Realizational Morphosemantics in L_RFG

Ash Asudeh
University of Rochester/
Carleton University

Daniel Siddiqi
Carleton University

Proceedings of the LFG'22 Conference

Miriam Butt, Jamie Y. Findlay and Ida Toivonen (Editors)

2022

CSLI Publications

http://csli-publications.stanford.edu/

Abstract

Lexical-Realizational Functional Grammar (L_RFG) is a theoretical framework that couples Lexical-Functional Grammar with the realizational, morpheme-based approach to word-formation of Distributed Morphology (DM). L_RFG has been developed to take advantage of LFG's strengths in modelling nonconfigurationality and DM's strengths in modelling complex nonfusional morphology and polysynthesis. We present some initial attempts at an L_RFG theory and formalization of *morphosemantics*, i.e. the morphology–semantics interface. We distinguish the domain of morphosemantics from the more general domain of lexical semantics: morphosemantics encompasses all and only aspects of meaning that affect the mapping from a semantic representation to a phonological representation. We focus on the phenomenon of *blocking* and, in particular, where blocking fails and a regular and irregular form are both possible. Modern approaches to this type of blocking face the following challenge: certain approaches essentially predict blocking to *always* happen, while others predict it to *never* happen. We focus on the former case: contemporary realizational approaches predict blocking to occur in situations where it does not, yielding a potential undergeneration of forms. We claim that the correct realizational approach to blocking requires reference to compositional semantics and show how this can be done in L_RFG. This correctly accounts for the distribution of four putative irregular/regular pairs that we take as case studies: *divinity/divineness*, *unkempt/uncombed*, *people/persons*, and *brethren/brothers*.

1 Introduction

Lexical-Realizational Functional Grammar (L_RFG; Melchin et al. 2020b, Asudeh et al. 2021, Asudeh and Siddiqi forthcoming) is a theoretical framework that couples Lexical-Functional Grammar (LFG; Bresnan et al. 2016) with the realizational, morpheme-based approach to word-formation of Distributed Morphology (DM; Halle and Marantz 1993).[†] According to the classification of morphological theories offered by Stump (2001), L_RFG is

1. *Lexical:* The lexicon is an inert list of mappings from formal properties to phonological representations (a.k.a. morphemes); and
2. *Realizational:* Morphology expresses syntactic categories and features and, in some theories, semantics.

In this paper, we present some initial attempts at an L_RFG theory and formalization of *morphosemantics*, i.e. the morphology–semantics interface.

[†]This work is part of an ongoing project led by Ash Asudeh and Dan Siddiqi; see lrfg.online. The project also involves Oleg Belyaev (Moscow State University), Bronwyn Bjorkman (Queen's University), Tina Bögel (University of Konstanz), Michael Everdell (University of Texas, Austin), Paul Melchin (Carleton University), Will Oxford (University of Manitoba) and our students, Veronica Burrage (Rochester) and Sam Turnbull (Carleton). We are grateful to all the project members for their participation and discussion, but especially to Mike, Paul, and Tina, who have thus far been our main collaborators. We are also grateful to the audience at LFG22 for their questions and comments. We thank Jamie Findlay, in particular, for extended informal discussion and careful editorial guidance. Last, but not least, we thank the two anonymous reviewers of this paper, who helped improve it greatly. Any remaining errors are our own.

The L_RFG framework has been described in a number of papers (Melchin et al. 2020b, Asudeh et al. 2021, Everdell et al. 2021, Everdell and Melchin 2021) and a book-length treatment is also underway (Asudeh and Siddiqi forthcoming). An L_RFG grammar defines a set of valid form-meaning pairs, based in part on a set of formatives (the Vocabulary) and a set of c-structure rules. The form incorporates prosody (fed by constituent structure, as in LFG) and the meaning incorporates information structure (fed by semantic structure, as in LFG). Here we want to focus on the morphology–semantics interface (i.e., *morphosemantics*) in L_RFG, although we won't have anything to say about the ι-mapping to information structure. This is schematized in Figure 1.

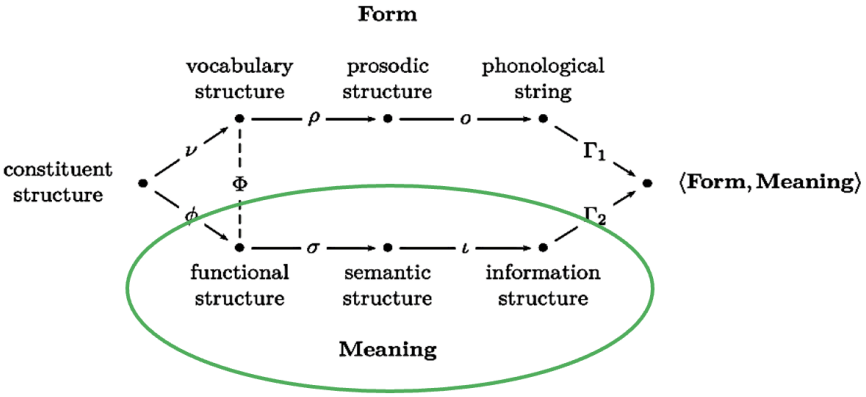

Figure 1: Morphosemantics in L_RFG's Correspondence Architecture

The paper is organized as follows. Section 2 introduces the framework of L_RFG. Section 3 looks at some problems at the morphology-semantics interface, in general terms. Section 4 provides details on L_RFG's exponence function, ν. Section 5 looks at L_RFG solutions to the aforementioned problems, which provides a sense of the general shape of L_RFG's treatment of morphosemantics. In particular, we offer partial analyses of four case studies:

(1) *divinity/divineness*
(2) *unkempt/uncombed*
(3) *people/persons*
(4) *brethren/brothers*

Section 6 offers some conclusions and prospects.

2 The L_RFG framework

L_RFG is a synthesis between Distributed Morphology as a theory of morphological realization and Lexical-Functional Grammar as a theory of syntax and grammatical architecture. L_RFG combines the strengths of the two frameworks:

1. Like LFG, it is a declarative, representational and constraint-based theory (without the bottom-up, phase-based derivations of Minimalism) that is ideally suited to modelling nonconfigurationality.

2. Like DM, it provides a realizational, morpheme-based view of word-formation and is good at modelling complex morphological structures including those found in polysynthetic languages, such as many North American Indigenous languages.

Nonconfigurationality and polysynthesis are two key phenomena that L_RFG seeks to model and explain, as in previous work on Ojibwe (Melchin et al. 2020b,a, Asudeh et al. 2021) and O'dam (Everdell et al. 2021, Everdell and Melchin 2021). However, this paper focuses on English examples, so these phenomena are not central concerns here.

2.1 Architecture and example

The L_RFG architecture was shown in Figure 1 above. L_RFG is syntactically similar to standard LFG, with changes to the c(onstituent)-structure tree and its relationship with morphosyntactic elements. The terminal nodes of c-structures *are not words*, but instead are *f-descriptions* (sets of f(unctional)-structure equations and constraints) and sets of Glue Semantics *meaning constructors* (terms that are used in the computation of compositional semantics). The categorial information in c-structure preterminals and the f-structural and Glue information c-structure are mapped to v(ocabulary)-structures via the exponence/correspondence function, ν.[1] Thus, a v-structure *expones* (i.e., realizes) the features in a c-structure terminal, given some c-structure category that labels a preterminal. Vocabulary structure is therefore a morphophonological level that maps from syntax and semantics to phonological form via prosodic structure.

Example (5) below is from Ojibwe (*Anishinaabemowin*, Algonquian). We use it to demonstrate the basics of an L_RFG analysis (for further details of L_RFG for Ojibwe agreement, see Melchin et al. 2020b). The L_RFG c-structure, f-structure and v-structure mappings are shown in (6). The relationship between terminal nodes and VIs is many-to-one, using the mechanism of *Spanning* (Haugen and Siddiqi 2016, Merchant 2015, Ramchand 2008, Svenonius 2016); i.e. one VI may realize features of multiple terminal nodes. The result bears some similarities to the Lexical Sharing model proposed for LFG by Wescoat (2002, 2005, 2007), but maintains, like DM, that the complex internal structures of words are part of syntax.

(5) gi- gii- waab -am -igw -naan -ag Ojibwe
2 PST see VTA INV 1PL 3PL
'They saw us, including you.'

[1] L_RFG inherits from Distributed Morphology the idea of partitioning morphological listedness into three domains (typically called 'lists'): syntactic, semantic, and phonological. The L_RFG enactment of this idea is through categories (c-structure), f-descriptions (f-structure), Glue meaning constructors (s-structure, semantics), and Vocabulary Items that realize the syntactic (c-structural and f-structural) and semantic information. The exponence function, ν, maps this information to a list of exponents, as conditioned by the listed items in the Vocabulary and the constraints on the mapping to prosody and phonology. Thus DM's syntactic and semantic lists are captured in the set of inputs to ν, as listed in the Vocabulary, and DM's idea of a phonological list is captured in the set of vocabulary structures that are the outputs of ν, again as listed in the Vocabulary.

(6)

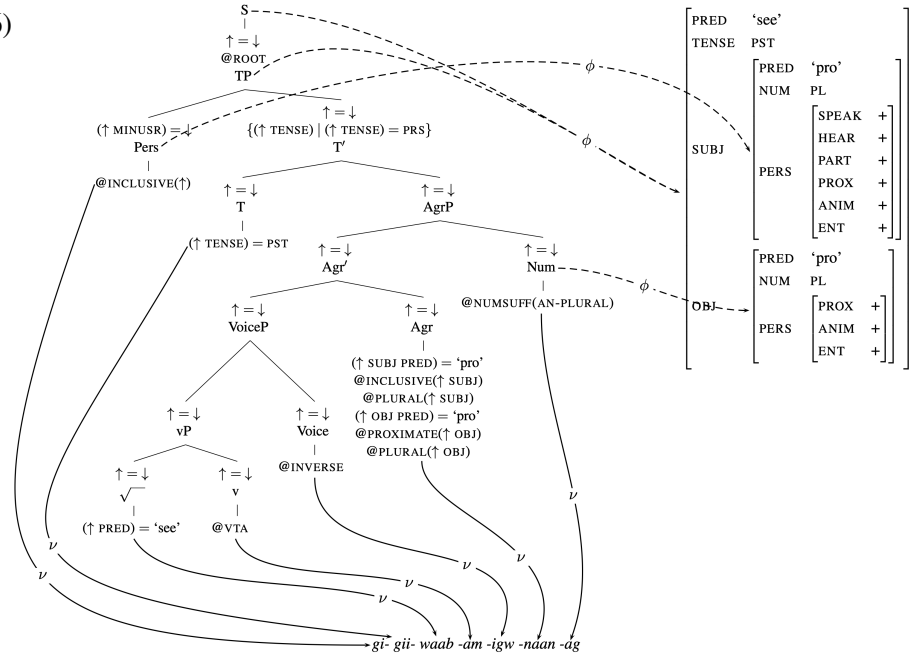

There is also some similarity between L_RFG's Vocabulary, a theoretical construct which it inherits from DM, and the lexicon in LFG. The Vocabulary is a set of Vocabulary Items (VIs). Each VI specifies a mapping from c-structural, f-structural, and semantic information to a v(ocabulary)-structure, which includes information about phonological form, phonological dependency, and prosody, as well as morphosyntactic and properly morphological information. This mapping is the exponence/correspondence function, ν. Here is the VI for Ojibwe *gii-* (PST) from (5) above, following Melchin et al. (2020b: 283):

(7) $\langle$ [T], $\Phi\{$ (↑ TENSE) = PST $\}$, $\lambda P \lambda t.P(t) \wedge t \prec u$: $((\uparrow_\sigma \text{TENSE}) \multimap \uparrow_\sigma) \multimap ((\uparrow_\sigma \text{TENSE}) \multimap \uparrow_\sigma)$ $\rangle$

$\xrightarrow{\nu}$ *gii-*

The relevant c-structure rule is as follows:

(8) $\text{T}' \longrightarrow$ T (↑ TENSE) = PST, AgrP ↑ = ↓

Lastly, the nonconfigurationality of Ojibwe is modelled as in standard LFG, by assuming that the root node is an exocentric S node. We do not have the space to explore the consequences of this further here, but see Melchin et al. (2020b: 277). The equivalent LFG representation to (6) would be the following:

(9)

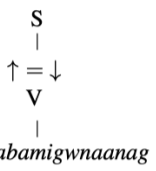

F-descriptions/templates as in (6)

In sum, the L_RFG representation in (6) contains the same f-structural information as the LFG representation in (9). However, L_RFG follows DM in using the syntactic structure, in this case c-structure, to model morphological structure and function directly; thus, the TP in (6) shows the morphological structure of the V node in (9) and the f-structure shows the morphological function, as in standard LFG. In contrast, in standard LFG the word *gigiiwaabamigwnaanag* would be licensed by a separate morphological component which also pairs the form with information about morphological function, which in turn forms the basis for the f-descriptions that define and constrain the f-structure.[2] For further discussion of different approaches to morphology in LFG, see Asudeh and Siddiqi (2022).

3 Motivation: Morphosemantic problems

An important preliminary question that we have to answer concerns how morphosemantics is distinct from general lexical semantics. We regard morphosemantics as encompassing all and only aspects of meaning that affect the mapping from a semantic representation to a phonological representation. In L_RFG terms, it is those meanings that condition the mapping to v-structure. The principle that governs this mapping, formalized in (29) below, is **MostInformative**$_s$.

Morphosemantic phenomena in L_RFG include the following:

1. Semantically conditioned morphology — morphemes which have semantic well-formedness conditions on their base. For example, *re-* is conditioned by the aktionsart of its base.[3]

   (10) *re-establish* (11) **re-believe*

2. Polysemy — morphemes which can appear in a wide variety of semantic and functional environments. For example, *keep* is used in many senses such that the grammar is sensitive to sense differences.

   (12) a. I always keep$_1$ my appointments.
   b. I always keep$_2$ my keys in my pocket.
   c. !I always keep$_{*1/2}$ my appointments [(e.g., written down)] and my keys in my pocket.

3. Lexicalization — complex morphological forms, consisting of seemingly productive morphology, that do not necessarily have the compositionally predicted meanings.

   (13) *antsy* (does not necessarily involve ants)

   (14) *lousy* (does not necessarily involve lice)

[2] A reviewer wonders if we are using f-descriptions as *representations* rather than as *descriptions* of representations, which would be a departure from standard LFG. We are not: the f-description for a particular VI is the input to a function Φ which returns the set of f-structures that are defined and constrained by the f-description. See Section 4 below for further information on Φ.

[3] We thank an anonymous reviewer for the second example.

4. Irregulars/regulars — extant pairs of regular and irregular forms where one form contains more specialized meaning

(15)	*brethren/brothers*	(17)	*divinity/divineness*
(16)	*unkempt/uncombed*	(18)	*cómparable/compárable*

In this paper, we consider putative varieties of 4.

This phenomenon is of particular interest to morphological theory because it represents a failure of *blocking* (Aronoff 1976). It is typically the case that irregular allomorphs block their regular counterparts. Indeed, this is exactly how we know that the relevant alternation is one of irregular allomorphy and not synonymy (for discussion, see Siddiqi 2021). For example, we know that *ran* is an allomorph of *run* precisely because **runned* is blocked. We are not sure that *raise* is an allomorph of *rise* (despite a similar irregular morphological process) because there is no allomorph that is blocked.[4] Similarly, it is debatable whether *people* is the plural of *person*, precisely because *persons* is licit in some environments (Arregi and Nevins 2013).

This blocking of regular allomorphs is typically inviolate, especially in high frequency words and also especially in uncontroversially inflectional morphology (*ate*/*went* always block **eated*/**goed*). It has been well-known since at least Aronoff (1976) that blocking of regular, fully productive forms can fail and that both irregulars and regulars can exist in parallel, in some circumstances. Aronoff (1976) has in some ways become the base way that we conceptualize blocking. In that approach, *glory* blocks **gloriosity* (and similarly *to sweep* blocks **to broom*) because the output of irregular morphological processes (such as suffixation of *-ity*) need to be stored in a certain semantic space. That space in the context of **gloriosity* is occupied by *glory*. However, perfectly productive and transparent morphological processes (such as suffixation of *-ness*) do not need to be stored because their meanings are completely predictable, so they cannot be blocked. Thus, *glory* does not block licit *gloriousness*.

In contemporary realizational morphological theory, failures of blocking are especially relevant. In realizational models, blocking falls naturally out of the core mechanism of the grammar. From a certain point of view, all morphemes are allomorphs of all other morphemes that they are in competition with. Allomorphy in these approaches is crucially conceptualized differently from the traditional Aronoff (1976) approach. Morphological forms *express* underlying meaning rather than contributing meaning. In some sense, *run* and *sneeze* are allomorphs of each other in expressing the underlying feature of being a verb. It is information beyond being a verb that determines which exponent is realized, such as expressing a meaning *run′* rather than *sneeze′*. For example, where the meaning *run′* is being expressed, *run* blocks *sneeze*, and vice versa mutatis mutandis. From the same point of view, *una* ('D.INDEF.FEM', Spanish) blocks *un* ('D.INDEF') in *una mujer*

[4] See the discussion of the English causative/inchoative alternation in Harley (2014) and responses in Siddiqi (2021).

('a woman') because the allomorph *una* expresses more information. Thus, the winner of any given competition effectively blocks the rest. While we have here couched this in *Lexical*-Realizational terms, the same logic applies to *Inferential*-Realizational models such as Paradigm Function Morphology (Stump 2001). In other words, this is a general point about realizational morphology. Put simply, realizational rules are fundamentally Paninian, meaning the most specific will apply in a given context, thus preferring portmanteau forms over general forms. This interaction of grammatical architecture and blocking is discussed at length in Embick and Marantz (2008).

Both these approaches to blocking are in some ways "wrong" in that they both incorrectly predict extremes. Realizational blocking expects blocking to *always* happen, and Aronovian blocking expects complete blocking of productive morphology to *never* happen. As this paper assumes a Lexical-Realizational approach to morphology, we here focus on the conditions under which blocking fails such that we get both a regular and irregular form. Thus, we seek to answer the following question:

(19) **Q:** If X, an irregular form, is a more specific form of Y, a regular form, how can X fail to block Y?

For example, if *divinity*/*curiosity*/*productivity* are more specific forms of respectively *divineness*/*curiousness*/*productiveness*, why are the latter not blocked by the relevant instantiation of the Paninian principle? We ground our analysis in the generalization that, in such cases (as discussed above), the irregular expresses non-compositional meaning and thus is not in competition with the regular. However, in contexts where the regular and irregular are deployed with the same intended meaning, the irregular should indeed block the regular. In L_RFG, exponence has access to compositional semantics — i.e., the *actual* semantics, rather than "feature semantics"/markerese or some other ad hoc syntactic markup — and any pragmatic place-holder variables (as in, e.g., Partee and Borschev 2003). This allows for a more precise and nuanced type of analysis, making L_RFG perhaps unique among Lexical-Realizational models of morphology.

4 L_RFG's exponence function: ν

In our previous work (Melchin et al. 2020b, Asudeh et al. 2021, Everdell et al. 2021), the exponence function ν mapped from a pair of arguments to a v(ocabulary)-structure, the exponent. However, since we are now turning our attention to semantics as well, we add a third argument to ν:

1. The first argument is a list of pre-terminal categories, typically of length 1, which are taken in the linear order in which they appear in the tree.
2. The second argument is itself a function, Φ, which maps an f-description to the set of f-structures that satisfy the description; i.e. $\Phi(d \in D) = \{f \in F \mid f \models d\}$, where D is the set of f-descriptions and F is the set of f-structures.[2]

[2]We thank Ron Kaplan (p.c.) for discussion of this point. Any remaining errors are our own.

3. The third argument is a set of *meaning constructors* from Glue Semantics (Glue; among others, Dalrymple 1999, Dalrymple et al. 2019, Asudeh 2012). For a recent high-level introduction to Glue Semantics, see Asudeh (2022a,b).

Meaning constructors are pairs of terms from two logics (the colon is an uninterpreted pairing symbol):

(20) $\mathcal{M} : G$

$\mathcal{M}$ is an expression of the *meaning language* — anything that supports the lambda calculus. G is an expression of *linear logic* (Girard 1987), which specifies semantic composition based on a syntactic parse that instantiates the general terms in G to a specific syntactic structure. The meaning constructors serve as premises in a linear logic proof of the *compositional semantics*.

(21) Alex likes Blake.

(22) Meaning constructors: $\mathbf{alex} : \uparrow_\sigma$
$\mathbf{blake} : \uparrow_\sigma$
$\lambda y.\lambda x.\mathbf{like}(y)(x) : (\uparrow \text{OBJ})_\sigma \multimap (\uparrow \text{SUBJ})_\sigma \multimap \uparrow_\sigma$

Note that $\lambda y.\lambda x.\mathbf{like}(y)(x)$ is η-equivalent to just **like**, but it is useful to use the expanded form to make the structure of the following proof more obvious. Note that in the proof we have instantiated the general Glue terms in (22) mnemonically.

(23)
$$\dfrac{\mathbf{alex} : a \qquad \dfrac{\lambda y.\lambda x.\mathbf{like}(y)(x) : b \multimap a \multimap l \qquad \mathbf{blake} : b}{\lambda x.\mathbf{like}(\mathbf{blake})(x) : a \multimap l}\ \multimap_\mathcal{E}, \Rightarrow_\beta}{\mathbf{like}(\mathbf{blake})(\mathbf{alex}) : l}\ \multimap_\mathcal{E}, \Rightarrow_\beta$$

The blue colour in the proof is not part of the representation, but highlights the meaning constructors in (22), which are contributed by the elements of sentence (21), as opposed to compositionally derived meanings, which are in black.

Here are two sample VIs, the first for the Ojibwe root $\sqrt{\text{SEE}}$ realized as *waab* ('see'), as in (6) above, and the second for the English equivalent *see*. Note that we use the η-equivalent form of the **see** function to reduce clutter.[3]

(24) Ojibwe
$\langle\ [\sqrt{\ }], \Phi\{(\uparrow \text{PRED}) = \text{'see'}\}, \{\mathbf{see} : (\uparrow \text{OBJ})_\sigma \multimap (\uparrow \text{SUBJ})_\sigma \multimap \uparrow_\sigma\}\ \rangle \quad \overset{\nu}{\rightarrow} \quad \textit{waab}$

(25) English
$\langle\ [\sqrt{\ }], \Phi\{(\uparrow \text{PRED}) = \text{'see'}\}, \{\mathbf{see} : (\uparrow \text{OBJ})_\sigma \multimap (\uparrow \text{SUBJ})_\sigma \multimap \uparrow_\sigma\}\ \rangle \quad \overset{\nu}{\rightarrow} \quad \textit{see}$

In a c-structure tree, this is represented as follows:

[3] We also henceforth ignore tense information. In order to be compatible with the meaning constructor for tense in (7) above, the meaning constructor for *waab/see* would have to deal with the tense variable, such that **see** : $(\uparrow \text{OBJ})_\sigma \multimap (\uparrow \text{SUBJ})_\sigma \multimap (\uparrow \text{TENSE})_\sigma \multimap \uparrow_\sigma$, where **see** $\Leftarrow_\eta$ $\lambda y \lambda x \lambda t.\mathbf{see}(y)(x)(t)$.

(26)
$$\sqrt{}$$
$$|$$
$$(\uparrow \text{PRED}) = \text{`see'}$$
$$\textbf{see} : (\uparrow \text{OBJ})_\sigma \multimap (\uparrow \text{SUBJ})_\sigma \multimap \uparrow_\sigma$$

Henceforth, we will show only the meaning language side of the Glue meaning constructors.

4.1 Conditions on exponence

Let V^i be the domain of the exponence function ν in some language L, i.e. the set of inputs to Vocabulary Items in L. We write $V^i(\alpha)$ to indicate the domain of some particular Vocabulary Item, α. We write $\pi_n(V^i(\alpha))$ to indicate the n^{th} projection of $V^i(\alpha)$. For example, $\pi_1(V^i(\alpha))$ returns the c-structure list in the first projection of the input to Vocabulary Item α.[4] The following conditions on exponence hold:[5]

1. **MostInformative**$_c(\alpha, \beta)$ returns whichever of α,β has the longest list of overlapping c-structure categories.[6]

   (27) Given two Vocabulary Items, α and β,
   $$\textbf{MI}_c(\alpha, \beta) = \begin{cases} \alpha \textbf{ if } \pi_1(V^i(\alpha)) = f \wedge\ \pi_1(V^i(\beta)) = g\ \wedge\ \textbf{span}(f, g) \\ \beta \textbf{ if } \pi_1(V^i(\alpha)) = f \wedge\ \pi_1(V^i(\beta)) = g\ \wedge\ \textbf{span}(g, f) \\ \bot \textbf{ otherwise} \end{cases}$$

The intuition behind **MostInformative**$_c$ is: whenever possible, prefer *portmanteau* forms. In c-structural terms, to be a portmanteau means that the list of categories in the VI must have length greater than 1.

So **MostInformative**$_c$ states that, in terms of lists of categories in Vocabulary Items, choose the VI that realizes the greater set of categories: prefer portmanteau forms on c-structural grounds.

2. **MostInformative**$_f(\alpha, \beta)$ returns whichever of α,β has the most specific f-structure in the set of f-structures returned by Φ applied to α/β's collected f-description.

   The proper subsumption relation on f-structures (Bresnan et al. 2016: chap. 5) is used to capture the intuition (below).

[4] This π is just standard notation for retrieving arguments to functions and should not be mistaken for a correspondence function.

[5] **MostInformative**$_f$, which is based on the subsumption relation between f-structure, is clearly related to the proposal of Andrews (1990). However, **MostInformative**$_c$ and **MostInformative**$_s$ have no correlates in that system, so the overall proposal is distinct.

[6] The function **span** is defined as follows:

(i)
$$\textbf{span}(\mathit{list}_1, \mathit{list}_2) = \begin{cases} \mathsf{first}(\mathit{list}_1) = \mathsf{first}(\mathit{list}_2) \wedge\ \textbf{span}(\mathsf{rest}(\mathit{list}_1), \mathsf{rest}(\mathit{list}_2)) \\ \mathit{list}_1 \neq \mathit{elist}\ \wedge\ \mathit{list}_2 = \mathit{elist} \end{cases}$$

(28) Given two Vocabulary Items, α and β,

$$\mathbf{MI}_f(\alpha,\beta) = \begin{cases} \alpha \textbf{ if } \exists f \forall g. f \in \pi_2(V^i(\alpha)) \wedge g \in \pi_2(V^i(\beta)) \wedge g \sqsubset f \\ \beta \textbf{ if } \exists f \forall g. f \in \pi_2(V^i(\beta)) \wedge g \in \pi_2(V^i(\alpha)) \wedge g \sqsubset f \\ \bot \textbf{ otherwise} \end{cases}$$

As expected, **MostInformative**$_f$ continues to prefer portmanteau forms, whenever possible, but this time on f-structural grounds. In terms of f-descriptions in Vocabulary Items, choose the VI that defines an f-structure that contains the greater set of features.[7] In f-structural terms, to be a portmanteau means that there must be an f-structure in the set returned by Φ of size greater than one. For example, neither Ojibwe *waab* in (24) above nor English *see* in (25) is a portmanteau form.

3. **MostInformative**$_s(\alpha,\beta)$ returns whichever Vocabulary Item has the more specific meaning.

   The proper subset relation on set-denoting expressions is used to capture the intuition (below).

(29) Given two Vocabulary Items, α and β,

$$\mathbf{MI}_s(\alpha,\beta) = \begin{cases} \alpha \textbf{ if } \pi_3(V^i(\alpha)) = f \wedge \pi_3(V^i(\beta)) = g \wedge \\ \qquad [\![\mathsf{M}_{\vdash}(f)]\!] \subset [\![\mathsf{M}_{\vdash}(g)]\!] \\ \beta \textbf{ if } \pi_3(V^i(\alpha)) = f \wedge \pi_3(V^i(\beta)) = g \wedge \\ \qquad [\![\mathsf{M}_{\vdash}(g)]\!] \subset [\![\mathsf{M}_{\vdash}(f)]\!] \\ \bot \textbf{ otherwise} \end{cases}$$

Once more, **MostInformative**$_s$ continues to prefer portmanteau forms, whenever possible, but on semantic grounds. In terms of meanings encoded in Vocabulary Items, choose the VI whose denotation is more semantically contentful. In order to do this, we introduce a function $\mathsf{M}_{\vdash}$ which returns the meaning language side of the conclusion of a Glue proof. We use the subset relation on interpretations (via the standard function $[\![\ ]\!]$) of meaning language terms in order to model semantic contentfulness as strict entailment. In other words, in terms of semantics, to be a portmanteau is to have a more specific meaning than the competitor.

In the case of VIs that are ambiguous and therefore have two or more disjoint meaning constructors, as in Section 5.1, or VIs that have optional meaning constructors, as in Section 5.4, each possible interpretation of α/β necessitates a separate competition. For example, we will see in Section 5.1 that *divineness* is ambiguous and so only competes with *divinity* on one interpretation, such that when

[7] Similar to other implementations of the Paninian principle, such as the Subset Principle in Distributed Morphology, **MostInformative** does not merely discriminate on the basis of the relative size of feature structures or sets of meanings; it delivers $\bot$ if they are not in a subsumptive/subset relationship, whether or not there is a size difference. This is working as intended, because morphological distributions have been found in which putative ties are resolved through some other principle. Bjorkman et al. (2021) call this *morphological upstaging*. For example, it is common for GENDER to not be expressed where PARTICIPANT is (e.g., French *il/elle*, *ils/elles*, but *je*, *tu*, *nous*, *vous*).

divineness appears in the same environment as *divinity*, it can only be the interpretation of *divineness* that does not compete with *divinity* that can be expressed.

Before wrapping this section up, we make two observations. First, $\mathbf{MI}_c$ and $\mathbf{MI}_f$ are *morphosyntactic* constraints, whereas $\mathbf{MI}_s$ is a *morphosemantic* constraint. Second, each version of **MostInformative** can result in a tie, represented by $\perp$. In sum, we have three conditions on exponence that concern the input to the function ν. They constitute a family of **MostInformative** constraints whose upshot is to prefer portmanteau forms on respectively c-structural ($\mathbf{MostInformative}_c$), f-structural ($\mathbf{MostInformative}_f$), or semantic grounds ($\mathbf{MostInformative}_s$).

5 The $\mathbf{L_R}$FG treatment of morphosemantics

Recall that the four cases we are looking at here all concern a specific kind of stem allomorphy. Namely, they are all putative cases of regular and irregular forms in competition, where both the regular and the irregular are grammatical. In these cases, a common view is that the irregular contains more specialized meaning than the regular, as a function of portmanteaus in the grammar (see, e.g., Aronoff 1976 et seq.) and whole word storage in processing (see, e.g., Baayen 1992 et seq.).

5.1 *Divinity/divineness*

The adjective *divine* is ambiguous; it either expresses a (highly positive) quality or some notion like "holiness". The ambiguity is demonstrated in the following example:

(30) This food is divine.

One reading of the sentence is that the food is very good.[8] The other reading is that the food is holy.

The regular affix *-ness* does not express any additional meaning: it simply serves as a deadjectivizing nominalizer. Thus, the regular/productive form *divineness* inherits the ambiguity of the root $\sqrt{\text{DIVINE}}$:

(31) This chocolate is divine but is affordable despite its divineness.

Example (31) has one relevant pragmatically acceptable reading, despite the ambiguity in *divine* and *divineness*, which might lead us to expect at least two more readings. On the acceptable reading, the chocolate is very good, but is affordable despite its high degree of goodness. An obvious alternative reading, in which *divine* is taken to mean holy, is pragmatically odd in the null context, because it does not accord with common world knowledge. Moreover, in a context in which the possible holiness of chocolate is in the common ground, *divinity* pragmatically blocks *divineness* due to implicature, as sketched below.[9]

In the dialect of interest, the irregular form *divinity* necessarily involves the notion of holiness. We therefore perceive a contrast between these two sentences:

[8] In what follows we do not model the degree of goodness, since it's not strictly relevant, but it can be modelled in the usual ways following Kennedy (1999, 2007)

[9] Another possible reading where *divine* is taken to mean holy but *divineness* is taken to mean very good is pragmatically odd, because it does not set up the required contrast for *despite*.

(32) This communion wafer is divine but is bland despite its divinity.

(33) !This chocolate is divine but is affordable despite its divinity.

Unlike (32), (33) is odd in the null context, because it attributes the quality of holiness to chocolate, which does not accord with common world knowledge.

The following examples illustrate that *divinity* does not contain a notion of goodness and that when goodness cannot be attributed (due to world knowledge), the form *divinity* meaning holy is favoured over the form *divineness* taken to mean holy, as shown in (34) and (35).

(34) Cthulhu's divinity/!divineness is terrible to behold.

(35) In *His Dark Materials*, Metatron is evil despite his divinity/#divineness.

In the relevant mythologies, neither Cthulhu nor Metatron is good.[10] For speaker/hearers who know this, (34) and (35) are therefore only acceptable with *divinity*, since otherwise the utterance would be interpreted as contradicting implicit/world knowledge, as in (34), or contradicting an explicit assertion, as in (35).

In contrast, (36a–b) show that the forms *divinity* and *divineness* can occur in the same environment, but in this case *divineness* can only mean goodness.

(36) a. Theologians have long been puzzled by why God allows evil to happen despite His divinity.

b. Theologians have long been puzzled by why God allows evil to happen despite His divineness.

The relevant reading of (36a) is that theologians have long been puzzled by why God allows evil to happen despite His holiness (i.e., surely a holy being would be good and therefore not allow evil to happen). This reading is pragmatically unavailable for (36b). The only relevant reading of (36b), despite the ambiguity of *divineness*, is that theologians have long been puzzled by why God allows evil to happen despite His goodness (i.e., surely a good being would not allow evil to happen). In other words, for speaker/hearers who have the form *divinity*, utterance (36b) with the same meaning is pragmatically disprefered to/blocked by utterance (36a), for reasons already touched on. By implicature, a speaker who intends specifically holiness should choose *divinity* because the alternative is ambiguous.

We take examples such as these as evidence for our stance that *divine* and *divineness* are ambiguous between meaning very good or holy, whereas *divinity* just means holy. The relevant meanings can be represented something like this:

(37) $[\![\textit{divineness}]\!] = \textbf{good}$ or $[\![\textit{divineness}]\!] = \textbf{holy}$

(38) $[\![\textit{divinity}]\!] = \textbf{holy}$

On pragmatic grounds, if the relevant property is holiness, as expressed by a meaning constructor with meaning **holy**, this is best realized as *divinity*, not *divineness*.[11]

[10] A speaker who does not know this about Cthulhu might therefore accept (34) as fine with *divineness*, should they believe that extreme goodness may be terrible to behold. This would be a situation similar to (36). But the case with (35) is different, because it asserts that Metatron is evil.

[11] Henceforth, we use italicized words as stand-ins for v-structures.

If *divineness* is chosen by the speaker, by implicature in some contexts the hearer will conclude that the speaker intended to communicate the property of goodness, not holiness, since had the speaker wished to specifically and unambiguously communicate the property of holiness, they ought to have chosen *divinity*.

In fact, the **MostInformative** constraints do not adjudicate the matter. In the relevant competition, between *divinity* meaning holy and *divineness* also meaning holy, **MostInformative**$_s$ returns $\perp$, since neither meaning is a proper subset of/strictly entails the other. There is no distinct f-structural information at play, so **MostInformative**$_f$ also returns $\perp$. Lastly, given the Pac-Man spanning we assume in *divinity* (see discussion below), neither *divinity* nor *divineness* is a better candidate for **MostInformative**$_c$, which also returns $\perp$. Thus, the *grammar* in fact does not adjudicate this case, only pragmatics.

The relevant Vocabulary Items are the following:

(39) $\langle\ [\sqrt{\ }],\ \Phi\{(\uparrow \text{PRED}) = \text{`divine'}\},\ \{\{\textbf{good} \mid \textbf{holy}\}\}\ \rangle \xrightarrow{\nu}$ *divine*

(40) $\langle\ [\text{n}],\ \Phi\{\ \},\ \{\ \}\ \rangle \xrightarrow{\nu}$ *-ness*

(41) $\langle\ [\sqrt{\ }, \text{a}, \text{n}],\ \Phi\{(\uparrow \text{PRED}) = \text{`divine'}\},\ \{\textbf{holy}\}\ \rangle \xrightarrow{\nu}$ *divinity*

Notice that (39) does not contain a category beyond $\sqrt{\ }$; in particular, it does not contain the category a. This is because there is evidence that this root is category-neutral. For example, *to divine the answer* does not mean to cause the answer to become divine, which is what we would expect if the VI for *divine* contained the category a.

Examples (42) and (43) show that if *divine* is interpreted as meaning **good**, which is inherited by *divineness*, then *divineness* can co-exist with *divinity*. In (42), we assume that *divine* spans [$\sqrt{\ }$,a] due to Pac-Man Spanning (Haugen and Siddiqi 2016, Melchin et al. 2020b), given that the VI for *divine* contains only a $\sqrt{\ }$ category.[12]

(42)

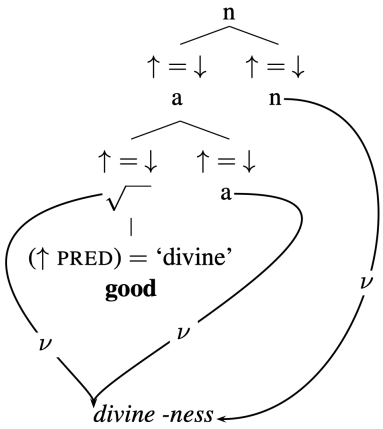

(43)

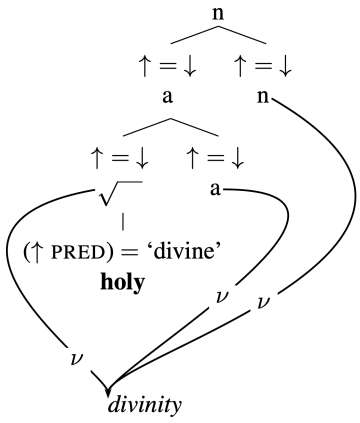

[12]We assume that Pac-Man spanning uses the same mechanism for comparing lists of categories as **MostInformative**$_c$. See the definition of **MostInformative**$_c$ and fn. 6 on page 10 above.

Goodness does not strictly entail holiness or vice versa (i.e., the good things are not a proper subset of the holy things, or vice versa). Therefore, **MostInformative**$_s$ has nothing to say about these two forms. Neither does **MostInformative**$_f$, since the f-structural contributions are identical. **MostInformative**$_c$ also returns $\perp$ in this competition, given the Pac-Man spanning of category a for *divineness*. Thus, both forms are allowed to be expressed by the grammar. However, since they mean different things, they do not create a pragmatic preference/blocking effect due to implicature.

5.2 *Unkempt/uncombed*

At first glance, the case of *uncombed* vs. *unkempt* seems parallel to *divine/divinity*. Indeed, this might be true for some dialects (such as Dan's!), for whom the meaning of *unkempt* entails the meaning of *uncombed*. These folks seem to be aware of the historical connection between the two forms. However, for most speakers of English, *unkempt* has a distinct root from *comb* (meaning its PRED feature is not [PRED 'comb']). Indeed, for these speakers, despite surface morphology, *unkempt* does not even contain negation:[13]

(44) $\langle$ [$\sqrt{}$, v, a], $\Phi\{(\uparrow \text{PRED}) = \text{'unkempt'}\}, \{\textbf{messy}\} \rangle \xrightarrow{\nu}$ *unkempt*

In this case, what we see here are two completely different c-structures: one which licenses the complex form *un-comb-ed* (45) and another that licenses the simplex form *unkempt* (46).

(45)

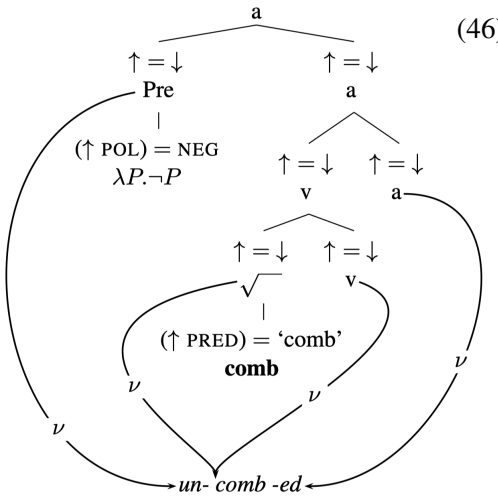

(46)

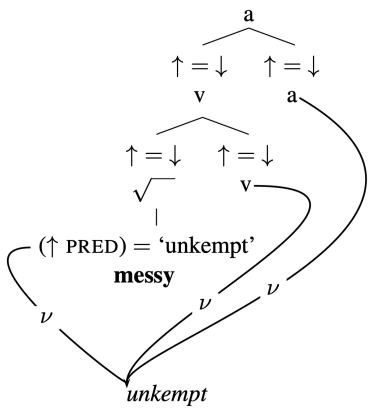

5.3 *People/persons*

Again at first blush, *people/persons* appears to be similar to *divinity/divineness*. While it is often argued that *people* is not actually a suppletive plural for *person* (see discussion in Siddiqi 2021), we set that debate aside and assume that *people* in fact does express [PRED 'person']. This case is particularly interesting: For

[13] We do, however, assume that speakers perceive *unkempt* as a derived adjective, as opposed to something like *blue*, which would have a simpler VI.

some speakers for whom *people* is the unambiguous plural of *person*, it is actually *persons* — the seemingly regular — which has specialized meaning! It appears only in highly formal contexts/registers.

(47) a. In cases of missing persons, the police search for missing people.

b. Indigenous people should contact the Indigenous Persons Bureau.

c. This room's capacity is 25 people, which is why there is a sign that says "Max 25 Persons".

This highly specialized meaning is evidence for the claim that *persons*, despite its seemingly regular morphology, is indeed a portmanteau (see Haugen and Siddiqi 2016). Further evidence for this claim is the fact that *persons* is legal in compounds (e.g., Missing Persons Department; see Siddiqi 2009 for discussion). Therefore, *persons* and *people* are in fact *both* portmanteau forms realizing the same c-structural and f-structural spans, as seen here in (48) and (49).

(48)

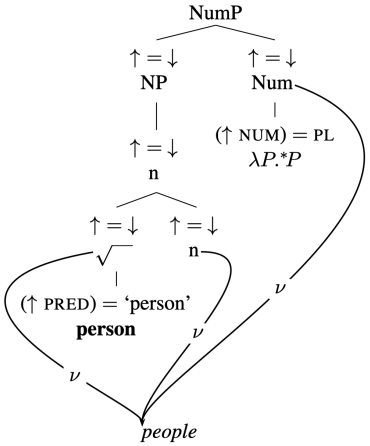

(49)

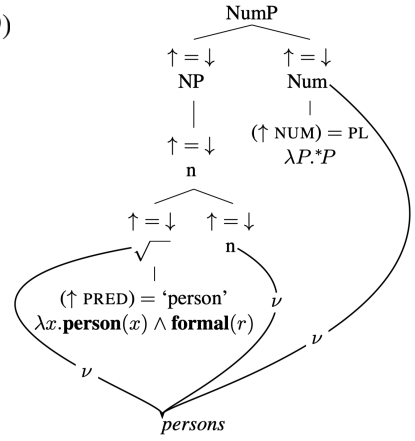

Thus, it is only **MostInformative**$_s$ that selects *persons* over *people*, and only in formal registers. We do not show the Vocabulary Items here, but they can be inferred from the c-structures in (48) and (49). We assume a mereological plural meaning, following Link (1983):

(50) $\lambda P.^* P$

We also assume, for the sake of explicitness, that the register variable, r, is a kind of Kaplanian contextual coordinate (Kaplan 1989).

5.4 *Brethren/brothers*

Let's lastly consider the case of *brethren/brothers*. Again, at first blush, we expect another *divinity/divineness* analysis. Instead we see that this requires a much more nuanced semantic and pragmatic account. Following Partee and Borschev (2003), we assume that a relational noun like *brother* involves a relation between the nominal entity and some other entity, such as a possessor. The meaning term for *brother* can be represented as follows:

(51) $\lambda y \lambda x \lambda R.\textbf{male}(x) \wedge R(x, y)$

Notice that, in an utterance where this is unresolved, the relational variable, R, is filled from context. In sum, (51) is the meaning term from the one *obligatory* meaning constructor for *brother*.

Of course, the relation **sibling** is always available in the null context. So we assume that there is a second, optional meaning constructor for *brother* whose meaning term modifies the term in (51) as follows:

(52) $\lambda \mathcal{R}.\mathcal{R}(\textbf{sibling})$

Thus, the interpretation of male sibling is available without context, but other interpretations are available if context and pragmatic knowledge supports them. In other words, as the term in (52) is optional, R in (51) can instead be instantiated contextually/pragmatically, for example as **close.friend** (where culturally appropriate, which is evidence of its pragmatic nature). Here are the meaning terms from the VI for *brother*:

(53) $\lambda y \lambda x \lambda R.\textbf{male}(x) \wedge R(x, y)$
($\lambda \mathcal{R}.\mathcal{R}(\textbf{sibling})$)

The optional meaning is thus available, and provides the interpretation in the null context. Alternatively, the pragmatic context fills in the R, such as in the case of **close.friend**. Indeed, *brother* can also be the singular of *brethren*, with the relevant meaning, as in the favoured reading, outside of other context, of a monk saying of another monk at the same monastery:

(54) My brother spoke out of turn.

In contrast, *brethren* obligatorily expresses the following relational meaning constructor in addition to the general meanings in (51) and (50):

(55) $\lambda \mathcal{R}.\mathcal{R}(\textbf{member.of.same.order})$

So *brethren* denotes the members of an all-male order. For speakers for whom the group must be a religious order, the meaning can be suitably further restricted.

We now have what we need to list the three VIs in this competition:

(56) $\langle$ [$\sqrt{\ }$, n], $\Phi\{(\uparrow \text{PRED}) = \text{`brother'}\}$,
$\{\lambda y \lambda x \lambda R.\textbf{male}(x) \wedge R(x, y), (\lambda \mathcal{R}.\mathcal{R}(\textbf{sibling}))\}$ $\rangle$ $\overset{\nu}{\rightarrow}$ *brother*

(57) $\langle$ [Num], $\Phi\{(\uparrow \text{NUM}) = \text{PL}\}$, $\{\lambda P.^* P\}$ $\rangle$ $\overset{\nu}{\rightarrow}$ *-s*

(58) $\langle$ [$\sqrt{\ }$, n, Num], $\Phi\{(\uparrow \text{PRED}) = \text{`brother'}\}$,
$\{\lambda y \lambda x.\textbf{male}(x) \wedge \textbf{member.of.same.order}(x, y), \lambda P.^* P\}$ $\rangle$ $\overset{\nu}{\rightarrow}$
brethren

In sum, as shown in (59) and (60), *brothers* is licensed because either the relationship is fully specified as sibling or appears underspecified, allowing for contextual specification of R. This underspecified R may resolve as sibling due to pragmatic forces, but it need not; it could resolve to close friend, among other possibilities.

(59)

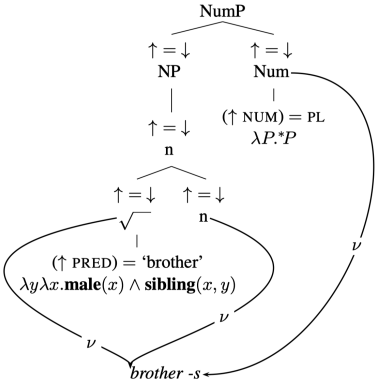

(60)

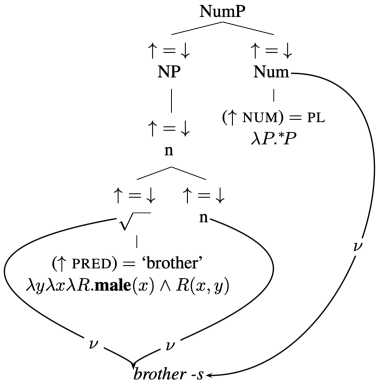

Brethren is disallowed in both (59) and (60) because of the absence of **member.of.same.order**. Thus, licensing of *brethren* fails despite the fact that **MostInformative**$_c$ would prefer *brethren* over *brothers*, because *brethren* is a portmanteau over Num. On the other hand, in (61), **member.of.same.order** is specified in the c-structure, so **MostInformative**$_s$ and **MostInformative**$_c$ together select *brethren* over *brothers*.

(61)

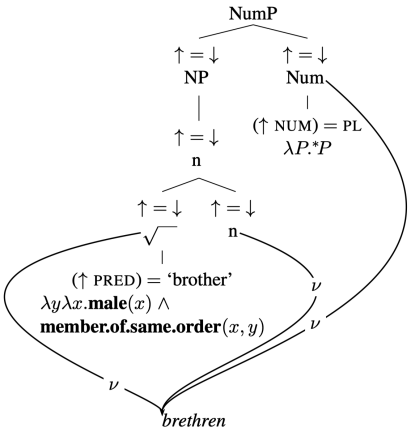

Since the $\sqrt{\ }$ node containing [PRED 'brother'] can come to have the meaning **member.of.same.order** through two means — overt specification and contextual specification — we make a correct prediction about morphosemantics here: The word *brothers* can be used with the same meaning as *brethren* when the meaning is contextually available, as when a monk might equivalently say (62) or (63).

(62) My brethren will make sure you are comfortable.

(63) My brothers will make sure you are comfortable.

However, the latter utterance could instead have other contextual meanings. Thus, if the monk wished to communicate specifically that the members of the order will ensure the addressee's comfort, *brethren* would be a better choice than *brothers*, because *brethren* has a more specific meaning.

6 Conclusion

Our goal in L_RFG morphosemantics is to use the actual compositional semantics to make morphological predictions. We use meaning constructors from Glue Semantics to accomplish this. This yields a nice result with respect to *locality*: meaning constructors are anchored to particular f-structures and thus only take scope over

their f-structural anchor.[14] We essentially get semantic locality for free: there simply is no question of being able to look "outside your domain" for a relevant feature, and therefore no need to place extra limits on processes for matching features and their probes, as in non-LFG-based DM. Our approach to capturing semantic specificity/information is akin to what may be familiar from event semantics: We leverage logical conjunction such that a term $\alpha \wedge \beta$ is necessarily at least as informative, and almost always more informative, than either α or β on its own.

In sum, L_RFG captures key ideas from the lexical-realizational framework of Distributed Morphology in principles like the **MostInformative** family of constraints. However, L_RFG formalizes these theoretical ideas in terms made available by the constraint-based framework of Lexical-Functional Grammar. Most importantly, since L_RFG has the capacity to use Glue terms in its Vocabulary Items, it can capture morphosemantics directly, in contrast to DM approaches built on the "Y model" (Chomsky 1981, 1995), in which syntax feeds LF and PF separately and there is therefore no direct interface between semantics and form. Lexical-realizational morphology has generally focused on problems of morphosyntax. We hope to have shown that the L_RFG approach to lexical-realizational morphology, with its new constraint-based tool kit, can also explain morphosemantics.

References

Andrews, Avery D. 1990. Unification and Morphological Blocking. *Natural Language and Linguistic Theory* 8(4), 507–557.

Aronoff, Mark. 1976. *Word Formation in Generative Grammar*. Cambridge, MA: MIT Press.

Arregi, Karlos and Nevins, Andrew. 2013. Contextual Neutralization and the Elsewhere Principle. In Ora Matushansky and Alec Marantz (eds.), *Distributed Morphology Today: Morphemes for Morris Halle*, pages 199–222, Cambridge, MA: MIT Press.

Asudeh, Ash. 2012. *The Logic of Pronominal Resumption*. Oxford: Oxford University Press.

Asudeh, Ash. 2022a. Glue Semantics. *Annual Review of Linguistics* 8, 321–341.

Asudeh, Ash. 2022b. Glue Semantics. To appear in Dalrymple (forthcoming).

Asudeh, Ash, Melchin, Paul B. and Siddiqi, Daniel. 2021. Constraints All the Way Down: DM in a Representational Model of Grammar. In *Proceedings of the 39th West Coast Conference on Formal Linguistics*, Cascadilla Press, forthcoming.

Asudeh, Ash and Siddiqi, Daniel. 2022. Morphology in LFG. To appear in Dalrymple (forthcoming).

Asudeh, Ash and Siddiqi, Daniel. forthcoming. *Constraint-based Morphosyntax*. Cambridge: Cambridge University Press, with Oleg Belyaev, Bronwyn Bjorkman, Tina Bögel, Michael Everdell, Paul B. Melchin, and Will Oxford.

[14] Scope-taking elements offer a more complex case, but they are still anchored in the relevant sense, because their Glue terms have the form $(X \multimap f) \multimap f$, where f is the f-structure that anchors the element's scope; i.e., f is the f-structure in which the scopal element serves as a GF. For more details see Asudeh (2022a,b).

Baayen, Harald. 1992. Quantitative aspects of morphological productivity. In Geert Booij and Jaap van Marle (eds.), *Yearbook of Morphology 1991*, pages 109–149, Dordrecht: Springer.

Bjorkman, Bronwyn, Cowper, Elizabeth, Hall, Daniel Curry, Hinds, Jennice, Koren, Louise and Siddiqi, Dan. 2021. Morphological Upstaging and 'Markedness'. In Angelica Hernández and Chris Plyley (eds.), *Proceedings of the Canadian Linguistic Association '21*.

Bresnan, Joan, Asudeh, Ash, Toivonen, Ida and Wechsler, Stephen. 2016. *Lexical-Functional Syntax*. Chichester, UK: Wiley-Blackwell, second edition.

Chomsky, Noam. 1981. *Lectures on Government and Binding: The Pisa Lectures*. Berlin: Mouton de Gruyter.

Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.

Dalrymple, Mary (ed.). 1999. *Semantics and Syntax in Lexical Functional Grammar: The Resource Logic Approach*. Cambridge, MA: MIT Press.

Dalrymple, Mary (ed.). forthcoming. *The Handbook of Lexical Functional Grammar*. Berlin: Language Science Press.

Dalrymple, Mary, Lowe, John J. and Mycock, Louise. 2019. *The Oxford Reference Guide to Lexical Functional Grammar*. Oxford: Oxford University Press.

Embick, David and Marantz, Alec. 2008. Architecture and Blocking. *Linguistics Inquiry* 39(1), 1–53.

Everdell, Michael and Melchin, Paul B.̇2021. Control the sentence, subordinate the pronoun: on the status of controlled versus non-controlled complement clauses in O'dam. In *Proceedings of the 39th West Coast Conference on Formal Linguistics*, Cascadilla Press, forthcoming.

Everdell, Michael, Melchin, Paul B., Asudeh, Ash and Siddiqi, Daniel. 2021. Beyond c-structure and f-structure: On the argument-adjunct distinction in O'dam. In Miriam Butt, Jamie Y. Findlay and Ida Toivonen (eds.), *Proceedings of the LFG21 Conference*, pages 125–145, Stanford, CA: CSLI Publications.

Girard, Jean-Yves. 1987. Linear Logic. *Theoretical Computer Science* 50(1), 1–102.

Halle, Morris and Marantz, Alec. 1993. Distributed morphology and the pieces of inflection. In Samuel Jay Keyser and Ken Hale (eds.), *The View from Building 20*, pages 111–176, Cambridge, MA: MIT Press.

Harley, Heidi. 2014. On the identity of roots. *Theoretical Linguistics* 40, 225–276.

Haugen, Jason D. and Siddiqi, Daniel. 2016. Towards a restricted realizational theory: Multimorphemic monolistemicity, portmanteaux, and post-linearization spanning. In Siddiqi and Harley (2016), pages 343–386.

Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry and Harvey Wettstein (eds.), *Themes from Kaplan*, pages 481–563, Oxford: Oxford University Press.

Kennedy, Christopher. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland.

Kennedy, Christopher. 2007. Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives. *Linguistics and Philosophy* 30(1), 1–45.

Link, Godehard. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In Rainer Bäuerle, Christoph Schwarze and Arnim von Stechow (eds.), *Meaning, Use, and Interpretation of Language*, pages 302–323, Berlin: Walter de Gruyter, reprinted in Link (1998: 11–34).

Link, Godehard. 1998. *Algebraic Semantics in Language and Philosophy*. Stanford, CA: CSLI Publications.

Melchin, Paul B., Asudeh, Ash and Siddiqi, Dan. 2020a. Ojibwe Agreement in a Representational, Morpheme-based Framework. In Angelica Hernández and M. Emma Butterworth (eds.), *Proceedings of the 2020 Canadian Linguistic Association*, Canadian Linguistic Association.

Melchin, Paul B., Asudeh, Ash and Siddiqi, Dan. 2020b. Ojibwe Agreement in Lexical-Realizational Functional Grammar. In Miriam Butt and Ida Toivonen (eds.), *Proceedings of the LFG20 Conference*, pages 268–288, Stanford, CA: CSLI Publications.

Merchant, Jason. 2015. How much context is enough? Two cases of span-conditioned stem allomorphy. *Linguistic Inquiry* 46(2), 273–303.

Partee, Barbara H. and Borschev, Vladimir. 2003. Genitives, Relational Nouns, and Argument-Modifier Ambiguity. In Ewald Lang, Claudia Maienborn and Cathrine Fabricius-Hansen (eds.), *Modifying Adjuncts*, volume 4 of *Interface Explorations*, Berlin: Mouton de Gruyter.

Ramchand, Gillian Catriona. 2008. *Verb Meaning and the Lexicon: A First Phase Syntax*. Cambridge: Cambridge University Press.

Siddiqi, Daniel. 2009. *Syntax within the Word: Economy, Allomorphy, and Argument Selection in Distributed Morphology*. Amsterdam: John Benjamins.

Siddiqi, Daniel. 2021. On the Taxonomy of Root Suppletion. In *Proceedings of the 39th West Coast Conference on Formal Linguistics*, Cascadilla Press, forthcoming.

Siddiqi, Daniel and Harley, Heidi (eds.). 2016. *Morphological Metatheory*. Amsterdam: John Benjamins.

Stump, Gregory T. 2001. *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press.

Svenonius, Peter. 2016. Spans and words. In Siddiqi and Harley (2016), pages 201–222.

Wescoat, Michael T. 2002. *On Lexical Sharing*. Ph. D.thesis, Stanford University.

Wescoat, Michael T. 2005. English nonsyllabic auxiliary contractions: An analysis in LFG with lexical sharing. In Miriam Butt and Tracy Holloway King (eds.), *Proceedings of the LFG05 Conference*, pages 469–486, Stanford, CA: CSLI.

Wescoat, Michael T. 2007. Preposition-determiner contractions: An analysis in Optimality-Theoretic Lexical-Functional Grammar with lexical sharing. In Miriam Butt and Tracy Holloway King (eds.), *Proceedings of the LFG07 Conference*, pages 439–459, Stanford, CA: CSLI.