular theoretical precedence is given to any one of these manifestations, since the property is no longer compositionally relevant apart from the word-form realising it. The plural forms *mice* and *dogs* are phonologically monolithic strings and syntactic atoms, listed individually in the lexicon as the contrasting neighbours of their singular forms, each bearing their associated denotations which provide the basic meanings to be composed in the larger sentences in which they may occur.

The number of items stored in the lexicon greatly increases, therefore, with one listing for each paradigm cell. Unlike for rules of exponence in the previous section, however, this multiplicity is only quantitative, with the postulation of no new kinds of theoretical entities or devices required, since the meanings of the word-forms are not derived in any way, but are instead stipulated from the off. This thereby maintains qualitative parsimony. The greater size and complexity of the lexicon may lead to other concerns, and these are treated in section 3.2 and its subsections.

A lexeme is the abstract theoretical entity around which paradigms are organised. Importantly, however, they are not necessarily supposed to have a concrete reality in the language itself, unlike roots and rules of exponence. Perhaps we attribute them some role in, say, the process of learning, or the production

and processing of novel forms (to be discussed later), but the crucial point is that they are never phenogrammatically manifested, and play no part in the process of interpretation. They are therefore also not directly tectogrammatically relevant, if syntactic categories are assumed to be borne by word-forms. Any potential meaning that could be attributed to a lexeme, for example the union of the denotations of all forms in its paradigm, is thus not compositionally relevant, but rather only lexically so, as it provides the common lexical meaning of all inflectional forms in a given paradigm, around which the specific, refined denotations are organised that actually play a part in an interpretation.

Somewhat orthogonal to the role of a lexeme is that played by inflectional categories in determining the scope of (grammatical) meanings that the word-forms of any paradigm may realise. The problem of morphological gaps exhibited by defective classes, such as the Latin deponent verbs referenced in section 2.2.1, can thus be addressed simply as a given lexeme failing to have an actual form or set of forms associated with a given subset of the theoretically possible paradigm cells. While these meanings may still be available semantically, given that there is no compositionality in operation within words, there is no conceivable process by which these word forms ought to be able to be derived, as was the crux of the problem articulated above. There is arguably a missing set of related forms for those defective paradigms, but this becomes a property of that section of the morphological system, and therefore systematic for that class (see the discussions of idiosyncratic stative meanings in Chichewa in section 3.1.2, or of ‘morphomic’ patterns in section 3.2.1 for examples of other phenomena that are treated analogously). Such a treatment is therefore more descriptively faithful to such phenomena, even if it does not provide an explanation *per se* (should there even be a synchronic explanation).

Importantly, while lexemes govern intra-paradigmatic relations, inflectional categories determine inter-paradigmatic relations, where different combinations of morphosyntactic properties define specific cells. The general theme of this complementarity, which is more directly treated in section 4, is that lexemes provide the (structures of) elements of the domain (for example lattices of the appropriate kind for nouns), which inflectional categories may systematically refine, in order to establish the appropriate interpretations of morphosyntactically and lexically related word-forms.

Many of the problems encountered in the above constructive theories simply disappear when morphologically complex word-forms are not decomposed. Firstly, there is no problem with non-concatenative structure, as it is taken to be monolithic, at least with respect to a compositional interpretation. This allows syntactic economy to remain exactly as would typically be required for sentential syntax - the phenogrammatical complexity of word-structure is no longer relevant to syntactic operations or the grammar in general. Moreover, at the tectogrammatical level, morphologically rich languages require no new syntactic categories to be derived for sub-word units, and no extra structure to be kept track of, as they are syntactically atomic. In fact, the converse arises, namely that morphologically impoverished languages become more syntactically complex, as they may often require more atomic elements to express the same

meanings.

Moving on to compositional transparency, the difficulty of not only settling on which sub-word units to associate meanings to, but also on what meanings should be associated to these parts, where one-to-one form-to-meaning relations are not exhibited, or forms are misleading, is again obviated, because no compositional import is attributed to word-internal structure. Solutions to some of these problems were proposed in the shape of abstract, phenogramatically unmanifested syntactic parts in section 2.2.2, or through an unparsimonious, distinct set of morphological-tectogrammatical rules which were extremely contentful, in section 2.2.3. No such attempt at a solution that may in turn lead to other problems is required, as the original problem does not arise when complex word-forms are not decomposed.

There are questions that arise in connection if no (compositional) semantic import is attributed to complex word-internal structure. How are the meanings of complex word-forms learnt? What determines their semantic contribution to the larger expressions that they may find themselves in? How is typological variation to be conceived of? Indeed, many generalisations appear to be missed if the parts of complex word-forms are not isolated in some way. For example, the suffix *-s* on nouns really does appear to imply plural meaning of the nouns it attaches to. The holistic conception of the lexicon that is associated with theories as this must be clearly articulated, if these questions are to be addressed to a satisfactory degree. This is the task undertaken in the subsequent chapters.

Another point concerns structural semantic complexity. Derivational morphology, with its semantically opaque and complex operations on meanings holds a secondary place in relation to inflectional morphology with respect to the broader morphological system. As was discussed in section 1.2.2, unlike inflectional forms, which are related within paradigms of a given lexeme, derivations are taken to create new lexemes, with their own paradigms. Given that derivational morphology can change a lexical item's syntactic category, where different syntactic categories have their own place in the broader inflectional system, the assertion that derivations result in new lexemes is sound.

Time and space do not permit an investigation of the nature of the connections amongst lexemes provided by derivational forms and how these may result in their specific interpretations and syntactic natures, for example in terms of argument structure and aspectual qualities (see Borer, 2012). It is nonetheless hoped that the conception of derivations as determining higher-order relations, across lexemes (or sets of word-forms, or paradigms - see footnote 1 of this chapter), than do inflectional categories, across cells in paradigms, suggests on the one hand a lesser degree of transparency in their semantics, and on the other, that there is still scope for some explanation in future research of the data concerning their syntactic relations (see section 3.1.3 for one potentially hopeful avenue in a word-based theory). Such a distinction is not maintained in the above constructive theories, where the roles of inflectional and derivational morphology are conflated in a compositional interpretation, and distinguished only through the differences in the meanings of the derivational and inflectional exponents, and not through any systematic contrast.

Lastly, it should be noted that an important typological prediction is made here. Morphologically rich languages, being taken here often to be able to express the same meanings with fewer compositional elements (namely whole word-forms) in comparison to their morphologically impoverished counterparts, are thus deemed to involve a conceivably simpler compositional process of interpretation. Further, the interpretations of the syntactic atoms in more synthetic languages may in turn involve functions of smaller arities, with less derived types and fewer lambda-abstractions, which all suggest lesser structural semantic complexity of stored lexical items. Moreover, the syntactic categories will be less derived, with fewer steps of composition required to arrive at a sentence's ultimate meaning (see section 4.2.3 for an example of how this might be exemplified by languages with pro-drop).

For overall cross-linguistic parity of complexity to be maintained —this being a fundamental presumption underlying empirical research into natural language semantics and linguistics in general, due to largely comparable rates of acquisition and processing of different languages —a commensurate source of complexity is demanded. It is proposed that that source is to be found in the network of meanings organising the lexicon. This argument is further outlined in section 3.2. One crucial contrast to constructive theories is that, in Inferential-Realisational theories, the meanings encoded in the lexicon themselves and how they are related by way of their encodings and the structure of the lexicon are seen to provide the source of typological morphological variation from the abstractive perspective. In constructive theories, the variation to be found in the diversity of phenogrammatical manifestations alone, which are ultimately beholden to a broadly analogous compositional process of interpretation cross-linguistically.

Table 2.3 summarises some of the main strengths and weaknesses of each kind of theory encountered in the preceding evaluation:

	Strengths	Weaknesses
M-I	Parsimony: one syntax for both word- & sentence-structure; concreteness maintained	Encounters serious difficulties for non-concatenative & fusional morphology
M-R	Non-concatenative & fusional structure treated better; a single syntax for words & sentences maintained	Opaque with two systematically distinct grammatical interfaces; the syntax is necessarily abstract
I-I	Explains syntagmatic morphological structure without recourse to syntax	Requires the storage of potentially highly contentful rules distinctly from basic lexical stems
I-R	Any and all problems regarding non-concatenative and fusional morphology are avoided; transparency & economy thereby maintained	Available generalisations concerning word-internal structure missed entirely; how can the morphological system be learnt?

Table 2.3: Summary of Evaluation

Chapter 3

Word-Based Morphology

A definite concept of ‘a word’ is of course essential to the word-based view of morphology taken here, to guide our understanding of how morphology is to interact with the rest of the grammar. A (consistently) word-based theory typically assumes that morphologically complex word-forms are ideal signs, in the sense of the triples articulated in section 1.1, ahead of whatever parts appear to make them up. It serves to reiterate the relevance of such signs to word-based theories before critique of them is treated below.

In the morphotactic dimension, the emphasis is on the primitive status of whole word-forms. They are not decomposable, in part due to the abundance of non-concatenative morphological structure, but also because it is whole word-forms that are related paradigmatically across the morphological system, and not their parts. Phenogrammatically, this amounts to a limitation on phonological representation, in that full word-forms can be listed individually as strings in the lexicon, the store of basic signs. This also speaks to the extended sense of (ii) (Compositional Transparency), where the need for non-concatenative, non-context-free operations can thus be eliminated in turn.

As concerns the morphosyntactic character of such theories, whole words are of primary grammatical significance. At no point is the association of morphosyntactic properties to a word-form taken to be directly dependent on its parts, but rather these parts may serve to discriminate which full forms have which grammatical status. This discrimination may be aided by patterns of relations amongst word-forms across the lexicon. For example, *dogs*, *pigeons*, *cataclysms* and so on are all related to *dog*, *pigeon* and *cataclysm* in the same way, namely that they are respectively plural and singular forms of the same nominal predicate. *Lice* and *mice* are related to *louse* and *mouse* similarly in the morphosyntactic dimension, but differently in the morphotactic dimension.

Tectogrammatically, no new syntactic categories need to be derived, as no direct grammatical relevance is attributed to sub-word units, which is to some degree more syntactically economical. It may be that a morphological derivation leads to the derived form being of a different syntactic category to the grammatical base, but this, again, is a result of the pattern of relations (amongst lex-

emes) in which the specific relation of the derivational base to the derived form sits. For example, all of the following word-forms are nouns: *goodness*, *sadness*, *arbitrariness*, *cantankerousness*. The words (lexemes) from which these forms are derived are all adjectives, and this therefore constitutes a suffixal pattern of derivational-morphological relations between nouns and adjectives. Another such pattern would be *length*, *warmth*, *breadth*, and so on, which constitutes a similar grammatical relation that is differently morphotactically constituted.

Semantically, the meanings that may be attributed to morphosyntactic properties are more stably associated to whole words than their parts. This is seen as a reflection of (i), where there may be great multiplicity in the exponence of a given morphosyntactic property, as well as fusion of several properties into one exponent. The potentially undesirable abstractions needed to account for such cases may simply be avoided if whole words themselves represent the foremost theoretical entity to which such meanings may be associated. The question of whether all morphosyntactic properties are of semantic import is saved for section 4, but there are at least a good many inflectional categories whose meaningfulness is not in doubt (number, tense, aspect, and so on).

3.1 In Defense of “the Word”

The claim that morphological words are in fact ideal signs has been refuted notably by Lieber (1992), and Marantz (1997). The typical conclusion of such arguments is the claim that there is “syntax all the way down”, internal to words as well as externally in sentences, as represented in the syntactocentric, morphemic theories discussed above. While an analysis of such theories is no longer pressing, being dealt with in section 2.2.1 and 2.2.2, a survey of some of the arguments against the concept of the word is undertaken here with respect to the above outline of an abstractive, word-based theory. In each case, it is assessed whether the arguments are entirely relevant to the specific claims made by the kinds of word-based theories of morphology adopted here, and where they are, a counter-argument, or another way of dealing with the same phenomenon is or suggested.

Importantly, these arguments were originally posed to refute Lexicalism (Chomsky, 1970), which makes different claims than does the word-based conception of morphology provided here. Lexicalism is in fact concordant with a constructive perspective on morphocentricity, such as is adopted in Incremental-Inferential theories, where the processes that derive words (typically in the lexicon) are seen as distinct from the rest of the syntactic processes by which sentences are constructed. This differs from the abstractive sense in which morphology is conceived of here, which augments the lexicon to include all possible fully realised morphological word-forms in the first place, as independently and systematically stored items, and denies a place theoretically to any set of processes deriving whole words, be they equivalent to or distinct from those that construct sentences.

The refutations, particularly that of Marantz (1997), are chiefly concerned

with the coinciding notions of words having: firstly, a special sound; secondly, a special meaning, and thirdly; a special structure. What is meant by special here is that there is some set of properties (phonological, semantic, or syntactic) that can only be attributed to one grammatical level, or entity or concept systematically, that being ‘the word’. This three-way distinction mirrors the conception of the grammar employed here, in terms of the signs that are its operands, that were introduced in section 1.1. If there are linguistic elements which fail to be special in one of these senses, even if they appear to be special in another sense, then it is claimed that the domain of the “word” cannot be coherently defined, and is thus of no direct theoretical relevance to the grammar.

As is evident from this line of argument, in its goal to refute Lexicalism, the conclusion offered that normal syntactic processes instead of any special morphological processes are what constructs words. Here, however, both of these suggestions are denied, and in their place, it is the network of relations amongst lexical items, which is organised according to the way in which those items coincide along all three of these dimensions in the broader morphological environment, that is taken to define the lexicon and morphology in turn. As is explained below, this lends a different perspective on the way that the storage may be conceived of with respect to signs, in that there is greater theoretical scope to include putatively non-ideal elements.

3.1.1 Special Sound

The refutation of words having a special phonological, or prosodic status centres around the fact that the prosodic domain of ‘words’ is observed not to align with their syntactic domain in certain cases. Two exemplary phenomena outlined in Lieber (1992) serve to highlight the point clearly. The first concerns the English possessive clitic suffix *-s* may attach to whole phrases, such as *the man I met’s dog*. While *met’s* is a word prosodically, the whole syntactic unit *the man I met* is what is interpreted as the possessor of the dog when composed with the clitic, with the prosodic word being irrelevant to the interpretation of the whole construction.

This is not a direct problem in an abstractive word-based theory, however, as cliticisation is not taken to be a morphological phenomenon, since the associated interpretation of the clitic suffix is not relevant externally throughout the lexicon, in that it does not relate word-forms to one another with respect to any morphosyntactic property. This differs from a genitive, where a close cousin to such a clitic, such as the German genitive singular suffix *-s* (as in *des Mannes* meaning “of the man”) forms a part of individual word-forms alone and not phrases, contrasting systematically with other case and number forms.

Instead, cliticisation is clearly a syntactic phenomenon, so the sentential syntax may work as is necessary to produce the relevant phrase. While *met’s* appears to be a phonological word in a prosodic sense, the evident divisibility suggests that the clitic could be stored individually in the lexicon with its syntactic category and meaning, to be operated on appropriately by the tectogrammatcs. The morphological theory taken here is agnostic with respect to

the storage of an element in the lexicon that is prosodically less than a word, in particular because it is of no broader paradigmatic import. The fact that this clitic appears not to be phonologically ‘special’, in being word-like, in the same way as it is so syntactically (read tectogrammatically), is ultimately irrelevant to the morphological conception of a word adopted here. The principles determining which constructions it may come to be a part of are syntactic, not lexical, since “met’s” is not a morphological unit.

Moving onto Lieber’s second argument, there may be discontinuous parts in a sentence, each of which appear to form a word prosodically, that need to be interpreted as a single unit. A clear example comes from the the Dutch verb *aan-geven*, meaning (amongst other things) “to let know”, “to indicate”, “to signal”,... whose parts themselves, *geven* and *aan*, mean “give” and “on” respectively. Observe the following examples:

- (1) De dirigent **geeft** de maat **aan**
 The conductor **gives** the beat **on**
 “The conductor **is signalling** the beat”
- (2) Ik geloof dat de dirigent de maat **aangeeft**
 I believe that the director the beat **on-gives**
 “I believe that the director **is signalling** the beat”

Its parts may be separate as in (1), but in certain syntactic constructions, it may also form a single unit phenogramatically, such as in (2). Moreover, when the parts have different meanings individually, away from such constructions, *aan* is a preposition roughly meaning “on”, and *geven* means “to give”. There is a sizeable class of such verbs in Dutch, whose (mostly prepositional) prefixes are separated with the same distribution, and many of whose meanings are similarly opaque (see footnote 1 of this chapter), such as *voor-zien*, meaning “to provide”, *op-slaan*, meaning “to save”, “to store”, *mee-maken*, meaning “to experience”, “to participate”, and so on.

This argument is more relevant to the morphological conception of the word, as it suggests that the third person present singular form of the lexeme *aangeven* item that can be stored in the lexicon in the appropriate paradigm cell as either *geeft... aan*, or *aangeeft*. Insertion of the former into a sentence may limit syntactic economy, as its non-concatenative structure requires a circumfixation operation, and this is unavoidable given that the meaning of such items can be decomposed no further, despite its phenogrammatical separability. This is therefore not as bad as the case of *WRAP* in section 1.1, as it only entails greater phenogrammatical complexity, with no corresponding increase in tectogrammatical complexity. The unseparated *aangeeft* could be seen as a special realisation of *geeft... aan*, where the parts inserted in the opposite order with the empty infix.

Besides this, it is routinely accepted that a similar mechanism is needed to cope with idioms whose meanings appear similarly ‘non-compositional’ despite the various discontinuous syntactic parts, each being prosodic words, that they may have. *Kick the bucket* would perhaps be the canonical such example. It is

thus conceivably not the greatest concession to extend such a mechanism to the case of separable prefix verbs.

Furthermore, word-and-paradigm theories (see section 3.2) often hold a place for what is termed ‘periphrastic’ morphology. This is where a paradigm cell, may contain more than one syntactic unit, at least phonologically-speaking. For example, the perfect tense part of the paradigm of the English lexeme **go** may be constituted by the periphrastic forms *has gone* for the third person singular cell and *have gone* for all the others. See table 4 for a clearer and more systematic example. In much the same way, the cells of the lexeme *aan-geven*, similarly for the whole class of such verbs with separable prefixes, systematically provide but one unit semantically even when the item comes to be separated phenogrammatically by the sentential syntax.¹ Importantly, no word-forms are decomposed into smaller meaningful compositional elements; rather, the periphrastic forms are no longer syntactically atomic.

In a related, albeit perhaps more striking manner to cliticisation, morphology is seen to come apart from the prosody and more generally the phenogramatics. Previously, the clitic was tectogrammatically atomic and had its own compositional semantic contribution, but it was not a prosodic word. This is no problem if prosodic units smaller than words can be seen as lexical items, the basic inputs to an interpretation. In these latter cases, however, separable affixes and periphrastic parts may several constitute full prosodic and phenogrammatically atomic word-like units, but only receive an interpretation (as a simplex) in combination with other such units. This presents no issue if several prosodic, phenogrammatically atomic words are stored with one meaning, even if they surface discontinuously in a sentence. This may decrease (phenogrammatical) syntactic economy, but it is perhaps a necessary trade-off to maintain compositional transparency, especially considering the difficulty of associating meanings to the parts.

The morphological conception of the word proposed here is thus seen to be more directly relevant the process of interpretation than the phenogramatics. Note, however, that this does not conflict with the general argument that morphologically complex items should not be decomposed because compositional transparency should extend as widely as possible into the phenogrammatical level of the grammar. It is only to rescue compositional transparency that morphology, in terms of the storage of basically meaningful items in the lexicon, is taken to supervene on the phenogramatics. There may be other morphology-internal, paradigmatic arguments for this supervenience, however.

The typical justification given to periphrasis is that periphrastic forms are

¹These prefixes are prepositional, and there are greater degrees of transparency between the prefix and the verb. The meaning of *aan-geven* may be primitive, but for other verbs, such as *weg-nemen*, “to remove”, or “take away”, which would be the literal translation of its parts, the parthood is clearly more semantically transparent. This is similar to the case of auxiliaries for composite verb forms, such as the future tense *will go*, where the auxiliary appears to carry the tense meaning. While these may be more conceivably decomposable than single morphological units such as *went*, or even *aan-geven*, for expository purposes, the stance is not softened to exclude periphrastic forms from receiving a single interpretation for now, but more subtlety and flexibility may perhaps be desirable for such cases.

systematically related to other paradigm cells whose items are in fact monolithic. The nominal system of Tundra Nenets, a Uralic language, provides one such example (Ackerman & Stump 2004), which is pictured below in Table 3.1 by the inflections of the noun *ti* meaning “reindeer”. There are three distinct numbers and seven cases, yielding 21 paradigm cells. Only in four of the dual case-forms is periphrasis to be observed, and this is systematic across the nominal system. Having attributed the appropriate theoretical import to paradigm structure, given that the non-periphrastic word-forms are not interpreted compositionally, but rather with respect to their paradigm cell, it is a simple extension to include the periphrastic forms into this scheme.

	singular	dual	plural
nominative	ti	texəh	tiq
accusative	tim	texəh	ti
genitive	tih	texəh	tiq
dative	tenəh	texəh n’ah	texəq
locative	texəna	texəh n’ana	texəqna
ablative	texəd	texəh n’ad	texət
prolative	tewna	texəh n’amna	teqm

Table 3.1: Tundra Nenets noun paradigm

Morphology *qua* the paradigmatic structure of the lexicon may be seen to provide the simplex interpretations of these potentially phenogramatically discontinuous units, while the syntax determines where these units surface in a sentence. In such a way, morphology, in terms of the listing of individually meaningful elements in the lexicon, may be seen to encroach on and override the typical domain of application of sentential syntax, in that the ordering of syntactically atomic elements is one step removed from an interpretation, being supplanted by the listings of the lexicon.

Apart from the paradigmatic organisation of the lexicon in Tundra Nenets, the four periphrastic forms above could be considered in fact to be not only phenogramatically, but also tectogramatically composite. We could assume that *texəh* carries the dual and lexical meaning, and that the other periphrastic parts carry the respective case meanings. This is compositionally transparent for each form in isolation. With respect to the broader morphological system, however, the transparency is inconsistent given the otherwise atomic forms inhabiting the neighbouring cells. In light of this, it affords greater theoretical consistency and uniformity to maintain that the periphrastic forms receive simplex interpretations. In such a way, the paradigmatic structure of the lexicon, can be taken to override isolated instances of apparent decomposability when they are appropriately contextualised with respect to the morphological system at large.

This distinguishes the morphological concept of a word further, as a lexical item inhabiting a paradigm cell, stored with its meaning, from both the prosodic and syntactic concepts of the word, which can form the atomic parts

of a sentence both pheno- and tectogrammatically. Broadly speaking, the semantic contribution of words in these latter two senses to the whole sentence typically coincides with the meaning of the former, namely when the word is a continuous string, but the concepts may be seen to come apart for languages with sufficiently complex morphological systems. Again, this is little different theoretically from the way in which the storage of items on the one hand, and syntactic and prosodic word-hood on the other are misaligned for idioms, except that in the case of periphrasis, it is the paradigmatic structure of the lexicon that explains the misalignment beyond any apparently non-compositional status.

It must therefore be possible for the morphological word, considered roughly as a paradigm cell, not to have the special status of being prosodically monolithic, given the discontinuity of periphrastic morphology, or of the class of verbs with separable prefixes, but this is no great concession given the way in which the transparency of interpretation can be maintained through the storage of such items systematically with respect to the broader morphological system. And to reiterate, the resulting reduction in syntactic economy is conceivably less than that required for *WRAP*-like operations at the level of sentential syntax.

As such, the fact that stored lexical items can both under- and overshoot phenogrammatical atoms does not necessarily refute the morphological concept of a word offered here. This requires morphological words to be conceived of as the systematically stored items of the lexicon with their basic meanings that are direct inputs to a compositional interpretation, whatever their ultimate manifestation. This does not contradict the above claims of compositional transparency being dependent on the phenogrammatology, however, as the lexicon, in terms of its paradigmatic structure, can be seen to take over the role of syntax in interfacing between forms and meanings.

3.1.2 Special Meaning

The lexical meanings attached to words, or lexemes, taken apart from whatever grammatical function they may have, are what is taken to count as their ‘special meaning’. As mentioned above, special meanings are not exclusive to words, with larger phrases such as idioms also having idiosyncratic meanings. This was dealt with previously as a matter of storage, where, if morphology concerns the systematicity of the lexicon, and the lexicon is simply seen as the store of basically meaningful linguistic elements, then any element with a basic, idiosyncratic meaning can be stored in the lexicon and dealt with by morphology.

As a fanciful example, **kick the bucket** may be seen as a lexeme, with its own paradigm providing particular forms such as *kicks the bucket* in the third person singular present cell, *kicked the bucket* for the past tense, and so on. Its forms would be simple to deduce if the paradigm of the lexeme **kick** is already known, although perhaps some extra analysis of its syntactic structure would be required to avoid the incorrect third person singular form *kick the buckets*, or past tense *kick the bucketed*.

Marantz’s argument progresses with the claim that specific kinds of syntactic

structures carve out a ‘locality domain’ for the association of a special meaning to a unit. This argument is not directly relevant to the task at hand, beyond providing a putative constraint on the number and distribution of syntactic atoms that could be stored all with one meaning.

The claim becomes more relevant, however, when this syntactic structure is claimed to be word-internal, with non-lexical sub-word units resulting in idiosyncratic meaning when composed with a lexical meaning. The Inferential-Realisational conception of morphology adopted here does not hold a place for word-internal syntactic structure, but the evidence may still be problematic in a different sense.

A cross-linguistic constraint is evidenced, where only stative and not eventive passives can be associated with idiosyncratic meanings. In English, one such case might be from the sentence “they were knackered”, where the meaning of *knackered* is the perhaps more idiosyncratic stative one of being tired, as opposed to the possible eventive meaning of having been worn out or damaged from the original verb *to knacker* (as in “his knees had been knackered by all those years playing football”). The morphological data evidencing this constraint comes from Chichewa, a Bantu language spoken in southern and southeastern Africa, and is originally taken from Dubinsky & Simango (1996). It is as follows, where AGR signifies agreement, PROG, progressive aspect, PASS, passive voice, and STAT, the stative:

- (3) Chimanga chi-ku-gul-idwa ku-msika
 Corn AGR-PROG-buy-PASS at-market
 “Corn is being bought at the market.”
- (4) Chimanga chi-ku-gul-ika ku-msika
 Corn AGR-PROG-buy-STAT at-market
 “Corn is cheap at the market.”
 [idiomatic reading of “buy” in the context of the stative suffix]

While the presentation of the above presupposes sub-word divisions, the reader is implored to imagine *chikugulidwa* in (3) and *chikugulika* in (4) as single words (as well as *kumsika*) with their respective meanings of “is being bought” and “is cheap”. A second example is provided in the original text of the passive *chinalimidwa* and stative *chinalimika*, meaning “was cultivated” and “was bountiful” respectively.

It ought to be mentioned that the stative is in general productive in Chichewa, but, as is noted in Dubinsky & Simango (*ibid.*), its effects appear more ‘lexicalised’ than for the passive. This is because the stative is seen to affect argument structure; the passive carries the implication that there is someone who carried out the relevant action, in line with the transitive verb’s typical argument structure, but this is not so for the stative. Moreover, causative and applicative formations, do not have stative forms, as they necessarily include an second argument (albeit with a different thematic role in the event).

The change in argument structure is taken to be a likely source for idiosyn-

crazy, as it allows for a qualitative difference in meaning to manifest, given the resultant opacity to the grammar, in comparison to forms of the verb with typical argument structure. Of course, everything is deemed lexical here, whether there is such an effect on argument structure or not. Type-theoretically, however, the lost argument place for stative forms of transitive verbs would render them equivalent to any intransitive verb or simple adjective, suggesting a potentially different set of lexical relations due to the difference in grammatical status.

The problem now is that it is relations amongst paradigm cells that may denote idiosyncratic meanings, as opposed to simply encoding such variation through individual morphemes bearing the potentially aberrant morphosyntactic properties. Importantly, the locus of this idiosyncrasy is therefore not a sub-word unit of any description —that being the original challenge to the word as being a locus of special meaning —but rather, the idiosyncrasy is located in the set of paradigm cells realising word-forms bearing the relevant morphosyntactic property.

The licensing of such idiosyncratic meaning is no longer syntactically specified in the view taken here; the lexicon takes up this role instead, by systematically associating the idiosyncrasy to the appropriate set of word-forms. Given that stative forms appear elsewhere in the language with the typical morphosyntactic property yielding the appropriate meaning systematically, these exceptions are conceivably better accounted for from a paradigmatic perspective.

In the syntactic case, it is unclear how this idiosyncratic meaning results from the morphology by way of a compositional interpretation - at some point in this process the whole word it is part of appears to dictate its interpretation ahead of any part. This inconsistency is not present in a paradigmatic explanation, where the idiosyncratic exceptions to a typical interpretation for stative forms become systematic in a sense, as they form a cluster, holistically, for the relevant lexemes, or sub-paradigms of those lexemes. This is loosely analogous to the case of deponents presented in section 2.2.1, where their idiosyncratic properties are to have active meaning despite passive morphology, and never to be able to have passive meaning. They again form of cluster of lexical items across the entire Latin verbal system.

3.1.3 Special Structure

The special structure of a word concerns the idea that complex morphological structure, whatever it may be, can receive a simplex interpretation. The original Lexicalist argument claims a distinction between readily syntactically interpretable inflectional morphemes, or other more productive derivational morphemes from the idiosyncratic grammatical character of non-productive derivations. Chomsky's (1970) example was the contrast between the productive, gerundal nominalisation “growing” and the non-productive derivation “growth”, where the latter no longer carries the syntactic properties of the verb; we can speak of “John's growing of tomatoes”, but not “John's growth of tomatoes”.

The above characterisation of an Inferential-Realisational theory extends a simplex, atomic interpretation to all morphologically complex words, and leaves no place for an explanation of their meanings in terms of a compositionally interpreted internal syntactic structure. This is a different and significantly stronger position than that of Lexicalism. The way in which interpretations are instead explained is holistically, with respect to the paradigmatic structure of the entire lexicon. Word-internal structure does therefore still serve a role, but it is in the discrimination of the simplex interpretation of the (in a different sense) morphologically complex word, as opposed to being of any compositional import.

This is most clear for inflectional morphology, where the paradigmatic structure is most relevant in providing the appropriate interpretation, as is outlined in section 4. Derivational morphology still provides the most recalcitrant cases, however. Marantz (ibid.) suggests that the derivational morphology of verbs such as *transmission*, *ignition*, and so on, does in fact carry the semantic implications of both the verbs that are their derivational base, with their aspectual and other qualities, and the nominalising suffix.

Of perhaps greater relevance is the claim that this is no longer a special word-specific structure, and that it is instead just the same as syntactic structure elsewhere. This is furthered by the observation that *kick the bucket* does not have the same aspectual properties as *die*, but rather carries with it those of a transitive verb and a direct object, and it cannot be used with stative aspect such as *‘he was kicking the bucket for three weeks’, unlike *die*, which supposedly could be (although the sentence ‘he was dying for three weeks’ seems somewhat strange to me at least, but nonetheless not entirely ungrammatical).

Let us not examine this particular argument thoroughly here, but only answer the question of whether, if these non-productive derivations do indeed carry the semantic content in some sense of their derivational bases and affixes, then how could this be accounted for in terms of the present theory. There is a significant body of work on this subject from other (mostly anti-lexicalist) theoretical standpoints; see, for example, Borer, 2012 (ch.5), or Borer, 2014. As for the theory provided here, this content would, of course, have to be recovered somehow from the network of the relations of the lexicon instead of from any internal syntactic structure. The following example serves to demonstrate how this might be seen to work.

Lieber (ibid.) gives the following sentence as an example of how morphological structure appears to be visible to the syntax (and therefore the interface with the semantics):

- (5) I still consider myself a California_in though I have not been there_i for years.

Seeing as *there* refers to California, seemingly a part of the derived demonym *Californian*, this derivational formation is assumed to operate through syntactic means, thus making the syntactic part meaning California available for reference. This is of course under the assumption that the anaphoric reference

requires some syntactic identity, which is certainly not a necessary assumption; there is also an argument to be made that there need be no direct linguistic reference at work at all, as sufficient contextual information is present to make a pragmatic inference to the referent (see Hankamer & Sag, 1976).

If we do accept that initial assumption to some degree, however, the first issue to note is the following. What happens for less transparently derived demonyms, such as *Mancunian*, or even suppletive ones such as *Scouser* - would Manchester and Liverpool not be made available for reference in these latter case? If we believe that, in the sentence

(6) I still consider myself a Scouser though I have not been there for years

there refers to Liverpool, then the syntactocentric analysis would force us to postulate an abstract decomposition, where *Scouser* decomposes into two abstract, unrealised morphemes, one meaning Liverpool, and the other being a function that takes a place and returns a person from it. The same is true, albeit to a slightly lesser extent for cases such as *Mancunian* amongst many others. This is likely to be problematic for familiar reasons, namely those outlined in sections 2.2.1-2.2.3.

The alternative preferred here is to accept that the relations amongst these lexemes in the lexicon somehow makes these not linguistically referents available. This is not out of keeping with theories such as Frame Semantics (Fillmore, 1976), where a part of a word's meaning is some 'frame' of other meanings, related through some cognitively relevant semantic domain. For example, an utterance involving the word *buy* would activate the frame of a commercial event, where the thereby related meanings of *sell*, *purchase*, *pay*, *cost*, and so on, are all elicited. Here, the mechanism of relation is provided by the derivational-morphological structure of lexicon, that relates lexemes (or paradigms) to one another. Moreover, these relations are typically transparently phenogrammatically manifested, save for perhaps a few suppletive cases. The general transparency is of course not taken to imply compositionality, however.

In much the same way, the meanings of the derivational forms originally discussed above are provided on the one hand by the set of meanings of other derivational forms that share the same affix. The cluster of nominalisations (or simply words) that bear the suffix *-ion* to which *ignition* and *transmission* belong, all share some kind of meaning given that suffix, perhaps that they are entities of some description, due to their nominal status. On the other hand, the derivational bases, such as *transmit* and *ignite*, may provide other aspects of the semantic content, such as that they denote events of specific natures, with certain aspectual qualities, involving some number of participants, and so on. These sets of connections result in *transmission* having a potentially different semantic contribution to some relevant synonym, for example *conveyance*, whatever that may be.

This is still somewhat analogous to the syntactic case of *kick the bucket*, except that there is a different source assumed theoretically for the association of the related content, namely the (derivational) relations of lexemes (by way

of their paradigmatically organised forms) throughout the lexicon. It is this difference that maintains the specialness of morphological structure and prevents it from being conflated with syntactic structure.

A final point of interest concerns what are known as morphological ‘bracketing paradoxes’. One example would be the word *unlockable*, where, if some internal structure is supposed, then it can be parsed to mean either “able to be unlocked”, or “unable to be locked”. Such cases are highly analogous to cases of syntactic ambiguity at the sentence level. This may then be taken to undermine the concept of ‘the word’ as having a special structure, because it suggests that word-internal and sentential structure are two sides of the same coin, in that they both involve equivalent structural ambiguity. The question that arises is how a word-based theory would account for the apparent ambiguity if no word-internal structure is involved in its interpretation.

While word-internal structure may not have compositional semantic import, it does still specify relations to other word-forms. There are therefore two competing sets of relations. Either *unlockable* is related primarily to the class of derivations which share the prefix *un-*, for example *unfounded*, *uncanny*,..., and also includes words such as *unbeatable*, *unfathomable*, where the alternative parse involves elements that are not part of the language (**unbeat*, **unfathom*), perhaps for reasons regarding aspectual properties of the base verbs, or just simply idiosyncrasy; or *unlockable* is related primarily to the class of derivations which share the suffix *-able*, such as *manageable*, *associable*, and so on.

Whichever set of relations is taken to be primary is what may be said to determine the meaning of *unlockable*. Its meaning is taken to be lexical and therefore in a sense primitive, however, so the question arises of how to cash out the evidently systematic differences in meaning. One way to do this, following Borschev & Partee (1998), would be to involve a richer lexical semantics in compositional interpretation, where a series of suitable meaning postulates captures the necessary conditions that derive from a lexical item’s definition. When composed in larger expressions appropriately, this has the result that the meaning of a sentence is no longer a proposition, but instead a whole theory (set of formulae) that describes the model of the situation expressed by a sentence.

In our case, there are two lexical meanings, that of *unlockable*₁ (*un-lockable*), and that of *unlockable*₂ (*unlock-able*), corresponding to the two possible sets of primary lexical relations. Each involves a specific set of meaning postulates that leads to a particular set of resultant inferences such that we get the desired result when each is part of a broader expression. When we have a sentence that in our logical language is something like *unlockable*₁(*x*), then it comes along with another formula, something like that *¬lockable*(*x*), and so on, until the maximally explicit meaning that it is not possible to put *x* in the state of being locked is attained. If it is instead the case that the sentence is something like *unlockable*₂(*x*), then we would have an accompanying formula that is something like *◇ unlock*(*x*), and again until the maximally explicit meaning is attained.

It is the meaning postulates which therefore provide the appropriate primary and secondary set of lexical relations, and thereby determine the respective scopes of negation and modality that appear to be the source of ambiguity.

This is in a sense a lexical decomposition, except that it is not manifested through a compositional process involving the internal structure of *unlockable* directly, but rather it is achieved through the relation of *unlockable* to other words through meaning postulates, which may in turn be composed in whatever larger expression the derivational form may find itself in (again, to yield a theory as its full meaning). A full analysis that explores the particulars of the semantics in detail is not presented here, as it is hoped this brief outline demonstrates a sufficiently compelling suggestion for such a non-decompositional explanation. It also hints towards a method to probe the semantics of derivational morphology from a word-based perspective in future research.

It appears therefore that the notion that words have a special structure can be maintained thusly, in that it is non-compositional and as such distinct from typical sentential syntactic relations. Instead, the internal structure of whole word-forms indicates their relations to other lexical items. The specific set of relations can then come to impact its ultimate meaning in larger sentences. At no point are meanings assigned to its parts individually, in the same way as is done to the parts of a sentence; rather, its parts (by way of their structure in the word) serve to discriminate the sets of related forms that in turn determine its possible (enriched) lexical meanings.

3.2 Words and Paradigms

With the morphological concept of the word adopted here defended, it is now pressing to engage more concretely with the word-based theory that is informing the discussion at hand. Word and Paradigm (WP) Morphology (Blevins, 2016; Blevins, Ackerman & Malouf, 2018; Blevins, 2006) provides the model of morphological analysis adhered to in the present investigation. It is pertinent to review its tenets before undertaking a semantic implementation that takes inspiration from it in chapter 4, following a sketch of its more general philosophical significance in providing a holistic alternative to compositionality in the determination of words' meanings.

WP theories are in essence the kind of Inferential-Realisational theory outlined in 2.2.4, that conceives of the (inflectional) morphological system as a network of relations amongst the whole word-forms inhabiting paradigms organised by lexemes and morphosyntactic properties. It was remarked in the discussion in 2.1 concerning the nature of paradigmatic 'rules', that they are conceived of not as rules *per se*, unlike paradigm functions that take the pair of a lexeme and a set of morphosyntactic properties and return the word-form inhabiting that cell. Instead these paradigmatic rules are more so relations across an already-given network of morphological patterns that licenses proportional analogies for the language learner. This can now be more clearly outlined, as it relates directly to the claims of WP theories on how they are distinguished from other morphological theories, even those that may on the surface appear to be close cousins.

Firstly, while we have been speaking up to now in terms of (inferential-

)realisational theories, WP would in fact deny that it is realisational in the full sense of the term, claiming instead that it is an *implicational* theory. For the above evaluation in terms of Stump's (2001) taxonomy of morphological theory, this subtle distinction was not all too important, as whole words are ultimately given primary morphosyntactic relevance in either case. It is nonetheless important for a discussion of paradigmatic relations conceived of as licensing analogies to word-forms from others, in relation to the issues of learnability and processing treated in the following two subsections.

The distinction is made because certain Inferential-Realisational theories, such as those that involve paradigm functions, are seen to fall prey to one aspect of the 'constructivist fallacy', relating to the abstractive-constructive distinction introduced in section 2.2.4. This 'fallacy' concerns the way in which the individual forms that realise paradigm cells are assumed in constructive theories to be derived in isolation from the rest of the morphological system, per instance of derivation. Inferential-Realisational theories that involve paradigm functions also make such an assumption, because they take it that word-forms are produced in a given syntactic context given bare lexical and morphosyntactic information to feed the paradigm function.

It may be that the specific Inferential-Realisational theory claims that the morphological spell-out procedure makes reference to other paradigmatic forms in computation of the output, either 'extensionally' through the assignment of a shared exponent, or more 'intensionally', by applying the same rule or set of rules as would be involved in the production of another form, but neither involves any necessary nor essential reference to other word-forms that themselves inhabit the paradigm cells. This undermines the robust status given to whole word-forms as the nodes in the network of paradigmatic relations. In such Inferential-Realisational theories, it is thus the coincidence of a number of morphosyntactic properties, and lexical information, taken apart, which are considered most directly theoretically relevant, ahead of the forms realising them.

While the above discussion of Inferential-Realisational theories avoided the constructivist fallacy (and any associated semantic issues), because of the adoption of a WP perspective, it is helpful to reinforce the primacy of word-forms ahead of grammatical features, which relates to the self-attribution of WP theories as being 'implicational'. An implicational theory emphasises the inferential component, where paradigmatic rules provide the patterns of implication that relate full word-forms by way of the structure of the morphological system. There are thus no hollow cells that only realise the word-forms per instance of derivation - a paradigm cell in the lexicon of an individual speaker may be empty before it is learnt by analogy across a pattern of implication involving related forms, but not thereafter. The language itself, abstracted from particular speakers, or to an ideal speaker with perfect linguistic knowledge, may be conceived of as having all its paradigms full.

One word-form implies another based on the paradigmatic structure of the lexicon, thereby lending theoretical primacy to the word-forms themselves. The paradigmatic structure organised by the morphological system may be equated with a network of patterns of implications amongst related forms, by which

learners are pointed towards what the correct forms of unencountered paradigm cells would be given knowledge of the word-forms inhabiting other related ones.

Importantly, any set of such relations may be used to provide the appropriate analogy, with there being no special status attributed to any ‘diagnostic’ forms. For example, singular forms are of no greater predictive value than plurals, because of some perceived greater simplicity. The benefit of this is not only to limit pre-theoretic assumptions concerning markedness, but also to allow for the network of implications to be cast in statistical terms, accurately reflecting natural language as it exists in reality for any speaker (or hearer). That is not to say such matters as markedness are denied outright, but rather only that the theory is agnostic towards them. There are a number of more practical consequences of this conception that are borne out in the following two subsections. These chiefly concern the relevance of such a theory to a language-learner and -user, demonstrating how such a conception of the lexicon is able to lend psychological plausibility to WP approaches.

This is particularly relevant given that two general assumptions around the psychological plausibility of the principle of compositionality concern the reality of language learnability and use. If the language as a whole is considered a (potentially) infinite object, then compositionality provides the language learner and user with a finite set of items and rules that can be learnt and according to which which novel complex expressions (such as whole word-forms) can be understood (and produced).

While these are the typical assumptions associated to the principle, a critical assessment of the at times unexamined arguments is undertaken in Pagin & Westerståhl (2010). In any case, the argument concerning the (potential) infinity of language may not apply for the morphological system as it is conceivably finite, since it the positing a recursive mechanism in the production of morphological expressions is in general less warranted than for sentences (a case where it is possibly warranted is mentioned in section 4.5). If the conclusion is indeed reached that the original arguments do not hold water, then there is at least a weaker argument, namely that compositionality may reduce the computational complexity of interpretation and communication in general, with these processes evidently taking place online and immediately.

Complexity provides another important target of WP approaches, as morphologically complex languages are conceivably more complex, simply due to the greater number of word-forms and the impact this has on the structure of the lexicon. This sense of morphological complexity is particularly relevant to a language learner, as it could be taken to imply variation in the learning process that is not in fact exhibited in reality given that language acquisition takes place over a roughly equivalent time-scale cross-linguistically. The ideas of WP approaches relating to these concerns are relayed in the following subsection.

In another sense, however, morphologically complex languages may be simpler, as they have the capacity to express the meaning of a sentence with fewer words, thus resulting fewer steps in the compositional process of interpretation, as well as the words involving potentially less derived syntactic categories and correspondingly less abstract semantic types. This would be relevant to ques-

tions concerning natural language processing (and production), where, cross-linguistically, humans are seen to communicate using natural language with the same degree of efficiency, in relation to the received notion that they are beholden to the same cognitive mechanisms.

Arguably, this would make little difference if a coarse-grained measure formal computational complexity is employed, as, assuming both morphologically more complex and simpler languages are beholden to the same kind of compositionality at the sentential level, the same degree of complexity would be implied for both, regardless of the then negligible variation in the number of atomic parts in each of their sentences. This sense of complexity may be considered not to be directly psychologically plausible, however, when articulated so in terms of a formal definition of complexity. In a more intuitive sense, the question remains, therefore, of whether there should be a more fine-grained measure of complexity relevant to morphologically complex languages in WP approaches, perhaps related to the paradigmatic structure of the lexicon. This is treated in the subsequent section (3.2.2). It is pertinent since compositionality is not entertained within whole word-forms, with their interpretation being simplex, so there is no such procedure to which to associate any conceivably processing effects of morphological complexity.

3.2.1 Learnability

The patterns and networks of relations amongst full word-forms are what provide the object of research for WP models. These patterns may be highly complex. This complexity is organised in some way, however, if it is to license analogical predictions of full word-forms from others. On the one hand, there is the organisation in terms of the space of inflection defining paradigm structure - this is an abstract theoretical means of describing the organisation. On the other hand, more concretely reflecting the way in which that abstract organisation is relevant to a language learner are the surface forms constituting any particular cells, and moreover the statistical distribution of these elements that the learner is confronted with in their linguistic exposure. The interpredictability of word-forms from others across the lexicon, by way of its paradigmatic structure abstractly, and in relation to its statistical distribution concretely, is what can offer a measure of the complexity of a morphological system that is relevant to the question of learnability.

Information theory is the tool employed to measure the complexity of different morphological systems (Blevins, 2013). This is because there are asymmetries amongst the relations across the forms that fill paradigm cells that could be exploited by a learner. Take the case of Russian nouns. Below is the paradigm of the second declension noun *dynja* meaning “melon”. It is a fact of the Russian nominal system that all three declensions share the exponent that marks the morphosyntactic properties LOCATIVE and PLURAL. This means that from the form *dynjax* alone, it would be impossible to predict any of its other forms of this lexeme, because this cell happens to give no information about them, in that it does not limit the alternatives inflectional classes to which *dynjax* may

belong. By contrast, the nominative and accusative singular forms are unique identifiers of this declension. As a result, these cells provide more information. A middling predictor would be the locative singular cell, whose form is shared only amongst the first two declensions.

	singular	plural
nominative	dynja	dynji
genitive	dynji	dynj'
dative	dyne	dynjam
accusative	dynju	dynji
instrumental	dynej	dynjami
locative	dynje	dynjax

Table 3.2: Russian second declension nominal paradigm

The acknowledgement of such asymmetries amongst the implicational relations between word-forms by way of their place in the broader paradigm structure of the lexicon allows the avoidance of any explanation in terms of *ad hoc* inflectional class features that are individually meaningless and hold no place grammatically other than to stipulate idiosyncratic form-to-meaning relations for items in particular inflectional classes. Moreover, if the interpredictability can be measured across the entire morphological system of a given language, then it is possible to compare the complexity of different languages.

We can thus consider a cell to be a random variable, C , where the differences amongst forms across inflectional classes provide the number of possible kinds of forms ($|C|$) that it can take. For example, in the Russian nominal system, the nominative singular cell has three values, while the locative plural has one. For expository purposes, one may simply assume a uniform distribution across inflectional class forms, with this giving the worst case complexity (as uniform distributions involve maximal uncertainty); frequency distributions could be factored in where available, and lead to further insights. (Shannon) entropy is the information-theoretic concept that intuitively captures the capacity for uncertainty reduction, or informativeness of a random variable. Its typical precise formulation is not entirely relevant for the present purposes, but a simplified measure of entropy (H) of paradigm cells can nonetheless be given here, with respect to $|C|$, or the number of different kinds of forms that could realise the cell across all inflectional classes. We see that for the Russian noun system, with three declensions, the locative plural, with one exponent across all classes takes the minimal value (a.); the locative singular is more informative (b.), with one form for the first two declensions, and another for the third; and the nominative singular is maximally informative (c.), having a different form for each declension:

- $H(C) = \log_2(|C|)$

- a. $H(C_{\text{LocPl}} = 1) = 0$

- b. $H(C_{\text{LocSG}} = 2) = 1$
- c. $H(C_{\text{NomSG}} = 3) = 1.58$

Entropy therefore increases as the number of different alternatives of word-forms realising the cell increases (and as the uniformity of their distribution increases). From this measure of cells, a joint entropy measure can be taken of all cells across a paradigm, yielding the paradigm's overall informativeness. This measure accounts for the possible interdependencies across paradigm cells. This ought to be no greater than simply the sum of all cell entropies, which the joint entropy can only be as great as if the cells are statistically independent (and it can never be greater).

This prediction is termed ‘the Joint Entropy Conjecture’ (Blevins, 2013) and is of course testable. It is widely observed to be true, such as for the above case of Russian, where the nominative singular and the accusative singular are interdependent, being the same across all paradigms in all inflectional classes. It is these interdependences that license analogies reliably due to their informativeness. If one knows both forms of one noun of a given declension, such as those of *dynja* above, and one knows the nominative singular of another noun in the same declension that shares this form, then its accusative singular form can be robustly deduced (as well as all its other forms, since the nominative singular predicts declension membership uniquely).

One important affordance of such techniques is that a WP morphologist can measure a different kind of complexity associated to morphological systems than what is apparent on the surface. Ackerman & Malouf (2013) develop this notion into the distinction between ‘enumerative’ and ‘integrative’ complexity. The former amounts to a measure of simply the total number of possible forms that may occur in any given language, which may of course be vast. If this is the kind of complexity assumed when probing typological variation, then, given the WP approach, it is difficult to explain how a language learner can learn languages of greater enumerative complexity within the same time-frame as their simpler counterparts. Even if the decomposition of word meanings were granted, the phenogrammatical variation of the encodings of the same underlying structures could remain problematic for learners.

Integrative complexity, by contrast, better measures what learners are actually confronted with, by accounting for the systematic interpredictability across the word-forms of a given language's lexicon according to its paradigm structure. This is quite conceivably exploited by learners to facilitate the acquisition process. The specific measures and results are reported in Ackerman & Malouf (2013), but to summarise, the relevant claim, termed ‘the Low Conditional Entropy Conjecture’, states that the ‘conditional entropy’ amongst related word-forms ought to remain significantly low cross-linguistically; in other words, a measure of the mutual predictability of the forms in a particular language given knowledge of other related forms remains sufficiently low throughout any morphological system. This means that the enumerative complexity of a morphological system can be extremely large, as is evidenced cross-linguistically, provided that its integrative complexity, quantified so, remains significantly low. And

this is indeed what is found by the authors in a small, but typologically various survey of languages of differing morphological complexity.

This is relevant to the present investigation in the following ways. Firstly, if compositionality is denied within words, there is conceivably be some other means by which language learners can learn new word-forms with the associated meanings given their existing linguistic knowledge. This is particularly important if, as in WP theories, the number of words increases in line with the complexity of the morphological system. How can all the forms representing the meanings that the language being learnt is capable of expressing be acquired in the first place by the learner, in a process of acquisition that is comparable cross-linguistically? This may be termed ‘the Paradigm Cell Filling Problem’ (Ackerman & Malouf, 2009).

Secondly, if compositionality does not apply to complex word-forms, then there is no immediate reason why morphological systems should not be completely suppletive, with all paradigm cells being phonologically unrelated, since, on the face of it, the enumerative complexity would remain the same. Systems only ever provide minimal number of suppletive forms, however. This is because suppletion, if widespread, would drastically increase the integrative complexity of the system, since there would be no interpredictability amongst morphosyntactically related forms. This would make such languages unlearnable. By contrast, the limitation on integrative complexity that is afforded by the phonological patterns amongst paradigmatically related word-forms, whatever their phenogrammatical complexity, is what renders morphologically languages learnable.

Thirdly, something that compositional explanations of morphological systems do not explain, which the notion of integrative complexity does, is so-called ‘morphomic’ patterns (Aronoff, 1994), which is when morphological forms are related, while their morphosyntactic properties are not obviously so. An example relayed in Stump (2018), originally from Haiman (1980) is of Hua, a Papuan language, where the second person singular and first person plural are consistently the only verb forms that share respectively non-default suffixes systematically across all 12 moods. This is because they feature in an implicative relation, as opposed to participating in pattern that involves a morphosyntactic property with its corresponding meaning. There is therefore no compositional explanation of this relationship; rather, it is morphology working independently from the rest of the grammar, organising a lexicon of whole word-forms in such a way that it may be learnt efficiently through the resultant implicational relations.

Now, since complex morphological systems are conceivably finite, not being a recursively computed system, it is arguably the case that the typical argumentation from compositionality need not apply to the relevant questions of learnability. Nonetheless, compositionality would provide a handy means to decrease the complexity of the acquisition process. As discussed above, however, this may only be true for highly iconic agglutinative morphological systems, where any deviations from this, due to non-one-to-one form-to-meaning relations, are conspicuously left out. Integrative complexity is a notion that allows

for questions of form-to-meaning transparency to be overlooked to some extent, and thus in no way privileges concatenative morphology. It provides instead a holistic measure of morphological systematicity, and thereby treats concerns that could arise when the principle of compositionality is abandoned in morphological analysis.

3.2.2 Processing

While the question of how novel morphologically complex words can be learnt (and subsequently understood and produced) by language users in a WP approach is best treated in terms of the above information-theoretic notions employed to probe learnability, the (synchronic) psychological plausibility of paradigms is reinforced by a certain line of psycholinguistic research. The reader is recommended to consult Milin, Kuperman, Kostić, and Baayen (2009) for a comprehensive summary of several of the results.

For the present purposes, it is pertinent to relay in general terms the finding that there are processing effects linked to the structure of paradigms. Importantly, paradigms are to be conceived of statistically, as was hinted at in the previous section, where individual word-forms provide token-frequencies, and the families they belong to —be they inflectional or derivational— type-frequencies. These figures can be correlated to the response latencies of participants in visual decision tasks involving words for which frequency data is known. It is generally taken to be the case that paradigmatic structure, conceived of statistically, results in lexical access effects, a concept which is at worst neutral with respect to the question of whether a word’s meaning is computed, although it intuitively appears more aligned with the claim made here that it is stored as a simplex with the whole word-form that denotes it.

The results of these experiments suggest that there are indeed ways in which the processing of language users depend in some way on the paradigmatic structure of that language’s morphological system, insofar as it is manifested statistically throughout the linguistic community. A set of simple experimental results from Schreuder & Baayen (1997) can help to give an idea of the kind of processing effects that are probed. Firstly, the processing of pairs of monomorphemic Dutch singular nouns with roughly the same token frequency, but varying stem frequencies (the cumulative frequency of their singular and plural forms) was tested through both a visual lexical decision task and a subjective frequency rating. For example, in the corpus they used, *akker*, meaning “field” and *gif*, meaning “poison”, had frequencies of 214 and 213 respectively for their singulars, but 404 and 0 for their plurals. It was observed that those with the higher stem frequency produced shorter response times and were rated higher.

This was followed by a comparison between groups of nouns of the same kind, with similar stem frequencies, whose morphological families (the set of derivational forms and compounds containing the original noun) differed in size, this being a measure of type-frequency. Variation in the cumulated token frequencies of the morphological family members of these nouns with comparable token frequencies was also tested, but no result was found, which led to the

type-frequency test. For example, *smart* meaning “sorrow” and *rente*, meaning “interest”, had similar stem frequencies, while the number of derived forms involving each was 3 and 23 respectively. It was found that words with larger family-sizes produced shorter response times and were rated higher. The experiments demonstrate that the domain of inflection appears to be sensitive to token-frequency, while for derivation, type-frequency is seemingly of greater importance. We do not wish to interpret what this could mean exactly for a WP model, beyond noting that there is some observable difference between the two morphological domains, that is reflected in the theoretical distinction built into WP approaches.

Something a WP model may be expected to account for is that there is a source of complexity other than a putative compositional semantic interpretation that could be relevant to the processing of morphologically complex whole word-forms. We have already observed that enumerative complexity is not relevant to the learning process, and that integrative complexity, which is conceivably relevant to learnability, remains low cross-linguistically. The source that would be relevant to processing, similarly to integrative complexity, ought to derive in some way from the holistic conception of morphology, where the system at large, through the paradigmatic structure of the lexicon, may be seen in some way to result in some effect. As noted previously, this may be taken to relate to the fact that morphologically rich languages appear to be able to express the same sentence meanings as their impoverished counterparts with fewer syntactic atoms and correspondingly less complex and rigid sentence structure, where freer word order is often exhibited.

Because the processing effects are probed statistically, through the distributions of paradigms and their members, a means of relating the observed effects of the (statistical) structure of the morphological system to a deterministic, and more abstract theory of interpretation is not readily apparent. Indeed, it would perhaps be unwise to interpret the results in such a manner. Nevertheless, for individual forms interpreted online in sentences, which may not even have been previously encountered and therefore not stored in the hearer’s individual lexicon, there should conceivably be some way for the distribution of elements throughout the morphological system to impact the process of interpretation of the sentence.

This is indeed what is observed, such as in the results of Schreuder & Baayen (1997) relayed above. Another notable result of this kind, that perhaps more directly relates to the theoretical issues surrounding paradigm structure, was found in Moscoso del Prado Martín, Kostić & Baayen (2004). The authors observed that, the greater the number of inflectional variants of a given word (or lexeme) —that is, the greater the size of a lexeme’s paradigm —the shorter the response latencies in lexical decision tasks were. This held true for lexical items in the same lexical (or syntactic) category as well. This was interpreted by the authors as an indication of the possibility for paradigm structure to influence lexical processing.

Another benefit of the statistical methodology employed by WP approaches is the capacity to probe derivational morphology in a similar manner to inflec-

tional morphology. Indeed, there have been many effects observed that concern the size of the derivational family of a given word, as in the second experiment in Schreuder & Baayen (1997) mentioned above, thus suggesting that derivational type-frequency can be of importance to the processing of derivational morphology (Blevins, 2013).

The question of meaning is not broached in these distributional morphological studies, and at any rate, the establishment of a clear relationship between the (statistical) structure of morphological systems and the semantics of morphologically complex words is perhaps an endeavour fraught with theoretical conflict. The insight remains that there is some significant psychological relevance to the paradigmatic structure of the lexicon. The psychological plausibility of such a theoretical stance reinforces the claim made here that a holistic conception of the lexicon in terms of its morphological organisation presents a principled starting point from which to explore and develop a non-compositional approach to an investigation of the interrelation between morphology and semantics.

3.3 Semantic WP - Morphological Holism

With the holistic nature of WP morphological theory now clear, the question remains as to what the semantic relevance of such a theory could be. Some sort of account of this is demanded if a robust case is to be made for complex word-forms, and not their putative parts, to be the most basic compositional elements to which meanings are assigned, as was argued for in section 2. To preface an actual semantic implementation of a word-based morphological theory for particular morphological phenomena, it is prudent first to probe the more philosophical claims of such a holistic theory.

The holistic character of the morphological system is posed here as a counterpoint to any theory that assumes compositionality within complex words, where that involves some morpheme-based notions in either the morphotactic or morphosyntactic sense. As the negative part of the argument has been made, this section offers the positive angle on what a morphologically holistic theory may entail semantically in general.

There are two senses of holism (Pelletier, 2012) to be distinguished that are relevant to the present discussion:

- **Ontological Holism** (*Wholism*): some properties can only be attached to entities that are not individuals.
- **Property Holism** (*Holism*): some properties of an object are defined in terms of the same type of property of some other object(s) and vice versa.

These are conceived of as the antitheses to atomism and compositionality respectively. They have been relevant to the above discussion in the following ways. Firstly, wholism is akin to word-basedness in the morphotactic dimension of morphological analysis, where it is full word-forms that are taken to be the theoretical entities to which meanings are attached. This corresponds to

what is termed ‘gestalt exponence’ (Blevins, Ackerman & Malouf, 2018), where fully realised complex word-forms are seen as a gestalt, namely a whole that is greater than the sum of its parts. That whole is taken to be the exponent of several morphosyntactic properties all at once, as opposed to any of its parts individually. Such a view is adopted above in light of the perceived phenogrammatical unity of whole words. This is a limited wholism, however, in that it is only relevant morphologically. Beyond words, sentences are conceived of largely atomistically, save for certain idioms and periphrastic constructions.

Secondly, and of greater importance to the semantic investigation below, is how holism provides an alternative to compositionality. If there is no compositionality within words, then some alternative theory is required to yield the meanings of complex word-forms, even if their individual interpretations are simplex; otherwise, any generalisations concerning morphosyntactic properties, as well as connections between derivational forms, are no longer available. Only the space of inflections is treated here, however, where the idea is developed that morphosyntactic properties are defined holistically: the meaning of a single word-form that bears a given morphosyntactic property is whatever it may be only because all other word-forms bearing that morphosyntactic property have their own appropriately related meanings. This ought to apply in some similar manner to the related meanings of derivational forms as well.

Again, it is important to note that this morphological holism is not of the usual kind. In a typical semantic holism, the meaning of any word is dependent on all other words in the lexicon, with this scaling up to sentence-level meaning. For example, the word *dingo* is related to *Australia*, and thus also to *continent*, and then to *globe* and so on, and so on, such that all relations amongst all words determine any individual words meaning.

Here, the holism is limited to the morphological domain of the grammar, as it is represented in the paradigmatic structure of the lexicon, thereby systematically constraining the relations amongst meanings. When words come to be used in sentences, compositional interpretation may proceed as usual, depending on the meanings of the most basic elements, namely fully realised word-forms. This boundary between holism and compositionality at the level of the word may shift somewhat from language to language, of course, in line with morphological complexity. This provides a novel theoretical basis from which to conceive of the typological variation between languages of different morphological complexity that is observable in sentence structure.

This limited morphological conception of holism, which could also be conceived of as a kind of molecularism within the morphological system, is not, therefore, subject to some of the typical critiques of an all-out holism. The first potential problem (Pelletier, *ibid.*) of a strong holism is that it implies that all sentence meanings depend on all other sentence meanings, as their interrelations scale up when words are combined into complexes that inherit the constituent words’ holistic meanings, and add new ones from the interrelations of that complex with other complexes. The conclusion is that every sentence meaning then depends on every other sentence meaning, and a language user must therefore already understand all of the (potentially) infinitely many sentences in the lan-

guage to understand any particular one. This is not an issue here, as both the limitation on holism to the simplex interpretations of the whole word-forms as organised by the morphological system at large, and the acceptance of a compositional interpretation at the level of sentences, means that the same conclusion cannot be arrived at.

A second issue is the question of the internality or externality of the lexicon. There are many problems with an internalistic conception of holism that relate to learnability. Firstly, just at the level of the lexicon, the extremely likely differences in learning experience will lead people to have different lexicons, in that the system of relations amongst words that constitute their meanings will be different. If this is the case, communication becomes impossible, as people will mean different things by the same word. For example, an oncologist and their lay patient no doubt have a very different understanding of what is meant by *tumor*, with the assumption being that this is because the networks of other meanings in which their respective meanings of *tumor* are situated will differ, and yet they should still be able to communicate about any such tumor efficiently and successfully. Pelletier (ibid.) claims that any notion of two people's meanings being 'similar enough' to communicate appropriately cannot be achieved, as it is necessarily circular. Time and space preclude an account of the arguments surrounding the issue of similarity, and other related issues, which are many; suffice it to say that the internalistic-holistic picture is subject to critique.

Some of the problems associated with an internalistic conception of a holistic lexicon may be avoided, however, if one instead adopts an externalistic stance, where the meanings of words are a socially shared, concrete matter of fact, regardless of how an individual learns them. The question of learning for an individual is thus conceived anew as the problem of attaining knowledge of this external entity. Language change is one issue that arises in relation to such an external conception of the lexicon, as any change in the holistic semantic value of a word changes the value of every other word. This would make it impossible to, say, read books from long ago. An analogous issue arises concerning the intertranslatability of different languages. Pelletier (ibid.) suggests that a clearer notion of this external reality, of the ontological status of the lexicon of the language, is required to provide a theory of learning, and could help to assuage worries about diachrony.

At least when it comes to learning, this is exactly what a WP model achieves. Not only is the external, holistic entity properly conceived of, as the morphological system of a language, which is for example documented by traditional grammars, but a theory of learning also follows naturally given the statistical conception of the lexicon, as is outlined above in section 3.2.1. The distribution of exposure learners receive to parts of the external whole guarantees sufficient information from which to learn it in its entirety. See section 4.5 for a discussion of the implications of a holistic semantic WP approach regarding diachrony and typology.

The original problem for learners of an internal lexicon also diminishes somewhat, because a metric of similarity is seemingly provided. Namely, despite conceivably negligible differences, particularly for especially infrequent lexical

items, the statistical distributions of paradigms provide a broadly constant measure of a typical learner's exposure. The differences in the acquisition process are thereby smoothed over for the most part. Learners attain full, internal lexical knowledge because they are confronted with a typical distribution of the same material. The systematic morphological organisation of the lexicon that is indeed the same for all learners further helps to diminish differences in the acquisition process, as learners may infer unencountered forms and their meanings from those they have already learnt. Two learners may have been exposed initially to two different forms of the same lexeme, but come to attain the same knowledge of the whole paradigm by analogy from other paradigms of the same inflectional class.

This means that there are no great problems with either an internalistic nor externalistic conception of the lexicon, and WP theories are therefore neutral with respect to this distinction. In fact, it may be said that they are explicitly agnostic thereto. What is important is that the meanings of whole word-forms, with respect to both lexical and morphosyntactic content, depend on those of others. It is the paradigmatic organisation of the lexicon with respect to these kinds of content —be it internal or external to language users' minds —that allows it to be learnt.

The final point concerns the fact that this is seen as holism with respect to the morphological system, but in fact only as molecularism with respect to the entire lexicon. Molecularism is often critiqued as an unstable halfway house between meaning atomism and holism, since more and more word meanings can keep being added to the relevant clusters of relations. The morphological system delineates the clusters very clearly here, however, where the meanings of word-forms are related systematically, with respect to shared lexical meaning, the inflectional categories encoded by a language, and also possible derivational formations. Only the former two are dealt with here, with an investigation of the potentially more complex holistic character of the derivational morphology left to future research.

As was briefly outlined in section 2.2.4, lexemes and inflectional categories organise the paradigmatic structure of the lexicon and thereby determine the clusters of related meanings. Relations in lexical meaning are internal to the paradigm of a given lexeme, while morphosyntactic relations are across like rows and columns of paradigms. Of course, where some morphological subsystem inflect for more than two categories, one can no longer speak of rows and columns, but this jargon is intuitive, at least.

That the lexical meaning of two inflectional forms in the same paradigm, or of the same lexeme are the same is (hopefully) uncontroversial. For example, it would be somewhat absurd to suggest that *mouse* and *mice* refer to two different kinds of entities.² It is thus the shared aspects of the meanings of

²That is, 'kinds' in the concrete sense of worldly concepts relating to things such as mice, and not in the abstract sense of semantic structures, *i.e.* that they may involve plural individuals, where the difference this latter sense may contribute to certain other notable effects that are not observed for less semantically weighty morphosyntactic properties. See the 'Discussion' section of Baayen, Lieber & Schreuder (1997) for a good summary of many of phenomena

clusters of forms expressing the same morphosyntactic properties that are the primary concern of the following investigation. The relevant questions concern how the semantic contributions of distinctions in inflectional categories come to be realised throughout the lexicon, according to the paradigmatic organisation of morphosyntactically related items.

Chapter 4

Morphological Case Studies

The following 3 sections deal with a series of common (inflectional) morphological phenomena of various natures. A formal semantic interpretation is given to whole word-forms that straightforwardly exhibit the relevant morphological patterns. It is hoped that these examples serve to illuminate some of the potential benefits of conceiving of the morphological system holistically, as an network of interrelated whole words.

Inflectional categories provide the guiding principles according to which domains may be carved up systematically, in order to mark out the denotation of any word-form bearing a certain morphosyntactic property. While only certain inflectional categories are typically deemed to be of semantic relevance, such as number, tense and aspect, if the morphological system is taken to interface with the semantics directly, in a holistic manner, then an attempt must be made to give some place in an interpretation to all inflectional categories.

The examples and a conceivable interpretation for each them - at least a starting point for improvement in future research - is presented in the following order. Number is dealt with first as it represents a morphological phenomenon with a clear semantic contribution, which has already been encountered for example in section 2.2. This is followed by 3 examples involving morphosyntactic agreement in some way. Grammatical gender is a morphosyntactic property, whose semantic status is somewhat dubious despite the examples given of agreement. First pronouns are treated, where anaphoric reference is mediated through grammatical gender, and then the more typically morphosyntactic case of adjectival agreement is discussed. This is followed by a brief account of null subjects, where the apparent compositional effects relating to verbal argument agreement observed in pro-drop languages are investigated, before considering whether that account may be extended to imperatives. Lastly, grammatical case provides a phenomenon whose semantic import is very unclear, if it is to have any at all. There are of course many stones left unturned here, of course, but the hope is that these examples are sufficient to provide an insight into the potential benefits and challenges that arise when trying to implement a WP theory in a formal semantic setting.

Throughout, an attempt is made to characterise inflectional paradigms as organising the word-forms in their cells such that their meanings accord with the morphosyntactic properties they bear appropriately. This typically amounts some sort of restriction on the domains of the functions they are interpreted as. These restrictions pertain to the way in which morphosyntactic properties can organise the domain.

4.1 Number

Returning to the case of the *mouse-mice* distinction, we take common nouns to denote sets of entities. The atoms of any common noun generates a lattice, following a Link-style (1983) analysis, whose non-atomic elements are in turn reified as elements of the domain. The singular forms denote the atoms alone, while the plural denotes the whole lattice, to allow for distributive readings when composed with the relevant verbs. As a recap, a distributive reading would be something like “Mary fed the horses”, where each horse is fed individually. The present analysis could be extended to a number system that has dual marking, where the two-element joins are also denoted appropriately. There are other subtler number distinctions that are not to be treated here, but hopefully they could be treated in a similar manner, with denotations being assigned directly to whole word-forms.

Common nouns themselves are typically taken to be of type $\langle e, t \rangle$, and thus denote sets. In the plural domain, we take AT to be the set of atoms, and SUM to be the set of non-atoms and let us take $PL = AT \cup SUM$ to mark the domain of plurals, as the entire lattice of relevant elements (in the entire domain of entities, there may be other elements such as unindividuated masses, such that it is not the case that $PL = D_e$). Let us then refine the domains of the functions that are the interpretations of *mouse* and *mice* such that they are of types $\langle AT, t \rangle$ and $\langle PL, t \rangle$ respectively, these types being functions that are subsets of all the $\langle e, t \rangle$ functions, namely the partial functions defined on the relevant domains of just atoms, and atoms and sums.

Every common (count) noun therefore has one denotation for its singular form and another for its plural form, just as *mouse* and *mice* do. These morphosyntactic properties are not concretely linguistic elements themselves. Instead, they emerge from the organisation of the morphological system, that relates nouns in this manner, where there is a holistic property across the class of nouns that relates denotations of related word-forms, where the singular includes only the atoms, while the plurals also contain non-atoms, or sums. It is this distinction in meaning that the forms discriminate. As common nouns are constants in the logical language, it is therefore the morphological system is organised in such a way that *mouse* and *mice* respectively bear singular and plural meanings coming from the same lattice.

Now, plurals tend to be more marked morphologically (Corbett, 2000), and this should correspond to them being more marked semantically, as is claimed in and supported by the analysis provided in Farkas & de Swart (2010). This is

not corroborated by the current analysis here, however. Two kinds of markedness are distinguished: ‘denotational’ and ‘conceptual’ markedness. The former concerns a subset relation, where the singular is at present more marked because its denotation is included in the plural’s. Conceptual markedness concerns conceptual simplicity, where the plural is conceivably more complex because it involves more semantically loaded join elements. By reifying the conceptually loaded join-elements as part of the meaning of the plural, its precedence in terms of conceptual markedness is upheld. The number distinctions provided here do fail, however, to account for the greater markedness of plurals denotationally.

Consider how a decompositional analysis would have to isolate either a singular or a plural morpheme. Let us ignore the now familiar phenogrammatical challenges and focus on the tectogrammatical and semantic difficulties. Observe the interpretation given to the plural feature in Farkas & de Swart (*ibid.*), where *P indicates a number-neutral property:

- $\llbracket \text{PL} \rrbracket = \lambda x \lambda *P [x \in \text{SUM} \cup \text{AT} \ \& \ *P(x)]$

The first thing to note is the greater structural semantic complexity involved here. When composed with a quantificational determiner phrase, the search for a witness proceeds to find one that is a plural or an atom, and that fulfills the relevant property denoted by the head noun. This requires two lambda-abstractions, over the property and over the individual. In our case, the extra specification is carried out implicitly, as the restriction is built into the denotations without the membership clause needing to be stated. When composed with a quantificational determiner phrase the search for a witness may proceed as usual, with the restriction already included. This therefore involves no greater structural semantic complexity.

Furthermore, the property itself arguably remains number-neutral in our analysis, as the number information is only built into the domain restriction on the function, or property that the respective number forms denote, where this restriction emerges from the organisation of the nominal system. For the mouse distinction, there is some property **mouse**, that is denoted, but whose ultimate reference is always refined appropriately through the relevant domain restriction on that function depending on which form is in fact uttered. The number forms therefore have distinct, but systematically related meanings, by way of the property shared by their denotations. When the features are reified, the number forms themselves essentially mean the same thing, namely just the property, but they have different compositional import by way of the inclusion into the compositional procedure of the feature. A philosophical question thus arises as to whether, when we utter a plural, there is an individual compositional element with the above interpretation that we in fact mean by what we say. The analysis presented here does not force us to concede that there is.

There are two other problems treated in Farkas & de Swart (*ibid.*). Firstly, there is that of exclusive readings of plurals, where atoms are not possible referents. Secondly, there is the data presented regarding Hungarian cases of singular forms interpreted as denoting elements from *SUM*. A full treatment of these phenomena is not presented here - only a suggestion that the current account is

not necessarily incompatible with treatments they provide. One of the problems for the exclusive reading however, is that we either require *mice* to have two interpretations, both the one initially provided that includes atoms, and another that does not. The domain of the latter interpretation could be restricted appropriately such that the exclusive interpretation of *mice* is a function of type $\langle SUM, t \rangle$.

This means we would have to drop the assumption that one paradigm cell is associated with one meaning, however. Ultimately, this is no less parsimonious than assuming two competing plural feature interpretations as in Farkas & de Swart (*ibid.*), provided some equivalent competition scheme could be put forward. Moreover the notion that paradigms organise the meanings of their forms is essentially maintained, except that it is instead possibly sets of meanings of individual forms that are organised. Indeed, this is simply a case of polysemy, since the plural meanings are related, where the restriction on the domain for the exclusive reading is to a subset of that of the inclusive reading.

As for the Hungarian data, the domain restriction for singulars could be dropped, as is done in Farkas & de Swart (*ibid.*), while it is maintained only for the plural readings. This would then render the singular indeed maximally semantically unmarked, as is desired to correspond to its morphological unmarkedness. The typical atomic restriction could then be taken to result from an optimality-theoretic system of constraints similar to what is presented in Farkas & de Swart (*ibid.*). How this would be implemented given the current set-up is beyond the scope of this investigation, but the current semantics does not appear to be incompatible. If it is indeed possible, however, this would solve the problem of the current account not corroborating the markedness of plurals denotationally.

In the present analysis, the inflectional category of number organises the system of common nouns holistically; Singular and plural meanings (or sets of meanings) are properties of contrastive sets of inflectional forms, in such a way that no particular forms are privileged. The phenogrammatcs and semantics of morphology have thus been kept apart in a highly desirable manner. There is no difference semantically between the *mouse-mice* and the *dog-dogs* distinctions, as they both contrast the respective meanings in an equivalent way. There is a difference phenogrammatically of course, with the latter suffixal pattern being overwhelmingly more prevalent. This phenogrammatical distinction ought to play no role semantically, however, and it does not here. In a decompositional theory, either there are two compositional elements denoting the morphosyntactic property with their respective phenogrammatical characters, one concatenative and one not - each with the same meaning - which is unparsimonious; or the inflectional morphemes are rendered abstract, and meanings are no longer borne directly by concrete linguistic elements. The kind of separation afforded in the present account between the inflectional class distinctions and their shared meanings is not theoretically problematic in either way.

The example of number provides a starting point for the following discussion of inflectional categories whose semantic contribution is perhaps not so readily apparent, if one can even be concocted. The guiding principle is that contrastive

sets of forms marking the distinctions of a given inflectional category carve out the domain in some systematic way, providing correspondingly contrastive sets of meanings.

4.2 Agreement

This subsection treats three phenomena all involving agreement. Agreement takes many forms, as there are many possible features that come to be agreed with by the relevant different word-forms of a given syntactic category. The at times apparently non-local effects are of particular interest with respect to a compositional semantics.

The first two cases deal with the agreement of grammatical gender, where this is not typically taken to be something of semantic relevance. This assumption is challenged here in connection to the idea that the paradigmatic organisation of the morphological system also organises the meanings of different word-forms with respect to the morphosyntactic properties they bear - with grammatical gender being one of those properties.

We begin with the anaphoric reference of pronouns, taking examples from Dutch, where agreement with the grammatical gender of the antecedent referent is seen to aid in the resolution of that reference. This is then extended to cases of adjectival determiner phrases, as they are observed to fulfill a similar role. There are certain obstacles that the goal of semantically implementing grammatical gender must confront, namely those that are taken to suggest that it is not of semantic relevance. These also inform the discussion.

Lastly, verbal argument agreement provides us with an example that could easily fall prey to an analysis that suffers from much of the problems that have been taken here to be best avoided, given the criteria outlined in section 1. The analysis provided here avoids these problems, by suggesting that there is no tacit compositional element in pro-drop languages providing pronominal meaning directly, and that it is rather part of the verbal meaning in some way already. This may challenge any assumption concerning typological uniformity of the compositional process of interpretation. Following this, imperatives provide another case involving null subjects, inviting a discussion of the applicability of the preceding analysis to this related phenomenon.

4.2.1 Anaphoric Pronominal Reference

Jacobson (2012) offers an account of English personal pronouns that incorporates (non-grammatical) gender features. To relay it briefly, a variable-free semantics is employed, where assignment function are not involved. Instead, pronouns are interpreted as the identity function on the relevant individuals, where I denotes the identity function on the singleton that contains some special element in the domain denoting the speaker. *He* is defined on male individuals, *we* on the set of plural individuals of which the speaker is an atomic element, and so on.

This requires the domain of entities to be refined such that these salient groups of entities can indeed provide the domains of the identity functions. The rest of the semantics, with respect to the consequences for a compositional interpretation of pronouns of type $\langle e, e \rangle$, is dealt with by appropriate unary type-shifting rules that are not of concern for the time being. What is important is that we now have subsets F, M of D_e (or rather of AT , for the case of the singular pronouns) from which we can refine functional types, for example $\langle F, F \rangle$ which the interpretation of *she* and *her* fall into, similarly to what we did above for number distinctions on common nouns.

In languages with systems of grammatical gender, such as Dutch for example, third person pronouns refer not only to male and female individuals, but also to entities marked with a particular grammatical gender. Common and neuter are the only two grammatical genders of Dutch, and operate largely independently from the male-female distinction for individuals. *Hij*, means either “he”, when referring to some single male individual, or “it”, when referring to some individual denoted by an antecedent singular noun of the common gender (that is not male). *Zij* either means “she” and refers only to females, or it means “they”, regardless of animacy or gender.¹ *Het* refers to an individual denoted by an antecedent singular nouns of the neuter gender, and it may also be used impersonally.

The resolution of reference between antecedent common and neuter nouns may be determined in line with the contrast in grammatical gender between *hij* and *het*, in much the same way as the contrast between *he* and *she* can do so in English for biological sex, or the contrast between *hij* and *zij* also in Dutch. This indicates that there is indeed semantic relevance to this distinction that is, on the face of it, only grammatical, without a direct ontological basis, especially since it is somewhat arbitrary which nouns are of which gender. This latter point should perhaps not deter us given the semantic evidence of antecedent reference resolution.

Observe the following two simple Dutch sentences:

- (1) Ik genoot van het brood en de
 I enjoyed PREP the-NEUT bread-neut and the-COM
 pindakaas. **Het** was warm en **hij** was smeug.
 peanut-butter-COM. It-NEUT was warm and it-COM was smooth.

The referents of the pronouns can be automatically disambiguated by the gender contrast. Now, while this example is quite contrived - never mind that the disambiguation would indeed also be aided by the coherent ordering of the referents across the two sentences, and likely also some prosodic accentual distinction between the pronouns (as well as the descriptors of course) - the point remains that the gender contrast alone can serve to disambiguate the antecedent referents.

¹In more formal and archaic Dutch, *zij* can be used to refer to particular nouns of common gender as well, which used to be feminine before the previously distinct feminine and masculine genders were combined to form the common gender, but let us disregard this here. Also, just as with English “she”, *zij* can also be used to refer to ships and possibly other vehicles of some description that the (not improbably male) speaker happens to be particularly fond of...

One could even imagine a set of circumstances in which this disambiguation could only take place through the gender distinction in the used pronouns. Perhaps the interlocutors had had lunch earlier that day, and the hearer had baked the bread and made the peanut butter himself. Now, suppose the hearer had just asked the speaker what they thought of it and that when the speaker replies with the above, the end of the first sentence, namely the expressions referring to the bread and the peanut butter are not heard. In this case, the hearer can still understand the second sentence by inferring the referents of the pronouns from their gender (again, never mind that they could probably work it out from the properties ascribed to the pronouns' referents).

When it comes to the interpretations of the Dutch pronouns, let us conceive of them in a similar way to Jacobson (ibid.). The difference here is that instead of only carving up the domain into M and F for male and female (atomic) individuals, we add novel categories. Let us define C and N as the subsets of the domain of entities, for which the nouns that denote them are of the respective genders (and which are atomic). This may appear dubious on the face of it for several reasons that are to be addressed below, but let us continue first to outline the interpretations of the pronouns.

While in English we may follow Jacobson and interpret² $\llbracket he \rrbracket = \lambda x_{x \in M}.x$, we now have in Dutch that $\llbracket hij \rrbracket = \lambda x_{x \in (M \cup C)}.x$, the pronoun can refer not only to antecedent individual males, but to also the denotations of antecedent singular common nouns. $\llbracket het \rrbracket = \lambda x_{x \in N}.x$ can give us the appropriate reference for the denotations of antecedent neuter nouns, ignoring the impersonal usage. Let us also ignore *zij* for now, as it is not relevant to the distinction in grammatical gender, and that its plural interpretation complicates things further.

There are two potential problems that such an analysis would be confronted with. The first is extremely challenging and has been ignored in the preceding couple of paragraphs. It is namely that a common and neuter noun can of course denote the same entity. Take the example of *lamp*, of common gender, and *licht*, of neuter gender, meaning “lamp” and “light” respectively. These can be used to refer to the same object. This is perhaps what typically leads to the assumption that grammatical gender is of little to no semantic relevance, and understandably so. Here, what has been observed with respect to anaphoric reference resolution implores us to attempt to challenge this.

We are thus presented with two possibilities: either C and N have a non-empty intersection, or they both are empty to begin with and are updated by the context. We can dismiss this first possibility, because it may end up being that words such as *ding*, of common gender, meaning “thing”, and *voorwerp*, of neuter gender, meaning “object”, each cover most, if not all of the domain, resulting in C and N each covering the whole of the domain of entities in turn, and thus not helping to disambiguate anything. Let us therefore explore the second option.

We thus require an update functions that takes an discourse-initially empty

²Jacobson actually interprets this as the identity function of type $\langle M, M \rangle$, but this difference is minor, being mostly notational, and is glossed over here, as the restriction on the domain of the identity function is equivalent

C or N , and returns a C' and N' that includes the denotation of the noun of common or neuter gender that has been uttered. Let us define this update function. It is a two place function that first takes some noun P ; then, dependent on the gender of the noun —something conceived of as being linguistically encoded by the morphological system —it takes as its second element C and N and does the following:

- C.** If P is a common noun, then $N[P] = N$, while $C[P] = C \cup \llbracket P \rrbracket$
- N.** If P is a neuter noun, then $C[P] = C$, while $N[P] = N \cup \llbracket P \rrbracket$

Observe (2), where the second alternative anaphoric pronoun is ungrammatical given the co-reference indicated.

- (2) Ik vond een lamp_{*i*} in de kelder. Ik denk, dat [hi_{*j*}/
I found a lamp-COM in the basement. I think that [it-COM/
*het_{*i*}] kapot is.
*it-NEUT] broken is.
“I found a lamp in the basement. I think it is broken.”

Now we can see with our above pronoun interpretations, beginning with an empty C and N respectively, that we predict the co-reference in the first example, and the ungrammaticality of the second, given the update to C alone, in line with **C.**, following the utterance of some P , namely “lamp”. N remains empty, as no neuter noun has been uttered. By contrast, had “licht” been uttered instead of “lamp”, C would remain empty, and N would be updated, in line with **N.**, so the first alternative anaphoric pronoun would be rendered ungrammatical and the second grammatical.

Defining the pronouns with the appropriate restrictions on the assignment functions is what allowed us to achieve this. Importantly, the lexicon organises the genders of nouns, holistically, with two opposing sets of common and neuter nouns. This information is what is taken to guide the update procedure that provides the subdomains to which the anaphoric reference of pronouns is sensitive.

Questions remain about how long the updates to C and N would remain, and how they might be overwritten. Indeed, in some discourses, the same object may be referred to with words of different genders at different stages. How this might come to affect the update procedure is not treated in full here, other than to suggest that an overwriting process would likely be required. Another conceivable problem is that a common and neuter noun with overlapping denotations come to update C and N in some discourse, but they are each used to refer to distinct objects that both fall in the intersection. In this case, it could be more prudent to update each with the singleton of each referent, as opposed to each of the nouns’ full denotations. This would require some interaction with quantifiers to be worked out, which is saved for future research. This is invited given the fact that quantificational determiner phrases often inflect for gender, such as the distinction in Dutch between the common and neuter definite articles *de* and *het*.

The second problem concerns nouns which denote male or female individuals, but are of a conflicting grammatical gender. All diminutives in Dutch are of neuter gender, but may be used to refer to individuals of a given sex. *Meisje*, meaning “girl” is neuter, but is referred to with the pronoun *zij*. This imposes some sort of order of precedence on the sections carved out of the domain, namely that biological sex trumps grammatical gender when it comes to antecedent reference. This is not only the case for neuters, however; *hond* means “dog”, and is of common gender, but were it to be used in a sentence with the knowledge that the dog denoted is female, *zij* could also be used to refer back to it in subsequent sentences.

This is not directly a problem, however, given that the restrictions on domains of the identity functions can be sufficiently modified to account for these discrepancies. For example, *het* would have to exclude male and female individuals. This is simply achieved by intersecting N with the complement of $M \cup F$. *Hij* could also be adapted analogously to exclude just F .

Not only does this extra set of specifications gel with the intuition that the semantics should reflect that sex is somehow more salient than grammatical gender, as is corroborated by the linguistic data, but it is also arguably a result of the way in which the domain is carved up. Every common noun has a grammatical gender, so every member of the domain is denoted by a noun of one of these genders. A much smaller subsection of the domain, however, is of a certain sex. This means that there is a precedence inherent in the informativeness of the sections of the domain: sex limits alternatives more so than grammatical gender does. Such salience and informativeness correspond to conceptual and denotational markedness respectively. This is corroborated by the fact that the female possessive pronoun *haar* is distinct from *zijn*, which is the possessive pronoun for all the entities denoted by common and neuter nouns (as well as male entities), suggesting greater markedness of the sex distinction than that of grammatical gender in Dutch.

Returning to the relevance of this kind of interpretation to the examples that Jacobson treats, it appears as though the original analysis could be easily adapted to include analogous Dutch examples. The problem that Jacobson attempts to treat is exemplified in (3):

- (3) Only SUE called her mother

The problem that arises comes from the interpretation:

- **[[her mother]]** = $\lambda x_{x \in F}$ [the-mother-of(x)]

Here, that x is then taken to be the most contextually salient female. When composing this interpretation into the above sentence, however, it comes to be the case that only females other than Sue are interpreted to be non-mother-callers, due to the restriction on the domain. What we actually want, of course, is that everyone else is also a non-mother-caller. Jacobson suggests that the trick to achieve is that we want to make non-female non-mother-callers available by dropping the original refinement in the relevant environments.

Firstly it is noted that sentences in which these kinds of problems crop up given the variable-free interpretations display focus-sensitivity. Following one popular analysis of focus, where focus-marked elements raise a set of focal alternatives, the trick that Jacobson employs is to suggest that upon computation of these alternatives, the restriction on the functions is dropped, such as in the following interpretations:

- $\llbracket \text{she} \rrbracket = \lambda x. x \in Fx$
- $\llbracket \text{she} \rrbracket^{FOC} = \{\lambda x. x\}$ for all x in D_e

The reader is referred to the original text to observe how this solution came to be manifested compositionally in the above sentence. What is important is that any expression involving the original domain restriction loses this under the focus operator, which scales up compositionally to result in the desired effect.

Regardless of how stipulative the change in meaning in the above focus computation may be - something Jacobson acknowledges - there is nothing that would prevent this being adapted similarly for grammatical gender domain restrictions in the case of Dutch pronouns. The question is what the actual focus-sensitive constructions would look like in Dutch and whether they follow a similar pattern.

As it has already been mentioned, the possessive pronoun *zijn* is actually used for both common and neuter nouns, so it is no longer relevant. In German, however, which distinguishes three grammatical genders, masculine, feminine and neuter, the problem would apply, as the possessive pronoun for feminine nouns differs from that for masculine and neuter nouns (which are again the same). Therefore, where it would in fact be relevant in German, we could adapt the analysis given here to German, taking there to be three subdomains of entities, M' , which consists of male individuals and is updated by the utterance of masculine nouns, F' , which includes female individuals and is updated by feminine nouns, and N , which is just updated by neuter nouns. Observe (4):

- (4) Nur die **Flasche** verlor ihren Deckel
 Only the **bottle-FEM** lost its lid

Suppose we are in some contrived situation where there is one focal alternative that is a jar, which could have also lost its lid, until the speaker reassures us that only the bottle suffered this fate. *Krug* means “jar” in German and is of masculine gender. Suppose also that some sentence mentioning the jar had been uttered prior to the utterance of the above sentence, for example if the hearer had just asked whether the jar or the bottle had lost its lid.

Now the problem is that our initial interpretation for *ihren Deckel* is one that restricts the domain of lid-possessors to members of F' , which includes bottles (or just the contextually salient one if the update procedure is sufficiently modified):

- $\llbracket \text{ihren Deckel} \rrbracket = \lambda x. x \in F' \text{lid-of}(x)$

Due to the focal stress applied to *FLASCHE*, the focus operator is applied in turn to *ihren Deckel*, and the restriction is lost. As a result, the jar, an element of M' thereby becomes another lid-possessor, affording the appropriate reading that it in fact didn't, while the bottle did, when composed with the rest of the sentence. Again, this is no more stipulative than Jacobson's account. What is maintained is in a situation without the raising of focal alternatives, the restriction is maintained, giving the desired semantics for grammatical gender the rest of the time.

It is also worth mentioning that the phenomenon of 'paycheck pronouns' that was originally treated works just as well when extended to grammatical gender features. First observe the interpretation Jacobson provides of the following sentence:

- (5) Every third grade boy_{*i*} loves his_{*i*} mother. Every fourth grade boy hates her.

What is important is that *his mother*, is interpreted as a function of type $\langle e, F \rangle$, given that we know mothers are female (ignoring that it should really be of type $\langle M, F \rangle$ due to the contribution of *his*). Namely, $\llbracket \text{his mother} \rrbracket = \lambda x. \text{mother-of}(x)$. Glossing over the details, in the second sentence, the interpretation of *her*, namely the identity function over female individuals, is type-shifted such that it is the identity function over functions of type $\langle \langle e, F \rangle \langle e, F \rangle \rangle$. This is composed with the remainder of the sentence, where the function *hates* denotes has been appropriately type-shifted, to arrive at the following:

- $\lambda f_{\langle e, F \rangle} [\text{every-fourth-grade-boy } \mathbf{z}\text{-hates } f]$

This may therefore take any contextually salient function of the appropriate type, and the mother-of function is just that. The important point is that the restriction on the range of the function to females is naturally inherited by the pronominal expression in virtue of the type shifting operations and the resulting compositional procedure.

Analogising to a similar instance in German let us inspect the following (again contrived, but sufficient) sentence:

- (6) Jeder junge Mann_{*i*} verliert seine_{*i*} Flasche. Jeder alte Mann behält
Every young man_{*i*} loses his_{*i*} bottle-FEM. Every old man keeps
sie.
it-FEM.

Here, the utterance of the feminine noun *Flasche* adds bottles to the subdomain F' . We thus appropriately restrict the range of the bottle-of function such that it is of type $\langle e, F' \rangle$. Nonetheless, assuming an entirely analogous set of type-shifts, the feminine pronoun *sie* originally denoting the identity function over F' comes to denote the identity function over functions of type $\langle \langle e, F' \rangle \langle e, F' \rangle \rangle$. The rest composes similarly such that the interpretation of the final sentence is looking to be saturated by some contextually salient function of type $\langle e, F' \rangle$, namely the bottle-of function.

The difference here is that, while the inheritance of the gender feature can be carried out entirely analogously to the above for the pronominal expression, it is the utterance of *Flasche* in the previous sentence that guarantees that the range, F' , will not be empty. Besides this, the compositional interpretation proceeds equivalently.

This analysis has shown that there appears to be little difference with respect to the compositional machinery between interpreting biological gender features and grammatical gender features on pronouns in different languages. What is different, however, is the necessity of what is essentially a pragmatic update function responding to what is said, and carving up the domain accordingly. This is because it is not possible to hard-code grammatical gender, unlike biological sex, whose basis in reality appears to allow for a static refinement of domains and ranges of functions in its terms. What this analysis suggests about grammatical gender, if it is to be relevant to the semantics, is that it organises nouns with respect to their potential to dynamically refine portions of the domain such that referring expressions can be target them anaphorically in turn.

4.2.2 Adjective and Determiner Gender Agreement

Morphosyntactic agreement is in general presumed to be some sort of phenogrammatical epiphenomenon of no direct semantic relevance. This is particularly the case of adjectival agreement, where the proximity to the noun the adjectives modify and come to agree with suggests a level of redundancy, as word order provides a clear constituent structure according to which the interpretation can promptly proceed. The following example indicates that this assumption is at least in part misfounded. This is because it demonstrates that adjectival agreement, here with the gender (and number) of the noun it modifies, allows for anaphoric reference, with the noun absent, just as was observed with pronouns above.

It comes from French, which has two grammatical genders, masculine and feminine. First suppose a scenario in which someone is introducing their two dogs to an acquaintance. They point to them in turn and say their names, as follows:

- (7) Il s'appelle Hugo et elle s'appelle Cléo. Il a deux ans
 He himself.calls Hugo and she herself.calls Cléo. He has two years
 et elle a six ans.
 and she has six years.
 "He is called Hugo and she is called Cléo. He is two years old and she is
 six years old"

Some time later, the interlocutors are elsewhere, away from the dogs, and they spy a chewed up bone lying on the ground. The conversation returns to the topic of their dogs; the original speaker explains:

- (8) Le jeune adore des os, mais la vieille
 The-MASC young(-MASC) loves PART bones, but the-FEM old-FEM
 en a peur
 of.them has fear
 “the young one loves bones, but the old one is afraid of them.”
 [The gender associated to *jeune* is in brackets because it has the same
 form for both genders (see footnote 3)]
 [PART means the ‘partitive’ article]

Now, of course, we know that *the young one* refers to the Hugo, and *the old one* to Cléo even in English, despite the fact that it lacks agreement, simply because this information was specified earlier.

Suppose, however, that the addressee had been distracted by the presence of the dogs, and did not hear the second sentence of the first interchange (and also that they are unable to estimate the age of a dog from sight alone). They are therefore unaware of their respective ages. At this point, if the second interchange had occurred in English, the hearer would not know which dog likes bones and which is afraid of them. The references are resolved in French, however, simply because the genders of the dogs are known, and the adjectives and determiners, which are uttered without any linguistic element expressing the entity they modify directly, encode the gender information. The hearer therefore learns of the dogs’ respective ages and feelings towards bones.

This example should demonstrate that agreement can in fact impact an interpretation. The question therefore arises of how to provide an interpretation that permits the reference resolution. The technique outlined in the section immediately above is adapted for this purpose.

Let us demonstrate with the singular forms of the adjective meaning “old”, whose masculine form is *vieux*, and its feminine form, *vieille*.³ Taking adjectives to be of type $\langle e, t \rangle$, namely first-order predicates, we restrict their domains similarly to above. *Vieux* is therefore of type $\langle M, t \rangle$, while *vieille* is of type $\langle F, t \rangle$. Therefore, when *la vieille* is uttered in the above example, the hearer can infer which entity is in fact old and is afraid of bones, because the compositional process of interpretation involves the information that entity which is old (and is afraid of bones) is also feminine (or female), thus allowing the reference to be resolved appropriately. This is because there is a linguistic element encodes that information to begin with, namely the adjective (and the determiner).

This picking out of the entity is of course carried out in combination with the relevant form of the determiner, which also agrees for gender. We could restrict the domain of the assignment function denoted by the quantifier analogously, or just add an extra clause stipulating that the bound variable be in F or M appropriately. This involves a typical interpretation of determiners that may

³ *Vieux* is also the masculine plural form, and when the subsequent noun it modifies begins with a vowel, it takes the form *vieil*, which is identical to the feminine phonologically speaking. The number and gender in this cases remains distinguished by the determiner that precedes the adjective, however, so we may gloss over these technicalities for the time being. This also applies to forms like *jeune*, where there is one syncretic form for both genders

conflict with the original goal of variable-free semantics due to the inclusion of assignment functions, although it is imagined that some work-around exists where some equivalent restriction can also be implemented. This is not to be explored here, however.

In the previous example, the entities are actually male and female, instead of simply being picked out by masculine and feminine adjectives. The point remains, however, that the adjectives have referred anaphorically. It is now pertinent to demonstrate that a similar example can be given that makes use of only grammatical gender. Observe the following (perhaps somewhat contrived) example from German:

- (9) Ich nehme zwei Bücher mit in den Urlaub, einen Reiseführer
 I take two books with in the holiday, a-MASC guidebook-MASC
 und ein Lehrbuch.
 and a-NEUT textbook-NEUT.

Das dicke bringe ich mit, um weiter zu studieren
 The-NEUT thick-NEUT bring I with, in.order further to study
 und den dünnen, um die Stadt gut kennen zu lernen.
 and the-MASC thin-MASC, in.order the city good know to learn.

“I’m taking two books on holiday, a guidebook_i and a textbook_j. I’m bringing the thick one_j to continue studying and the thin one_i in order to get to know the town.”

Now, the latter co-indexed phrases refer to the different books mentioned in the first sentence because the adjective and the whole determiner phrase it comes to be a part of bear the appropriate gender marking (ignore the likely common knowledge about the respective lengths and uses of different kinds of books). The mentioning of the different (kinds of) books added their denotation, or just their referents to the appropriate sub-domain; guidebooks, denoted by a masculine noun in German (or just the particular guidebook) were added to M' and textbooks (or just the particular textbook), denoted by a neuter noun in German, were added to N , such that they could be referred back to in the second sentence. Again, if it ultimately proves more useful, we could instead add just the individual books picked out by the entire determiner phrases including the nouns, but let us not consider this for now, as what is done here is sufficient for this example.

Returning to the main theoretical point, in both cases, if we were to try to isolate a linguistic element that encodes this information itself, two problems arise. The first is the familiar phenogrammatical quandary of non-concatenative, or even simply irregular morphology. The differences between the respective masculine and feminine forms of the French adjectives *beau* and *belle*, roughly meaning “nice”, *bref* and *brève*, meaning “brief”, should provide sufficient reminders of the likely pitfalls.

The second relates to the question of the extent to which we choose to tolerate the reification of grammatical gender. The update procedure by which

the relevant subdomains are configured suggests that this is not a true reification as in the case of plurals - it is something that happens throughout a discourse. What this means with respect to the organising capacity of the lexicon with respect to this inflectional category, is that the classification of nouns with respect to gender is not *a priori* considered a part of the system, which we thankfully avoid. Instead, it is something that can aid communication in some way, by restricting the domains of functions to discourse-relevant entities, in a systematic manner.

This discourse-relevant benefit arguably reflects a view of grammatical gender proposed in Dye, Milin, Futrell & Ramsar (2017), where the distinctions marked, for example, on determiners and adjectives reduce the space of possibilities for the coming nouns to be uttered. Once *le vieux* is uttered in French, or just *le*, there are only half as many nouns that the hearer may expect to follow. This conceivably aids processing, and provides a functional role for grammatical gender distinctions. This role appears to be reflected to a certain extent in the non-linear compositional semantics deriving from the above account, particularly across successive stages of a discourse where the computation of anaphoric reference is made significantly less complex given the domain restrictions.

The effects are conceivably also more local, however, although it may hard to conceive of an order of interpretation in the same way as the ordering of words in the cognitive processing of the speech stream. Assuming the quantificational interpretation of a determiner phrase searches for a witness to the properties involved in the restrictor, then given that the main property is denoted by a common noun of a particular gender which comes to be processed first because the noun phrase is computed before being composed into the determiner phrase, that search space is reduced once the noun is uttered.

Consider the English phrase *an actress*, where the female form of the word, contrasting with the male *actor*, conceivably reduces the search space to the contextually salient actress amongst all females, instead of all individuals (although this specific example is asymmetric unlike for the adjectives considered here). The functions denoted by the appropriately inflected adjective forms also come to have their domains limited by whatever noun of a given gender that acts as their syntactic head. This conceivably reduces their informational load of processing as well. This is of course highly speculative, but it does appear to follow given the analysis presented here.

The update procedure would likely need to be further refined to cash out these results such that the hierarchical interpretation procedure is appropriately interdependent with the linear processing of speech stream. It is hoped and assumed this could be achieved in further research, but it is beyond the limited goals of the outline provided here that aims only to suggest that some such method is conceivable and may ultimately be beneficial.

This kind of functional explanation is important because it challenges the assumption that morphosyntactic agreement is semantically redundant. The case for redundancy is arguably strongest when we have a whole determiner phrase where the non-noun elements in it agree, despite being right next to the noun. The ability for agreement to disambiguate anaphoric reference in

such a non-local manner already demonstrates part of its semantic import. To ground the functional explanation with respect to an according compositional interpretation procedure would provide further means to refute the otherwise quite plausible case for semantic redundancy. This is all the more true if it concerns grammatical gender agreement, whose semantic import is on the face of it highly dubious.

What is important theoretically here with respect to the distinctions between word-forms that display different kinds of morphosyntactic agreement is that each adjective and determiner form is taken to inhabit a particular paradigm cell, which determines its interpretation with respect to the relevant restriction on its domain that is paradigmatically encoded by that form. This is a product of the holistic understanding of the morphological system adopted here, where the set of all forms of a given syntactic category that bear a certain set of morphosyntactic features, for example masculine singular, share in the same restriction to their domains.

4.2.3 Null Subjects

The final example of morphosyntactic agreement involves verbal argument agreement, where verbs take a certain form depending on the person and number of their arguments. What is interesting with respect to a compositional semantics is that languages with considerable verbal argument agreement typically allow for the pronouns (or nouns) with which the verbs would agree to be dropped, a phenomenon known as pro-drop.

As was demonstrated with the example from Greenlandic in section 2.2.1, this can be both subject and object agreement simultaneously. For simplicity's sake, the following (perhaps somewhat contrived) analysis comes from Spanish, where only subjects are agreed with. *Como el queso* means “I am eating the cheese”, where *como* alone somehow comes to mean “I am eating”. This is contrasted with *yo como el queso*, meaning the same, but with the pronoun *yo* present, meaning the same as “I”.

Now, a typical analysis may assume that in the case of pro-drop, there is some ‘null’ pronoun, equivalent to *yo* that bears the first person singular meaning and acts as the verbal subject, despite it not being articulated. To assert that a phonologically unrealised element is the bearer of this meaning appears dubious, and this additionally renders the inflection on the verb semantically redundant. We will not consider this option any further.

Another common analysis would suggest that the suffix carries the first person singular meaning individually, providing the verbal argument. The suffix is borne by the verb, which would seem to invite an analysis that composes them first, and then the fully inflected verb with the object. However, when it comes to the interpretation, given that the object is the external argument of the verb, and the subject the internal argument, as is depicted below, there is a clear, and perhaps undesirable tension between the phenogrammatics and the tectogrammatics. This kind of decompositional analysis is of course something we would like to avoid in the first place anyway, so let us now consider the whole-word

option, and see if these problems can be avoided.

We would like it to be the case that the information about what the verb takes as argument is not something that is given by a distinct compositional element to the verb, which itself would thus bear only the lexical meaning in turn. What this means is that it should somehow already be part of the meaning of the verb that it takes the speaker as its argument. As before, we restrict the domain of the verb for each of its inflections to the appropriate members of the domain.

To make the example even simpler, let us consider the interpretation of the intransitive verb *duermo* meaning “I sleep”. We characterise the type of the function it denotes as $\langle SP, t \rangle$, where *SP* is a subtype corresponding to the domain consisting of the singleton containing the special element that is the speaker $\{SP\}$. Now suppose *duermo* is uttered. There is no longer any compositional element that itself means the speaker. We therefore have to assume that the argument is saturated by the hearer given the context of the utterance. When the pronoun is added, such as in the sentence *yo duermo*, also meaning “I sleep”, the domain restriction is rendered redundant, as the argument is provided by a distinct compositional element.

This analysis is perhaps less satisfying for a third person form, which - ignoring number distinctions - can take as argument any member of the domain (typically besides the speaker and the hearer). This is arguably no worse on the whole than assuming that the argument is present as some extra step in the process of composition, however. Indeed, we assume that the reference of, for example, “it” can be resolved efficiently by English speakers in the given context, taking it to denote simply the identity function over individuals that are not known either male nor female, and typically are also neither the speaker nor hearer, to follow Jacobson. It thus seems no greater a leap in faith to believe that speakers can also saturate verbal argument slots equivalently off the cuff with some contextually salient referent.

What is important here is that *duermo*, as the form that inhabits the first person singular paradigm cell of the lexeme meaning, involves this particular domain restriction, while the other paradigm cells have their own appropriate ones. This is a feature of the verbal system of the entire language, where every verb’s paradigm organises the domain restrictions on the meanings of its forms in an equivalent manner.

What is interesting about this analysis is how it attempts to provide a solution to the problem of pro-drop. By removing a step in the procedure of compositional interpretation, one can avoid the postulation of null elements, dubious phenogramatics, and diminished compositional transparency. The information that would have been encoded by the extra compositional element becomes part of the meaning of the verb itself, in virtue of its paradigmatic relation to other like forms.

Typologically, this leads to an interesting prediction that there is variation in the compositional process of interpretation, where languages with richer verbal morphology require possibly fewer steps to arrive at the same meaning compared to their morphologically impoverished counterparts. This could suggest varia-

tion in processing. Alternatively, it hints at some kind of interplay and perhaps trade-off between the (enumerative) complexity of the verbal system and the compositional process of interpretation. See section 4.5 for further comments regarding typology.

One might be led by such the preceding analysis to suppose that an analogous treatment could exist for another phenomenon that routinely involves null subjects, namely that of imperatives. Imperative null subjects are typically readily interpreted as the addressee, as in (10), or some group of addressees, as in (11), potentially including the speaker, as in (12), or they may quantify over some such group, as in (13). This is corroborated by the fact that in many languages, imperative morphology resembles that of second person indicative verb forms, although such syncretism need not imply anything more than a shared diachronic origin, with their interpretations being observably distinct synchronically. Imperative sentences carry a ‘directive’ meaning, where the interpretation usually assumed is that the speaker is directing the addressee (or addressees, or some of the addressees) to make the proposition denoted by what is said true.

- (10) Make your bed!
- (11) Play a duet!
- (12) Let’s go!
- (13) Someone fetch me a dishcloth!

In terms of inflectional categories, however, the imperative is considered a mood, and may therefore be considered more semantically (and pragmatically) loaded than for the indicative verbs in the pro-drop examples above. This is suggested by the directive meaning of imperative clauses, and is corroborated by the associated deviation in syntactic structure - something that is similarly observed for other semantically and pragmatically loaded clause types such as interrogatives. This indicates that the above treatment of verbal argument agreement, in terms of the inflectional categories of person and number, might be insufficient to cash out the semantics appropriately, because the taking of the addressee as the null subject argument for imperative verbs, must work differently in some way. Indeed, this is corroborated by (14), a sentence in Bhojpuri, an Indian language, taken from Zanuttini et al. (2012):

- (14) tebulwa: sa:ph rahe!
 table-NOM clean-NOM be-IMP.3SG
 “make it the case that the table is clean!” (lit. “the table be clean!”)

What we see here is that the imperative form of the verb, with directive meaning, telling the addressee to do something, in fact displays agreement in person and number with the syntactic subject of the sentence, namely the table. There appear to be two arguments, somehow, as is clearer in the (non-literal) translation provided. This indicates that a more nuanced semantic analysis is required, which will not be attempted directly here. The point of this example is to demonstrate that, while null subject imperatives may appear analogous to

cases of pro-drop, the phenomena may come apart in languages with sufficiently complex morphological systems, like that of Bhojpuri.

Putting the directive part of the verb meaning in (14) to one side momentarily, and focussing on the declarative part of its meaning instead, the person and number agreement on the verb could be treated just as for typical verbal arguments as above, where the interpretation of the copular form inflected for the morphosyntactic properties 3SG is appropriately excluded from taking the speaker, addressee, and any plural individuals as its argument. This argument specification is then made redundant when composed with the syntactic subject, namely the table, but that may not be so elsewhere.

The question remains as to how to factor in the directive meaning of the imperative mood of the verb. It is suggested in Zanuttini et al. (2012) that a second person feature elicited by the imperative is inserted into the syntax in a different, special way, compared to the ordinary way in which the third person singular feature is, resulting in the verbal agreement. In a language like English, the former would block the latter, but in Bhojpuri, this is not the case. The special way in which the second person feature elicited by the imperative is inserted would then allow for the meaning of the entire clause to be cashed out appropriately, while leaving the verbal-subject agreement intact, and not modifying that subject meaning.

This sketch of a treatment in Zanuttini et al. (*ibid.*) gives us an idea of what would have to be incorporated into the word-based analysis. Whatever the ultimate semantics, there needs to be an added aspect of the imperative form's meaning that includes an argument slot for the addressee as the special subject, while somehow leaving the copular verbal-subject agreement unscathed, and all the while incorporating the in some way modal effect of the directive meaning.

There is a conceivable advantage that would result the application of a word-based theory that captures all of those aspects of the meaning as part of the verb's meaning in one fell swoop. The special insertion argued for in Zanuttini et al. (*ibid.*) seems to imply some sort of non-locality at the syntax-semantics interface, where the two agreement features are inserted at different points in the syntactic structure. This would likely be required for the directive, modal aspect of the meaning involving the addressee to take wide scope over the declarative meaning of the proposition to be made true. Non-locality is undesirable, however, as a violation of (iii) (see section 1), and also therefore of the context-freeness of the interpretation procedure.

If it is possible to give an interpretation to the imperative form, such that the (underdetermined) directive meaning is part of the verb to begin with, and the addressee is automatically saturated as the imperative subject from the context, similarly to what was done for the preceding analysis of pro-drop, albeit with greater complexity, then the rest of the sentence could hopefully be composed to achieve the desired meaning. It may nonetheless be difficult to type the meanings appropriately such that they can be composed, but it is hoped this would be achievable in such a way that the non-locality implicit in the original analysis no longer presents an issue.

4.3 Recalcitrant Case

Grammatical case poses a considerably greater challenge to the attempt made here to give some kind of interpretation to words bearing certain morphosyntactic properties. There is typically a split drawn between ‘structural’ and ‘semantic’ case, where the former is taken only to be involved in the establishment of syntactic relations, while the latter may carry some meaning. For example, in languages with nominative-accusative alignment, these both being structural cases, for transitive verbs, the subject is typically marked with the nominative, and the object with the accusative. On the other hand, a language such as Finnish, with a highly extensive case system, certain ‘semantic’ cases may come to fill the roles that adpositions (pre- and postpositions) do in other languages, such as the ‘comitative’ case marking accompaniment, the ‘inessive’, being inside of, the ‘adessive’, being at, or in the presence of, and so on.

The distinction is not always so strict, with cases being co-opted to perform specific, perhaps unexpected functions in certain contexts. For example, in certain Slavic languages, it is common for copular expressions to involve the instrumental. In the Polish sentence *jestem mężczyzną_{inst.}*, meaning “I am a man”, the word for “man” is in the instrumental case, which, elsewhere, typically marks an object that the subject has used in some way to achieve something, often translated into English using the prepositions *by* or *with*. This is only the tip of the iceberg when it comes to quirky phenomena involving case, but they go beyond the limited goal here which is to hint at a possible role for case morphology in interpretation.

While the distinctions between semantic-case forms are conceivably little more difficult to treat compositionally than prepositional phrases, structural case presents a much more challenging phenomenon, whose semantic import is not clear from the outset. We see this reflected in the phenomenon of case agreement, where it is never involved in helping to determine reference, unlike other the morphosyntactic properties treated in the previous two sections. Rather, case appears to simply determine argument relations. Take the following (again somewhat contrived) German sentences as an example:

- (15) Ein alter Mann fing einen kleinen Fisch. [...]
 An-NOM old-NOM man-NOM caught a-ACC small-ACC fish-ACC. [...]
 Ulf schlug den alten wirklich hart.
 Ulf hit the-ACC old-ACC really hard.
 “An old man caught a small fish. [...] Ulf hit the old man really hard.”

The reference of the old man is one and the same in both sentences, despite the fact that the expressions denoting him bear different cases. The adjectives also agree for gender, which is what licenses the absence of the noun in the second sentence. Moreover, *den alten*, referring to the old man in the second sentence, and the phrase denoting a fish in the first sentence are both in the accusative case, but this entails no semantic connection. We can thus safely dismiss the notion that these purely grammatical cases are involved in determining refer-

ence. What we cannot yet dismiss, however, is that they are not relevant to the semantics in any sense.

One potential role for case follows from Dowty (2007). The **Categorial Theory of Argument Structure** proposes that an analysis of the argument structure of multi-place predicates follows from the order in which the carried functions they are interpreted take their arguments. This provides definitions of a (grammatical) subject, direct and oblique objects, and other complements (PPs, AdjPs, infinitival complements, etc.): the subject is last to be taken, direct objects are taken penultimately, and so on. This results in correct predictions about morphosyntactic case phenomena that reflect syntactic agreement and government.

The question for us is, if the morphological system organises meaningful distinctions amongst related forms, what semantic role could case then play? It appears as though case helps to determine the order in which many-place predicates take their arguments. The first thing to point out is that this is not a purely syntactic phenomenon; the event of a dog biting a man is of course different to that of a man biting a dog, and in many languages, the relevant linguistic expressions distinguishing these events employ different case forms on the respective nouns to achieve this ahead of —and in certain languages entirely instead of —word order.

Dowty relegates both morphosyntactic properties and word order to phenomenatics, being without direct relevance to the compositional process of interpretation, beyond perhaps encoding the appropriate parse of the tectogrammatical structure from which meanings are ultimately computed. There is, however, a notable difference between the resolution of the order in which a many placed predicate takes its arguments through word order, and through case marking. In the former case, it is indeed highly plausible that the word order encodes the tectogrammatical structure according to which the interpretation proceeds, because of the locality of interpretation. This accords with the statistical findings of Futrell, Mahowald & Gibson (2015), where languages tend to display what is termed ‘dependency length minimisation’, which is tantamount to a limitation on non-local constituent structure.

Case marking, along with other morphosyntactic markings, often allows for this trend to be bucked, where the assumption is that this is because dependencies are explicitly marked, so the constraint of locality is no longer as relevant. Indeed, the seminal example of a non-context-free linguistic expression comes from a phrase involving cross-serial dependencies in Swiss German, the resolution of which can be determined through case markings (see Shieber, 1985).⁴ In line with Dowty, the relevance of this finding is not typically taken to go beyond the establishment of syntactic dependency, if case is only to be of indirect relevance to a process of interpretation.

⁴Another notable example of cross-serial dependencies comes from Dutch (see Hugbrechts, 1976), a language that has for the most part lost case distinctions besides in a number of fossilised phrases. The recency with which the cases have been lost in Dutch, as well as certain other features of its word-order, perhaps indicate that some syntactic sensitivity to dependencies that would have been afforded by the overt marking of case distinctions remains.

The perhaps overly ambitious and speculative goal here is to gesture towards some account of how different (structural) case forms could in some way be of different semantic import. Clearly, we have to extend our notion of meaning to some extent in order to at least have a go. The resulting question is thus: how might we make an attempt to implement case distinctions directly into the semantics, determining the order in which many-place predicates take their arguments independently from an intervening tectogrammatical parse? Again, we wish to avoid the case distinctions being solely relevant to determining an intervening parse.

Observe the following German sentences, both meaning “the woman loves the man”:

- (16) Die Frau liebt den Mann.
The-NOM woman-NOM loves the-ACC man-ACC.
- (17) Den Mann liebt die Frau.
The-ACC man-ACC loves the-NOM woman-NOM.

Let us take *liebt* to be a function of type $\langle\langle e, t \rangle, t \rangle, \langle e, t \rangle$. Now, with DPs being of type $\langle e, t \rangle, t$, we cannot specify that verbal predicate is applied to the denotations of the nouns in a certain order, as the subject in the nominative case *die Frau* is in fact a function that takes as argument the predicate, a function of type $\langle e, t \rangle$, the VP with its direct object already applied. Given such a semantics, this will always be the case for the final argument taken by any predicate.

We have to abstract from the asymmetric process of individual applications of verbal arguments, therefore, since we are unable to say the nominative is the verb’s internal argument, as it in fact takes the interpretation of the rest of the verb phrase of the sentence as arguments. There is nonetheless an order imposed on what the order of functional applications in the process of interpretation is.

Any attempt to encode some sort of linear ordering in the compositional process of interpretation gives a new sense to ‘mode of combination’ as it is typically understood in the principle of compositionality. Instead of a hierarchical syntactic structure, derived from constituencies implicit in the sentence meaning with its original word order, case determines a linear mode of combination, directly reflecting the currying of multi-place predicates.

On the one hand, to somehow make this a part of our semantics would be unparsimonious, as within the constituents that are taken as arguments by multi-place predicates, hierarchical syntactic structure remains. The linear mode of combination would thus be relevant only to the verbal syntactic domain, or possibly other domains where the relevant syntactic element received an interpretation as a curried function. For these reasons, we may wish to abandon such an attempt.

On the other hand, suppose we can ground the phenomenon directly semantically, where this mode of combination avoids such a linear parse, distinct somehow to the rest of the tectogrammatcs. Ideally, it would fall out of some kind of distinction between the interpretations of case forms, perhaps in relation

to multi-place predicates. It could then be significantly less unparsimonious and stipulative.

Now, we take it that determiner (or noun) phrases pick out individuals in some way. This involves quantification in the semantics of course. In a language with case marking, each such phrase is marked for case. Suppose that what case does is to index the introduction of individuals into the process of interpretation, by way of the quantifiers, such that a sequence of those individuals is attained.

Let us then suppose that Dowty's theory of argument structure constrains the point at which an individual may be added to such sequence. Conceiving of the sequence as a stack, an individual referred to by a nominative phrase may be added first, followed by that referred to by an accusative, that is placed on top of the first individual, so to speak, and so on for the different case forms. This stack then determines the order in which the individuals are taken from the queue and taken as arguments by multi-place predicates, in correspondence with the order in which they had initially been carried out.

Any attempt to implement this process compositionally in a type logic would no doubt be difficult. The reader is referred to Dekker's (2012) system of Predicate Logic with Anaphora, where a loosely related dynamic treatment of discourse anaphora employs a similar idea, involving a sequence of individuals that impacts the satisfaction of formulae involving the relevant existential operators. For case, however, the putative sequence would be constructed clause-internally, and the interpretation is therefore not only not contextual, as it does not relate to anaphoric reference, but it is also unlikely to involve logical connectives in any typical manner, due to the locality constraints on case, such that its specification seems always to lie within the domain of a single verbal predicate.⁵

A potentially more helpful extension to our semantics would be to include events, harking back to Davidson (*e.g.* 1967). In such a theory of interpretation, sentences introduce not only individuals such as humans into discourse, but also events. A clear motivation for this is to have a referent for the anaphoric pronoun in the following sentences: "We played football yesterday. It was fun.". Additionally to introducing the event, certain special relations may be introduced into the semantics, corresponding to what are often termed thematic (theta-) roles, through which canonical kinds of relations to the event are encoded; following the Neo-Davidsonian approach, there may be agents, who carry out an action, patients, on whom an action is carried out, experiencers, instruments,

⁵There are cases of one predicate taking more than one co-ordinated objects, and more than one co-ordinated predicates taking one object, such as in the following sentences:

- (i) I bolted down the stairs, my breakfast, and the door.
- (ii) I washed, dried and ironed my dress shirts.

In either case, in a language that marks case distinctions, the object (or objects) would be in a particular case —often the accusative for direct objects —and that would determine when it (or they) argument by the verbal predicates (or predicate resp.). This co-ordination is not operating at the level of clausal co-ordination, but rather on smaller constituents, which suggests indeed that it is not equivalent to logical conjunction operating on truth values, and as such that it is more compositionally involved.

locations, and so on. These relations are exactly what appear to be determined by case markings, both structural and semantic.

Let us begin simply with the above sentence, meaning “a woman loves a man”, where we give its translation into a logical language of event semantics as:

- $\exists e \exists x \exists y (L(e) \wedge W(x) \wedge M(y) \wedge \text{AG}(e, x) \wedge \text{PT}(e, y))$

where AG and PT denote the special relations to the event of agents and patients respectively. To avoid the complications of determiners, let us consider the equivalent sentence in Polish, where (18) represents just one of the possible word orders:

- (18) kobieta kocha mężczyznę.
 woman-NOM loves man-ACC.

We give the following interpretations, ignoring tense and other currently irrelevant factors:

- $\llbracket \text{kocha} \rrbracket = \lambda e_v. L(e)$
- $\llbracket \text{kobieta} \rrbracket = \lambda R_{\langle v, t \rangle} \lambda e. \exists x (W(x) \wedge R(e) \wedge \text{AG}(e, x))$
- $\llbracket \text{mężczyznę} \rrbracket = \lambda R \lambda e. \exists x (M(x) \wedge R(e) \wedge \text{PT}(e, x))$

In such a way, the composition can proceed to yield the interpretation of the full sentence, with the event argument itself ultimately saturated in the given context. What is important here is that the paradigm of a noun (or determiner, adjective,...) can then organise the association of roles to specific case forms.

There are problems that such an analysis encounters. Firstly, there is the problem of verbs that take non-standard cases. The above example of the Polish copula taking the instrumental case is one amongst many. In German, some verbs such take direct objects in the dative case, such as *helfen* meaning “to help”, or *glauben*, meaning “to believe”, while a few take their objects in the genitive case, such as *gedenken*, meaning “to commemorate”, and *bedürfen*, meaning “to require”.

If the theta-roles are hard-coded in the interpretations of case forms, then this leads to problem, as, if we expect these direct objects to be patients semantically, and the dative and genitive do not carry that theta-role as part of their meaning, then the sentences involving such verbs will receive incorrect interpretations. It may be that a more refined account of the event semantics of these verbs reveals that the theta-roles of their syntactic objects are not ‘patient’. Otherwise, there would have to be separate dative and genitive meanings stored, where the case forms are of the appropriate theta-role, and then it is only in combination with these particular verbs that those meanings are selected. Another option could be to let the specific verbs modulate the meanings of the non-standard case forms, but this is non-local and therefore arguably undesirable.

Another interesting point is that this provides conceivably just as great of a challenge, if cases are considered to constitute individually meaningful features, as if cases are taken only to organise the meanings of different whole case-forms. An accusative case feature would just as reasonably be considered to carry the meaning of ‘patient’ as any accusative case form does, along with the specific noun (or determiner or adjective) that is inflected. Furthermore, an analogous problem exists for verbs that use specific prepositions, such as *listen to*, *look at*, *consist of* and so on, where the prepositional meaning is lost and the preposition serves only to mark what semantically appears to be a patient, in much the same way as for an otherwise dative or genitive meaning.

What is perhaps interesting is that some analysis in terms of the earlier hierarchy of cases providing a sequence of individuals could work better here to treat these idiosyncratic case usages, where, genitive and dative forms, for example, still provide arguments that are taken by their verbs ahead of the nominative subjects. Further development of that idea would be required first to explore any specific potential benefits when applied to these examples, however.

We now see that case, despite not altering the reference of distinct case forms, may still be considered to be of some compositional semantic import. This is either (or both) through the ordering of arguments to be taken by multi-place predicates, or (and) through the association of thematic roles to noun or determiner phrases in an event semantics. This is another sense in which a morphology can be seen to interact with the interpretative process. Importantly, this interaction is taken here to take place through the way in which the structure of the lexicon organises the basically meaningful elements that provide inputs to an interpretation.

4.4 Typology and Diachrony

As has been alluded to throughout, broader implications would seem to fall out of the theoretical stance adopted here that are relevant to the (morphological) variation of languages, be that synchronic, through typological variation amongst different languages, or diachronic, through systematic change in one particular language over time. Let us first inspect the more constrained case of the latter, before leading onto a discussion of the scope of possibilities involved in the former.

Consider the case of diachronic morphological simplification. We may characterise this theoretically as the loss of morphological markings that partially or completely impedes the expression morphosyntactic properties through the previous systematic contrasts amongst word-forms in and across the paradigms of a given inflectional subsystem (such as that of nouns, or verbs). This entails a collapse of part of the paradigmatic structure of the lexicon. Now, these trajectories of change are extremely complex and may only be thoroughly analysed in light of the data collected and processed by philologists. There are clear overarching facts, however, such as the loss of the case system in English when compared to its forebears. Let us continue to spell out the general theoretical

point.

While morphological simplification is commonly attested, it is worth mentioning change may proceed in the other direction as well. This is evidenced by instances of morphological complexification, through processes of grammaticalisation that bind a series of compositional elements into one unit over time, which may also be termed ‘morphologisation’. To provide an example, Habicht (2000) documents the development of the modern Estonian comitative case suffix *-ga* from the previous adposition *kaas*. The noun system of Estonian was thereby made more complex, corresponding to an increase in the number of cells of its paradigms.

Returning to simplification, if the paradigmatic structure of the lexicon organises the aspects of semantic import shared across the simplex interpretations of related whole word-forms, then the collapse of parts of or whole paradigms results in the relevant meaningful distinctions being lost, and word meanings may change as a result, likely being generalised in some way. Consider the present tense indicative sub-paradigm of the Spanish *comer*, meaning ‘to eat’, that was referenced in section 4.2.3, as displaying pro-drop. If the morphotactic distinctions that discriminate the different person and number forms morphosyntactically were to be lost in some future state of the language, then the encodings of the systematic restrictions on the arguments that we take to provide the semantics of these distinctions would no longer be linguistically present, and it is likely that they could therefore no longer feature an interpretation.

Now compare a language like French, whose providence is shared with Spanish, where former person and number distinctions on present indicative forms have been levelled to a significant degree.⁶ Pronouns are obligatory in French, unlike Spanish. The aim of this comparison is not to suggest that in a language without verbal person and number inflections, pronouns are always going to be obligatory, as this is simply not the case, considering many languages, such as Japanese, regularly drop pronouns despite lacking the relevant inflections; rather, it is to suggest that given the particular diachrony of French, where it was at one point in a state analogous to Spanish, the levelling of the relevant person and number distinctions in verbal morphology could have been a factor in prohibiting pro-drop, because the capacity for verbs to restrict the verbal subject arguments inherently was lost, and so too the licensing of pro-drop.

There are presumably other available explanations of this change, that do not depend on the semantic WP theoretical position outlined here; and, in light of specific philological data, perhaps such an explanation could be ruled out. Such an explanation may nonetheless demonstrate a broader point about the holistic conception of morphology adopted here, namely that, unlike in syntactocentric and more generally constructive morphological theories, where there

⁶The levelling is most apparent in phonological environments where *liaisons* revealing otherwise ‘covert’ inflections are not present, namely not before vowels. For example, the second and third person singular present tense indicative forms *es* and *est* of the verb *être*, meaning ‘to be’, are phonologically identical, pronounced as the vowel /e/, unless they precede a word beginning with a vowel, in which case the final consonant is pronounced, yielding /ez/ and /et/ respectively

is presumed to be some compositional process of interpretation that applies to morphologically complex word-forms, an abstractive word-based account of the interaction between morphology and semantics instead assumes variation in the process of interpretation of languages that differ in morphological complexity. This is a significant prediction (/assumption).

Moving onto typology, the juxtaposition of typological extremes of morphological complexity, from a maximally analytic, or isolating language at the compositional pole, to ‘polysynthetic’ languages at the holistic end, may serve to best illustrate this point. These terms derive Sapir’s (1921) typology, which is perhaps somewhat dated, but may still be used to demonstrate the point efficiently. To elaborate on ‘polysynthesis’ briefly, let us recall the opposition between analytic and synthetic morphology, where, to quote Sapir, the former “does not combine concepts into single words at all (Chinese) or does so economically (English, French)”, while in the latter “the concepts cluster more thickly”.

Polysynthetic languages are extraordinarily synthetic, to the extent that, following the definition provided in Fortescue, Mithun & Evans, (2014), they display ‘holophrasis’. This is where a single highly complex word-form is sufficient to express the meaning of a whole sentence in other morphologically simpler language, in particular through “including all bound core pronominals”. They must also be able to allow more than one lexically ‘heavy’ morpheme in the holophrastic verb, where these might have meanings that are individual words in other languages, such as “judging”, “wishing for”, “lacking X”, and so on. An example of a highly complex holophrastic word-sentence in West Greenlandic is (19), taken from Fortescue (1984), where INTR.PART means ‘intransitive participle’, and IND means ‘indicative’.

- (19) Nannu-n-niuti-kkuminaR-tu-Rujussu-u-vuq.
 polar.bear-catch-means.for-good.for-INTR.PART-enormous-be-3S-IND
 “he [a dog] is really good for catching polar bears with.”

Now, before delving into the comparison, there may be serious doubts as to whether the extremely complex and productive morphology of polysynthetic languages could possibly encode meanings holistically and thus non-compositionally. Indeed, the above authors mention theoretical possibility for verbalising and nominalising affixes in Yupik and Inuit languages to be applied recursively, which would be troublesome for a holistic stance, given issues surrounding learnability. One thing to note is that the high productivity of polysynthetic languages is most notable in the derivational domain, which we consider to be apart from the inflectional system that more so organises the paradigmatic structure of the lexicon. Another would be that the tallies of basic lexical stems, derivational elements and inflectional affixes rarely exceed a few hundred, such as the respectively 400, 171, and 128 in the Whapmagoostui dialect of East Cree, an Algonquian language (relayed in Dorais, 2017 from personal communications with specialist Vincent Collette). It is instead the range of possible combinations of these limited elements that provides the greater source of complexity. In any case, let us park the matter and suspend

any potential disbelief for the time being in order to continue the discussion.

In syntactocentric theories, there is conceivably no, or little difference between the interpretation procedures involved to yield the meanings of sentences for analytic and polysynthetic languages, save for phenogrammatical variation. Consider such an interpretation of a sentence that is expressed by many words in the analytic language, and but one in its polysynthetic counterpart, such as [45]. In the former, individual lexical items constitute the basic building blocks for a compositional interpretation, whereas in the latter, it is instead a number of lexical, derivational and inflectional morphemes that are the inputs, but they are likely involved in roughly the same process, providing an analogous set of basic meanings, and being combined in an equivalent manner tectogrammatically.

In the account provided here, the ultimate meaning of the sentence remains the same, as does the interpretation procedure for the analytic language. The difference is that the interpretation of the single, morphologically complex word-form is not compositional. In fact, ideally, it already bears the sentence-meaning itself. The reason why this would be so is due, firstly, to its relations to other paradigmatic forms, where the inflectional system of the lexicon discriminates the meaning of this word by comparison to others, secondly, to its relations to other lexemes' meanings by way of its derivational affixes, and lastly, because of the meaning of its lexical stem.

Now, if the extremely large (and in certain cases theoretically infinite) lexicons of polysynthetic languages are to be learnable, the space of possible word-meanings must be constrained in some way. Talmy (2000) provides a conceivable limit for the semantic scope of holophrastic forms, where they may only denote one 'macro-event', that represents the "conceptual integration of a complex event". There is one 'framing' event, and potentially other 'co-events' that can be expressed through certain other affixes, or the initial choice of lexical base. There may be only one subject, however, amongst several other constraints. Such a construct could be taken to constrain the holistic structure of the lexicon, by systematically limiting the what the kinds of semantic relations amongst lexemes and inflected word-forms.

This is all highly speculative, of course, and no possible semantic implementation that would involve these ideas is attempted here. It may turn out to be the case that the morphological complexity of polysynthetic languages is simply too great to assume a that a holistic approach to their semantics is possible. The goal of this exposition is to suggest that, with respect to a theory of interpretation, morphological variation should perhaps be taken more seriously. If we take morphology to be directly semantically relevant in the word-based sense elaborated throughout the above, then we may well expect drastic consequences cross-linguistically for the resultant interpretation procedures that theorists would posit to be involved in the determination of the meanings of whole sentences. This is because the initial atomic inputs to those procedures will vary greatly in terms of their basic meanings.

The task of finding data that could help to probe whether this prediction is indeed borne out in reality is left to future research. To ignore this possibility, however, may well be considered an injustice in some sense, given the

awesome scope of variation in morphological complexity exhibited across the world's languages. If anything, it is arguably more presumptuous to claim that these diverse languages are all beholden to a broadly analogous process of interpretation, with a constant set of inputs, and that morphological variation is therefore but a phenogrammatical quirk. It may be that no-one is making quite such a claim, but the assumption of morphological decomposition could conceivably lead one in this direction.

In suggesting variation in the interpretation procedure, no further claim is made regarding the differences in the ultimate meanings that can be expressed by different languages; it is only in the means by which speakers of different languages arrive at common sentence-level interpretations —namely the basic lexical items that serve as inputs —that variation is assumed. This is not problematic, in the same way that claiming variation in the ultimate sentence-level meanings could be, as that may suggest differences in cognitive make-ups of different cultural communities, amongst other potential philosophically suspect implications.

The more modest, yet still significant scope for variation proposed here, in terms of the atomic forms stored in the lexicon with their basic meanings, is perhaps a more suitable locus of linguistic variation, striking something of a middle ground between explaining morphology away as a phenomenon that lacks direct semantic import, and presuming that different linguistic communities have entirely different kinds of meanings. Such an intermediate view is consonant with Apresjan's (1974) suggestion that a language's lexicon reflects a culture's *naivmaja kartina mira*, or their "naive picture of the world". This informed Borshev & Partee's (1998) venture into lexical semantics, which was referenced in section 3.1.3.

Returning to Dowty's (2007) suggestion that the study of 'natural language compositionality' probe *how* and not *whether* different natural languages are compositional, the answer may be that, when it comes to certain expressions, namely morphologically complex word-forms, certain languages are in no way so. No language expresses all possible meanings with single words, however. A typological question that follows from this would then be: "to what extent are languages compositional (and to what extent are they, by way of their lexicons, holistic)?"

Chapter 5

Conclusion

Prior to concluding, it is worth very briefly recapitulating the contents of the above. Following a short introduction, a guiding conception of the grammar was provided, along with a set of related theoretical criteria. The relevance of morphology to this conception was then discussed. That set up an evaluation of a number of alternative kinds of morphological theory with respect to the adopted view of the grammar. A word-based view of morphology was then settled upon, which invited an analysis of the arguments surrounding the concept of ‘the word’. This was followed by a more precise account of the specific Word-and-Paradigm morphological theory, which outlined some of the more practical concerns around its plausibility associated to learnability and processing. The philosophical relevance of such a position, in terms of a holistic conception of the lexicon, was then elaborated. In the final chapter, a formal semantic implementation of those ideas was attempted for a small set of (inflectional) morphological phenomena, some of whose semantic import may have initially been questionable (and may remain so). Lastly, a discussion of the consequences of such a theoretical stance with respect to the issues of diachronic and typological morphological variation was presented.

In sum, the goal has been to explore the relationship between morphological theory and the principle of compositionality, in terms of what the syntactically atomic, most basic meaningful units that provide the input to a compositional interpretation may be. The result of the investigation is that whole, morphologically complex word-forms are the leading candidate, ahead of whatever parts they may appear to be decomposable into. This is in large part because of the problems that arise in attempting to isolate those parts both in terms of form and meaning, given the conception of the grammar adhered to. It is the organisation of the lexicon as a whole, which captures systematically the grammatical import of syntactically atomic, morphologically complex word-forms.

The claims may have been put forward rather polemically at times throughout. This is in part because significant challenges were made to some of the typical assumptions concerning how the meanings of apparently complex expressions are arrived at. It may be the case, however, that, when analysing the

grammars of individual languages, such as those with highly iconic strongly agglutinating morphological structure, the decomposition of complex word-forms does not encounter any such difficulties. Perhaps, in such cases, the stance could be weakened; one position need not be taken to hold for all languages, with more subtlety surely invited given a focus on particulars. Nonetheless, the arguments here were made in sweeping generality, and hence their zeal.

The position developed here seems to be quite novel, especially in terms of the close relationship envisaged between morphology and semantics articulated in terms of a holistic conception of the lexicon. This implicates a number of conceivably disparate fields of study, from philosophy to linguistic typology, and could present a number of possible fertile avenues for future research.

Bibliography

- [1] Ackerman, F., & Malouf, R. (2013). Morphological organization: The low conditional entropy conjecture. *Language*, 429-464.
- [2] Ackerman, F., & Stump, G. (2004). Paradigms and periphrastic expression: A study in realization-based lexicalism. *Projecting morphology*, 111-157.
- [3] Apresjan, J. D. (1974). *Leksičeskaja semantika [Lexical semantics]*. Moskva: Nauka.
- [4] Aronoff, Mark (1994). *Morphology by Itself*. Cambridge: MIT Press
- [5] Baayen, R. H., Lieber, R., Schreuder, R. (1997). The morphological complexity of simplex nouns.
- [6] Barker, C. (2002). Continuations and the nature of quantification. *Natural language semantics*, 10(3), 211-242.
- [7] Blevins, J. P. (2006). Word-based morphology. *Journal of linguistics*, 42(3), 531-573.
- [8] Blevins, J. P. (2013). The information-theoretic turn. *Psihologija*, 46(4), 355-375.
- [9] Blevins, J. P. (2016). *Word and paradigm morphology*. Oxford University Press.
- [10] Blevins, J. P., Ackerman, F., Malouf, R. (2018). Word and Paradigm Morphology. Ch. 13 in Jenny Audring, and Francesca Masini (eds), *The Oxford Handbook of Morphological Theory*, Oxford Handbooks.
- [11] Borer, H. (2012). In the event of a nominal. *The theta system*, 103-150.
- [12] Borer, H. (2014). Derived nominals and the domain of content. *Lingua*, 141, 71-96.
- [13] Borschev, V., Partee, B. H. (1998). Formal and lexical semantics and the genitive in negated existential sentences in Russian. In *Formal Approaches to Slavic Linguistics (Vol. 6, pp. 75-96)*. The Connecticut Meeting 1997.

- [14] Corbett, G. G. (2000). *Number*. Cambridge University Press, Cambridge.
- [15] Curry, H. B. (1961). Some logical aspects of grammatical structure. *Structure of language and its mathematical aspects*, 12, 56-68.
- [16] Davidson, D., Rescher, N. (1967). The logical form of action sentences. 1967, 105-122.
- [17] Dorais, L. J. (2017). The Lexicon in Polysynthetic Languages. Ch. 8 in Michael Fortescue, Marianne Mithun, and Nicholas Evans (eds), *The Oxford Handbook of Polysynthesis*, Oxford Handbooks
- [18] Dowty, D. (2007). Compositionality as an empirical problem. *Direct compositionality*, 14, 23-101.
- [19] Dye, M., Milin, P., Futrell, R., Ramscar, M. (2017). A functional theory of gender paradigms. In *Perspectives on morphological organization* (pp. 212-239). Brill
- [20] de Swart, H., Farkas, D. (2010). The semantics and pragmatics of plurals. *Semantics and pragmatics*, 3, 6-1.
- [21] Dubinsky, S., Simango, S. R. (1996). Passive and stative in Chichewa: Evidence for modular distinctions in grammar. *Language*, 749-781.
- [22] Fertig, D. L. (2013). *Analogy and morphological change*. Edinburgh University Press.
- [23] Fillmore, C. J. (1976). Frame semantics and the nature of language. In *Annals of the New York Academy of Sciences: Conference on the origin and development of language and speech* (Vol. 280, No. 1, pp. 20-32).
- [24] Fortescue, M. D. (1984). *West Greenlandic*. London: Croom Helm.
- [25] Fortescue, M., Mithun, M., Evans, N. (2014). What are the limits of Polysynthesis?. In *International Symposium on Polysynthesis in the World's Languages*.
- [26] Futrell, R., Mahowald, K., Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33), 10336-10341
- [27] Habicht, K. (2000). Grammaticalization of Adposition in Old Literary Estonian - Estonian: *Typological Studies IV*, Tartu (Tartu Ülikooli eesti keele õppe-tooli toimetised 14), 19-58.
- [28] Haiman, J. (1980). Hua. Hua, 1-608.
- [29] Halle, M., Marantz, A., Hale, K., Keyser, S. J. (1993). Distributed morphology and the pieces of inflection. 1993, 111-176.

- [30] Hankamer, J., Sag, I. (1976). Deep and surface anaphora. *Linguistic inquiry*, 7(3), 391-428.
- [31] Harley, H., Ritter, E. (2002). Person and number in pronouns: A feature-geometric analysis. *Language*, 482-526.
- [32] Huybregts, M.A.C. (1976). Overlapping dependencies in Dutch. Technical report, University of Utrecht. Utrecht Working Papers in Linguistics.
- [33] Jacobson, P. (2002). The (dis) organization of the grammar: 25 years. *Linguistics and Philosophy*, 25(5/6), 601-626.
- [34] Jacobson, P. (2012). Direct compositionality and ‘uninterpretability’: The case of (sometimes) ‘uninterpretable’ features on pronouns. *Journal of Semantics*, 29(3), 305-343.
- [35] Lieber, R. (1992). *Deconstructing morphology: Word formation in syntactic theory*. University of Chicago Press.
- [36] Link, G. (1983). The logical analysis of plurals and mass terms: A lattice-theoretical approach. *Formal semantics: The essential readings*, 127-147.
- [37] Montague, R. (1970). Universal grammar. *Theoria*, 36(3), 373-398.
- [38] Milin, P., Kuperman, V., Kostic, A., Baayen, R. H. (2009). Paradigms bit by bit: An information theoretic approach to the processing of paradigmatic structure in inflection and derivation. *Analogy in grammar: Form and acquisition*, 381, 214-252.
- [39] Moscoso del Prado Martín, F. M., Kostić, A., Baayen, R. H. (2004). Putting the bits together: An information theoretical perspective on morphological processing. *Cognition*, 94(1), 1-18.
- [40] Pagin, P., Westerståhl, D. (2010). Compositionality II: Arguments and problems. *Philosophy Compass*, 5(3), 265-282.
- [41] Partee, B. (1984). Compositionality. *Varieties of formal semantics*, 3, 281-311.
- [42] Pelletier, F. J. (2012). Holism and compositionality. *The Oxford handbook of compositionality*, 149-174.
- [43] Radkevich, N. (2010). On Location: The Structure of Case and Adpositions.
- [44] Sapir, E. (2004). *Language: An introduction to the study of speech*. Courier Corporation.
- [45] Schreuder, Robert; and Baayen, R. Harald (1997). How complex simplex words can be. *Journal of Memory and Language* 37,118-139.

- [46] Shieber, S. M. (1985). Evidence against the context-freeness of natural language. *Linguistics and Philosophy*, 8(3), 333–344.
- [47] Steele, S. (1995). Towards a theory of morphological information. *Language*, 260-309.
- [48] Stump, G. T. (2001). *Inflectional morphology: A theory of paradigm structure* (Vol. 93). Cambridge University Press.
- [49] Stump, G. T. (2018). Theoretical Issues in Inflection. Ch. 4 in Jenny Audring, and Francesca Masini (eds), *The Oxford Handbook of Morphological Theory*, Oxford Handbooks.
- [50] Talmy, L. (2000). *Toward a cognitive semantics* (Vol. 2). MIT press.
- [51] Zanuttini, R., Pak, M., Portner, P. (2012). A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language & Linguistic Theory*, 30(4), 1231-1274.