Morphology and the Web of Grammar: Essays in Memory of Steven G. Lapointe

C. Orhan Orgun and Peter Sells (eds.)

September 25, 2003

**CENTER FOR THE STUDY
OF LANGUAGE
AND INFORMATION**

Contents

September 25, 2003

Preface

U. C. DAVIS LINGUISTICS FACULTY

Here we have an unnumbered introductory chapter with roman-numeral page numbers which will be the: Preface by U.C. Davis Linguistics Faculty.

Morphology and the Web of Grammar.
C. Orhan Orgun and Peter Sells (eds.).

Acknowledgments

The editors would like to thank all those involved.

1

When Morphology . . . Disappears

GREG CARLSON

1.1 Introduction

Steven G. Lapointe's work on morphology includes particularly keen attention to meaning, and the meanings that morphemes contribute to the sentence as a whole. But I'm going to focus on some instances where there is no morphology, or, rather, an absence of it. In the instances I'll be discussing, it is the absence, rather than the presence, of morphology that is correlated with certain patterns of semantic interpretation.[†] At least, this is the initial appearance. In the end, though, I will pose the question of whether this quite plausible idea is correct, and offer an alternative hypothesis whereby the presence of morphology is what makes for certain interpretive possibilities, and not the its absence. This is Lapointe's thesis.

Since the patterns of data are to be found across languages, there are certain assumptions that will need to be made. While a good many languages have undergone significant formal study (the morphology, phonology, and, to a lesser extent, the syntax), comparatively speaking only a few have received the detailed semantic study that is necessary to draw firm enough conclusions to be of value. Nevertheless, I am going to generalize from those well-studied instances, and piecemeal evidence that points in the same direction, draw-

[†]I wish to thank Peter Sells and an anonymous reviewer for comments which improved this work considerably. Explorations of this work were presented to audiences at the University of Washington, the University of Antwerp, the University of Norway, Hebrew University in Jerusalem, the University of Texas/Austin and at the University of Rochester. I have benefited enormously from a large number of audience comments. I also wish to thank Ariel Cohen in particular for his thoughtful comments and, more importantly, questions and criticisms. This material is based upon work supported by the National Science Foundation under Grant No. 0082928.

Morphology and the Web of Grammar.
C. Orhan Orgun and Peter Sells (eds.).

ing conclusions about the character of morphology and natural language in general which, given our present state of knowledge, must be regarded as tentative but, I hope, worth thinking about.

1.2 A brief overview of phenomena

The question posed in this paper surrounds the interpretation of, primarily, indefinite direct objects. It is well-known that indefinites may be interpreted specifically, and non-specifically, even if we are not clear on what that means exactly (for one attempted clear characterization, though, see Enç (1991)). The specific/nonspecific distinction is a subtle but omnipresent ambiguity with all (or at least most) occurrences of indefinites when they are interpreted nongenerically (no one to my knowledge has (yet) detected a similar divergence of interpretation among generically interpreted indefinites). However, there are certain construction types which appear to favor the nonspecific interpretation.

Consider initially data from English. English is a language with articles that normally must attend a count noun head in the absence of quantification. But, as everyone knows, the article may be omitted (or null) if the count noun appears in the plural.

(1) a. John saw a cat.

b. *John saw cat. (unless *cat* is interpreted as a mass term)

c. John saw cats.

(1c) contains what I am calling an absence of morphology. This absence is associated with a non-specific interpretation only, in contrast to (1a), which allows both a specific and a nonspecific interpretation. The non-specificity of the bare plurals is supported to a large extent intuitively, but it is also clear that they have narrowest possible scope (there are some well-known exceptions to this I set aside—see Chierchia, 1998), and no wide scope readings, a hallmark of nonspecific readings.

Another instance where there is a lack of morphology is the appearance of "bare singular" count nouns in languages with articles and plurality marking. In these instances a noun phrase is doubly deprived, so to speak, lacking both the required article and the plurality marking that is "normally" required of count nouns in the absence of an article. A good example of this is found in Norwegian (Borthen, 1998, 2003), which allows bare singulars fairly systematically:

(2) a. Jeg kjører *bil*.
I drive car
'I drive a car.'

b. Petter spiser helst med *skje*.
Petter eats rather with spoon
'Petter would rather eat with a spoon.'

c. Jeg har bestilt *billett*.
I have ordered ticket
'I ordered a ticket.'

As Borthen clearly shows, these bare singulars have the semantic properties of bare plurals in terms of being uniformly non-specific weak indefinites; there are no wide-scope or specific readings for them. Bare singulars in other languages with articles and plurality marking have been shown to have an identical semantics, as in Albanian (Kallulli, 1999) Brazilian Portuguese, investigated in detail by Schmidt and Munn (1999) and Munn and Schmidt (1999), and Hungarian (Farkas and de Swart, to appear).

Besides a lack of determination and plurality, a lack of case-marking (in languages with systematic case-marking) also appears to be associated with weak, nonspecific indefinite readings of noun phrases. Enç (1991), for instance, goes to some lengths to show that when Turkish objects appear without case-marking, the result is uniformly nonspecific. Hindi has a similar construction. The following examples are from Mohanan (1995):

(3) a. ilaa baccõ-ko k^{h}ojtii rahtii hai
ila-Nom children-Acc search-Hab Prog be-Pres
'Ila keeps searching for the/some children.'

b. ilaa bacce k^{h}ojtii rahtii hai
ila-Nom children-Ø search-hab Prog be-Pres
'Ila keeps children-searching (i.e., performing the act of searching for children).'

The lack of the "usual" accusative marking on the object noun (which Mohanan initially characterizes as null nominative marking) is associated with only the indefinite, weak, narrow-scope reading. These interpretations are further elucidated in Dayal (1999).

Another instance of "lack of morphology" found in the literature is the study by Cheng and Sybesma (1999) regarding Cantonese, a language without case-marking, plurality, or articles. Cantonese, like Mandarin and many other languages, has classifiers. However, in Cantonese these classifiers may on occasion be omitted, to the effect of eliminating wider-scope specific or definite readings that may appear in the presence of the classifier:

Perhaps the most extensive body of literature which illustrates the kind of "missing morphology" focused on here is the vast amount written on noun-incorporation. I make no attempt to summarize the literature here (though see van Geenhoven, 2002, and Gerdts, 1998 for excellent overviews of the literature and pertinent issues). As van Geenhoven (2002) points out, "Noun

incorporation is a cover term for a wide variety of constructions across languages . . . " (p. 261). However, what is a characteristic of such constructions is that when the noun is "incorporated", any attendant morphology is normally stripped away, including case-marking, plurality and articles, in privative comparison to the fuller syntactic expression of the same meaning. Detailed semantic investigation of the semantics of incorporated constructions yields the observation that only weak, non-specific, indefinite readings may be associated with an incorporated nominal (e.g., van Geenhoven (1998) and Bittner (1994) for West Greenlandic). These detailed observations are regularly supported by comments of varying degrees of explicitness throughout the incorporation literature, whether theoretical or descriptive, that the semantics is of the same character. A few quotations from works on specific languages gives a flavor for this:

> (When incorporated,) " . . . nouns become non-referential" Sullivan (1984)
>
> "[the incorporated form] refers to habitual, permanent, chronic, specialized, characteristic or unintentional activities or states, or localized events, with the noun being generic, nonreferential, or indefinite." de Reuse (1994)
>
> "Incorporation is used when objects are non-specific and non-salient." Spencer (1985)
>
> "An incorporated noun often refers to a generic or unspecific class" Mithun (1984)

Such observations may be multiplied nearly at will. However, it remains unclear where the outer boundaries of "incorporation" lie. For instance, the Scandinavian and the Hindi examples cited above may be analyzed as instances of incorporation, and Mohanan explicitly argues this to the case for Hindi in her article, as do Asudeh and Mikkelsen (2000) for Danish. But Massam (2001) argues that there is a category of "pseudo-incorporated" nominals, based on evidence from Austronesian. The outer boundaries are uncertain (one could even look at English bare plurals as "incorporated" were one of a mind to do so), but there are hints which indicate that not all instances of missing nominal morphology with the requisite semantics should be analyzed as "incorporated".

For example, in some languages there is the phenomenon of "doubling", where an incorporated nominal interpreted as, say, the direct object of the verb, may appear in the presence of another (apparent) direct object:

(4) Kikv rabahbot wa-ha-its-a-hnini-'ki rake-'niha
this bullhead FACT-M.sg.SUBJ-fish-buy-PUNCT my-father
'My father 'fish-bought' this bullhead.' (Mohawk, Baker, 1995)

Whether the unincorporated "doubling" direct object is an argument, as commonly assumed, or an adjunct, as argued by Chung and Ladusaw (in

press) for Chamorro, is presently not the point. Rather, the observation is that doubling, when it occurs, may only occur with truly incorporated nominals. Lack of doubling is not a diagnostic for lack of incorporation, but such phenomena suggest that not all instances of an absence of morphology should be analyzed as incorporated.

There are two other instances of "missing morphology" that I wish to suggest are also a part of this general pattern, though they have received less documentation in the literature. One instance is that of a lack of direct-object agreement markers on an apparently transitive verb, in languages which have direct object agreement markers. For instance, Perrot (1972) reports that in Swahili, verbs agree with definite objects but not indefinite ones.

(5) a. U- me- leta kitabu?
2sg Perf bring book
'Have you brought a book?'

b. U- me- ki- leta kitabu?
2sg Perf 3sg bring book
'Have you brought the book?'

Aissen (2003) takes another point of view, arguing that tense is the controlling feature, so we must look again at this phenomenon, at least in Swahili. However, another somewhat more secure is example is drawn from Palauan, though here the pattern is more complex. Woolford (1996) reports that in the perfective aspect, with non-human objects, verbs agree with only with specific singulars, and not with indefinites, or indefinite plurals. In the imperfective paradigm, verbs do not agree with any objects. However, when an object is non-human, specific and singular, a preposition is inserted before the definite singular NP. So while there is more going on here, the connection between lack of agreement and indefiniteness, and between NP marking and specificity, is in evidence. What is not so clearly established in this literature is whether the indefinites are non-specific only.

The other phenomenon is "clitic-doubling", illustrated by the Spanish example below. A lack of such doubling is associated with nonspecific indefinite readings. Suñer (1988) notes two things in Spanish. First, with bare plurals, one cannot get clitic doubling—this appears to be the same with Albanian bare singulars, discussed in Kallulli (1999).

(6) Les dejaré todo mi dinero a los pobres.
To-them I-will-leave all my money to the poor
'I will leave all my money to the poor.'

(7) *Les donaré todos mios bienes a museos.
To-them I-will-leave all my money to museums

Also, clitic doubling cannot occur with nonreferential definites:

(8) *Lo alabrán al niño que termine primero.
Him they-will-praise the boy who finishes first
(bad on the reading where winner is yet to be determined)

There is one vital issue that I have left unaddressed in the above, which will largely remain that way. This is the question of whether "missing morphology" as characterized is truly the absence of morphology, or whether such instances are actually exemplars of null morphology—that is, actual morphology with phonologically null expression. So, for example, the fact that Slavic languages, and many others, generally lack articles is not something I would consider "missing morphology". Rather, the strategy here is to examine morphology contrastively; that is, if something that looks like it "normally" takes a certain type of morphological marking, and in certain instances lacks the expected marking, I am taking it as "missing" rather than as "null". This assumption may misfire in some instances (e.g., whether the Hindi objects mentioned above are null Nominative, rather than genuinely without case-marking), but I am going to assume that at least the bulk of the example are genuine instances of a lack of morphological marking. For a recent review of some of the phenomena discussed here, though with a very different set of concerns in mind, see Aissen (2003).

Let me summarize the discussion to this point. We have seen a series of phenomena that are associated with a certain type of semantic reading for, chiefly though not exclusively, direct object noun phrases. The few illustrative examples given are indicative of much more general patterns that can be found in a wide variety of unrelated langauges (see Farkas and de Swart for a listing of "stable" semantic patterns). In these instances, it is the lack of morphology that is associated with this nonspecific reading, and not its presence. So, the question is, do we have here an instance of **no** morphology bearing meaning? If this were so, I'd take this as contrary to Lapointe's view of morphology, where it is forms that contribute meanings, and not the absence of form. In the remainder of this article, I am going to present some ideas about what might be going on in the instances reviewed, and, if the ideas have merit, we are going to see that it is the presence, and not the absence of forms that bears the required meanings—or rather, the required functions.

1.3 Three levels of meaning

The point of view taken here is sketched out in Carlson (2003a), which is an attempt to deal with the question of why weak, nonspecific indefinite noun phrases are curiously the only instances in Diesing's (1992) framework which are not required to be raised out of the Verb Phrase and into the IP domain (to use Diesing's terminology). The account to be given, in contrast to Diesing's, is a semantic account, and not a syntactic one: it is just that weak non-specific

indefinites have a type of meaning which is "compatible" with verb phrase meanings, all other types of noun phrases being uninterpretable within that domain.

Recall that in the discussion above, the noun phrases were divided, implicitly, into not two but three classes: (1) truly incorporated weak nonspecific indefinites (2) nonspecific indefinites that are not incorporated. and (3) all others, generally perceived as "regular" noun phrases. It is only this last class that is generally associated with fullest expression of morphology, the incorporated instances regularly lacking the morphology. The unincorporated instances may, or may not be associated with the fullest expression of morphological marking, depending on the language. Nevertheless, there is no clear semantic distinction between the first two classes (incorporated, and unincorporated nonspecific indefinites)—were there a clear difference, one could argue for or against syntactic/lexical incoporation on the basis of meanings along with the formal arguments, but no one as yet has found a secure way to do this. I think the following will help us understand why.

1.3.1 Overview

At the most general level, meanings of sentences featuring an episodic verb are built up endocentrically from the verb (other major phrases may have a similar organization, Carlson, 2003b)). At the lowest level, there are lexical meanings of verbs. These are meanings that are associated with words, and not with phrases. By "word" I intend a descriptive term which means a lexeme which may exhibit productive derivational morphology, but excludes clitics and inflectional material that may appear with it, often obligatorily, within a sentence. It has long been noted that (at least monomorphemic) lexical meanings differ in some important ways from the kinds of meanings that full sentences may express. One important intuition is that lexical meanings are "universal", that is, they make no implicit reference to particular individuals or states of affairs. It is only when lexical items are situated in full syntactic constructions that reference to particulars emerges. This is even the case for proper names. The meaning of a noun like "Alice" seems to be something like "the property of being named 'Alice"'; the meaning of the noun phrase (or DP) "Alice", on the other hand, makes reference to some particular person. Similarly, a verb like "sneeze" delivers only a type of action, whereas in a sentence like "Alice sneezed" it becomes a component of an expression which (plausibly) makes reference to a particular act of sneezing.

One level up from the lexical category V is the level of VP (or V′) meanings. These are meanings that are associated with phrasal expressions, and not lexical expressions. That this level has distinct constraints on its meanings is not so readily observed, yet it is the cornerstone of the reinterpretation of Diesing's Mapping Hypothesis outlined in Carlson (2003a). The one area

of the literature where this phrasal intuition becomes apparent (though not in the phrasal vs. sentence-level distinction I am according it here) is in work on "categorical" vs. "thetic" judgments (e.g., Kuroda, 1972, Ladusaw, 1994, Basilico, 1998). However, the guiding intuitions there do not line up exactly with the way it is approached here. However, I am taking it that the phrasal level, too, shares with lexical items a lack of reference to particulars. It is at this level where "nonspecifics" can be found—they are nonspecific precisely because they occur at a level of interpretation where, as with lexical items, specific reference is not (yet) possible. Constituents of this level do not include such elements as tense, modality, speech act information, sentential negation, quantification, or reference to particulars. Also lacking at the phrasal level are the elements that make a sentence generic. The putative operator GEN (Krifka et al., 1995) is a part of sentence-level meaning, not phrasal level meaning.

At the highest level are IP or sentence meanings. The sentence-level meanings are what truth-conditional semantic theory is designed to represent. These meanings are assumed here to be of the structure usually presented (whether it is as presented in Chierchia and McConnell-Ginet, 2001), or the dynamic semantics presented in Kamp and Reyle, 1993). The key point is that sentence-level meanings are built up around notions of reference, truth, and predication, with its chief calculus being that of relations and entities standing in those relations, with quantification and intensionalizing elements. However, the lexical and phrasal levels are defined without reference to any such notions.

1.3.2 Verb meanings

The types of meanings associated with verbs may, to those who have examined verb meanings in detail, appear rather disappointing. Things will hopefully become at least somewhat less disappointing as we proceed. The basic idea behind the semantics of episodic verbs is that they denote an "event type" (or, an "eventuality", Bach, 1986, as this includes episodic verbs of all Aktionsarten, not just the subcategory of telic eventualities commonly also called "events"). This denotation is not found in the domain of actual events, and this denotation is neither true nor false. Event types do not have "argument slots" in them. Perhaps the best way to think about them is that they represent an understanding of how to perceptually classify something encountered as an event of a certain sort, being it a running, a jumping, or a pouring. But in the formal theory it is just an event type. Let us then define a set Σ as that set, with its members being E_1, E_2, $\ldots$. If V is a verb, the interpretation of that expression, designated by $[\![V]\!]$, is going to be some member of Σ.

The members of Σ collectively define an upper lattice, or a complete join semilattice, based on the "part-of" relation $\leq$ (its converse is $\geq$). Intuitively,

an event-type which is more specific than another event type is related by $\leq$. Thus, for instance, ⟦run⟧ $\leq$ ⟦move⟧, ⟦shatter⟧ $\leq$ ⟦break⟧, and ⟦state⟧ $\leq$ ⟦communicate⟧, as the first are more specific instances of the latter. On the other hand, ⟦bake⟧ and ⟦nudge⟧ are not so related (even if, in a *specific* act of baking, one might nudge the pan into the oven). This $\leq$ relation is intended to mean that the first member is automatically redescribable as the second under all circumstances. So, while a part of an act of building, for instance, might be an act of sawing, one cannot automatically redescribe an instance of sawing as a building. This is so even if one, in a perceptual sense, identifies an act of building by noticing an act of sawing. The limit to this is when a specific event type appears necessarily a component of accomplishing another. For instance, suppose that playing chess always involves thinking; even so, it is not the case that ⟦think⟧ $\leq$ ⟦play-chess⟧, as any act of thinking is not redescribable as playing chess. Rather, this relation between thinking and chess-playing are going to appear in the propositional semantics.

One might note that 'play-chess' is not (formally speaking) a member of the category V in English. This is because, I assume, English does not have object incorporation. But many other languages do have incorporation, and things like 'play-chess' *can be* among the lexical items of the language. A verb like this is obviously some composition of the meaning of 'play' and the meaning of 'chess', and is not atomic. Thus it is necessary to say something about how the two are combined at this, the lexical level.

The meaning of a noun is what we might call informally a 'property', drawn from a set of properties Π, each member of which is P_1, P_2, ... ; there is no need to assume that Π and Σ are disjoint, but for the sake of simplicity I will assume it for the time being. Like the verb meanings, these also form a semilattice, so that ⟦cat⟧ $\leq$ ⟦mammal⟧ $\leq$ ⟦animal⟧, and so forth. And while ⟦cat⟧ does not stand in the '$\leq$' relation to ⟦dog⟧, both stand in the '$\leq$' relation to ⟦mammal⟧ and to ⟦animal⟧. At this, the lexical level, there are no expressions which denote individuals; only the universals are represented here. If one represents a proper name such as 'Bob', ⟦Bob⟧ at this level of interpretation is going to represent the common-noun usage of the name only (e.g., the property of being named Bob), and not the usage making reference to a particular individual bearing that name.

Now, suppose we have an incorporation construction, which I take to be a lexical meaning, of the form N V. What is ⟦N V⟧ then? Intuitively, it is going to be a more specific event type than that denoted by V alone. Thus, ⟦N V⟧ $\leq$ ⟦V⟧. Thus, eating cake is going to be a more specific instance of eating. There is one constraint on this derived verb meaning. Namely, that if the V combines with a noun N that is related by $\leq$ to another N′, then ⟦N V⟧ $\leq$ ⟦N′ V⟧. So, again, eating cake will be a more specific event type than eating food.

So, the lexical meanings are defined, at this level, solely by their place in

the lattice of types and properties in relation to other lexical meanings of the same syntactic category. This would seem to miss something crucial about lexical meaning such as what the event-type denoted by "jump" has to do with the action of jumping. I am going to say nothing about the perceptual processes involved in classifying something perceived as a jumping; what I'm modeling here is the outcome of such processes, a categorization arrived at—in the case of major lexical items typically on the basis of encyclopedic knowledge (Borer, in press). But the other reason is that we are not finished yet, and at the sentential level such a connection will be made.

The approach to meanings of major lexical items taken here is reminiscent of, and to a certain extent, consistent with, ideas that also appear in the syntax literature which take at least roots of lexical items as noncategorical (see for instance, Marantz, 1997; Pesetsky and Torrego, in press; and Borer, in press). Assuming that "noncategorical" means that these elements cannot enter into the combinatory syntax, this is reflected in the assumptions made here about their semantics. The properties of nouns, and the eventualities of verbs are, as far as the semantics and syntax are concerned, atomic, and lack any properties at this level which allow them to combine and recombine with other elements of the syntax.

So, one is left wondering exactly how the N and the V are specifically combined in the incorporated structure illustrated above if $[\![\mathrm{N}]\!]$ and $[\![\mathrm{V}]\!]$ lack combinatorial properties. So far all we have done to this point is to identify a new node in the lattice of Σ, and localize it with respect to other members of Σ. We have not identified the exact mode of combination. That is because, I am claiming, this "mode of combination" cannot be specified until we go one level up—to the phrasal level, where the elements begin to be syntactically combined.

1.3.3 The V′ level

The immediate projection of the verb is the V′ level. This is the level at which the arguments are added, at least those that appear in situ in the verb phrase. In the interest of simplicity, I am simply going to assume that all such arguments of the verb are syntactically present at this one level—that is, we are not necessarily restricted to binary branching. If one prefers a binary branching or a more articulated V′ structure, a similar semantics can easily be provided. I am also assuming that a VP-internal subject may be present. An understanding of what types of elements may appear at the V′ level may be gained from an examination of Diesing (1992). In her system, though many types of noun phrases may appear to be surface constituents of the V′, at "logical form" (and in some instances, also at the surface structure), such noun phrases as specific indefinites, quantified noun phrases, proper names, and definite noun phrases, are not a constituent of the V′, but must be "raised" to a higher posi-

tion in the syntactic representation, the IP level. What can remain (it appears) as a constutuent among the types of noun phrases is the sole instance of weak, indefinite nonspecifics.

The mode of combination at the V′ level depends upon a verb associating with a series of thematic roles (Th_1, Th_2, ...), and these thematic roles in turn being "assigned" to the syntactic arguments. A part of what it is to be a "verb" in this view is that it is inserted under a V′ node that has an appropriate listing of thematic roles associated with it, which are listed in the lexicon with the verb. Most verbs have several alternative sets of roles, just as they have several alternative subcategorizations. So if a verb V takes thematic roles #4 and #7, whatever they might be, then it can be inserted under a V′ thus:

(9)
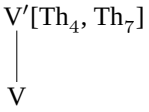

These thematic roles are associated with arguments—noun phrases and prepositional phrases, in the main—via a direct linking. Graphically, we might put it this way, though coindexing would have the same effect:

(10)
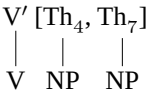

To use a more concrete example, the sentence "A cat saw a mouse", at this level, would have the verb "see" associated with, say, an Experiencer role (Th_5) and a Stimulus role (Th_6), intuitively encoding the way the participants participate in the event. A preferable, but more complex representation would put the thematic roles in terms of Dowty's (1991) theory of proto-roles, but I am again going to assume a set of discrete roles for the sake of simplicity. Following suggestions in Carlson (1984), the thematic roles are construed as functions on NP meanings to return sets of events—the set of events that (the meaning of) the NP "participates in" in this particular way.

The semantics given above for both verbs and nouns at the lexical level lacks the necessary structure for accomplishing this. We need to introduce something like token events, on the one hand, and something like individuals and groups of individuals, on the other (as weak indefinite NP's have forms like "several books" and "twelve people"). I am going to generally follow algebraic definitions found in property theory (e.g., Bealer, 1982) which provide a means of combining properties with one another and defining outcomes of those combinations. This is accomplished via reference to their extensions, or instantiations.

The extension of a V meaning, ˇ⟦V⟧, is going to be some set of event-type instantiations; I'll call these "token events", though bearing in mind that these are not real events of the world, at this level. Let $\mathcal{E}$ designate a set of token events $\{e_1, e_2, \dots\}$ such that for any V, ˇ⟦V⟧ is a subset of $\mathcal{E}$. These subsets preserve the lattice structure of the V meanings in the usual way (e.g., if one verb meaning is a part of another, then its extension will be a subset of the other's extension, etc.). Now, as I have said above, $\mathcal{E}$ does not consist of what we might think of as real events; it is rather more like the set of event-like things we can talk about. It does not matter that, for example, a mouse has never in fact devoured an elephant for that to have token event instances, since it is something we can talk about even if it never has happened, or ever will happen. ˇ⟦a mouse devour an elephant⟧ will be some subset of $\mathcal{E}$ just as surely as ˇ⟦a dog eat some food⟧. (In these examples, I intend the nonspecific readings of the indefinite noun phrases only.)

The extensions of N meanings are provided for at the NP (as opposed to DP) level, where ˇ⟦NP⟧ = ˇ⟦N⟧ in the instance where an NP consists only of an N (i.e., there are no modifiers or other elements combining with the N). It may seem that the natural thing to do would be to have these extensions be sets of individuals, but the problem with this is that whether a given individual falls under the extension of a noun is contingent. It depends on the particular time and place (and world) one is in. So, for instance, whether a particular person is a "student" or a "professor" obviously depends on the time of evaluation—the same person may be one at one time, the other, at another. So the best we could do in defining an extension for a given noun N is to think of the extensions as being individuals-while-they-are-an-N. But such objects are not individuals as the term is commonly construed. At this point in the interpretation, the "individuals" are just property-instantiations that can be distinguished from one another. Thus, corresponding to the domain of Π, there is a set of instantiations of noun-meanings which we will call $\mathcal{N} = \{i_1, i_2, \dots\}$ such that for any noun N, ˇ⟦N⟧ is some subset of $\mathcal{N}$. The extensions set-theoretically preserve the lattice-structure of noun meanings found in Π. Further, I assume that the elements of $\mathcal{N}$ themselves form a lattice structure which allows reference to pluralities (i.e., groups).

The thematic roles operate on these extensions. They have two functions: (1) to map each individual onto the token events it participates in, in the particular way defined by the thematic role, and (2) to intersect that result with the set of token events defined by the extension of the verb. The result of this combination, given the constraints mentioned above, will be a more specific event-type, i.e., an event-type whose extensions will be some subset of the extensions of event-type defined by the verb alone. Let's use a simple example. Suppose the intransitive V "freeze" is inserted into a one-argument $V'[Th_2]$, with Th_2 being "Patient". The extension of the N "plants" at the NP level is

going to be a subset of $\mathcal{N}$. Th_2(⟦plants⟧) is going to designate the set of token events in $\mathcal{E}$ in which a plant (i.e., some member of ˇ⟦plant⟧) functions as the Patient. This is going to include many non-freezing events (as when a plant blooms, is potted, etc.), some freezing events where it is a plant that freezes, but it will not include all freezing events (where some non-plant freezes). Intersecting this set with the set of freezing events is going to yield just that set of freezing events where some plant or other functions as the Patient.

Note that, at this level, the result is just a set of events (which corresponds to an event-type). The verb at this level is not a function; it does not have any argument "slots" of the type V(x, y), and the meaning of "plants freeze" is neither true nor false. It simply designates an event-type. But what we do have now is the machinery to go back and define "mode of combination" for an incorporated structure (let's say of the form [N V]).

A preliminary comment before illustrating. First of all, languages vary in terms of what arguments may be incorporated. All, to my knowledge, allow for patient-like direct objects to be incorporated, and while some allow (for instance) instrumentals, others may not; some allow unaccusative subjects, others may not. There is a good deal of linguistic variation to be found. Further, as a given language may allow for different alternatives to be incorporated, one cannot select a unique "default" relation and apply that at the lexical level. So what I am going to do is to select the simplest case and illustrate, that of direct-object incorporation.

Recall that the meaning assigned to something like ⟦fish buy⟧ is a property which finds its place in the lattice-structure of event-types. All we know about it is that it is more specific than ⟦buy⟧, and would be more specific than ⟦creature buy⟧ and less than ⟦bullhead buy⟧. Now, at the V′ level, we need to define ˇ⟦fish buy⟧. Let us assume (as is the case in the Mohawk example (4) above) that "fish" is interpreted as the Theme, rather than, say, as the Instrument. What we need to say is that ˇ⟦fish buy⟧ is going to be the same as Th3(ˇ⟦fish⟧) intersected with ˇ⟦buy⟧. Here is the schema:

In a construction of the form:

(11)

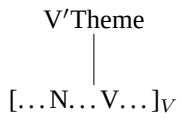

(where '. . . ' indicates possible morphemes),
ˇ⟦V′Th$_n$⟧ = ˇ⟦V⟧Theme(ˇ⟦N⟧)

This is going to be exactly the same result as the syntactic construction "buy (a) fish".

Actually, it is going to be the same result in terms of the extensions. But the result is not quite the same. This is because "fish-buy" is going to be guar-

anteed to denote a node in the latticework of V meanings in Σ and what this operation does is to define the extension of that node. On the other hand, the unincorporated form (e.g., "buy a fish"), while it does have the same extension, does not find any denotation in Σ. Quite simply, this is because it does not correspond to a lexical V meaning. It will turn out that the incorporated form and the corresponding full syntactic expression are truth-conditionally equivalent, so one can have no fish-buying without also having bought a fish (and vice-versa), so we need to look beyond truth-conditional meaning to see if there are any differences this corresponds to. In fact, I believe there is a commonly-expressed intuition in the literature that supports this non-truth-conditional distinction. Consider the intuition that the following authors are trying to express:

> " ... incorporation provides the lexicalized expression of a typical activity." (Axelrod, 1990)
>
> "[the incorporated form] refers to habitual, permanent, chronic, specialized, characteristic or unintentional activities or states, or localized events, with the noun being generic, nonreferential, or indefinite." de Reuse (1994)
>
> "Incorporation is used when the event is of greater interest than the participants." Spencer (1985)
>
> "...the noun no longer refers to an individuated specific or unspecific participant, and thus the whole clause shows a lesser degree of transitivity." Mosel and Hovdhaugen (1992)

The intuition is often expressed as saying that incorporated structures express a "typical activity", as if the activity expressed has achieved a certain degree of expectedness, or familiarity. Mithun (1984) uses the terms "institutionalized" activity and "atomic" concept. If we think of lexical meanings of verbs as encoding "typical activities", and phrasal, periphrastic expressions as lacking that status, then we have a formal means of expressing this distinction.

We also have a way of expressing the combinatory "gappiness" of incorporation. This is the fact that, at least in some languages, certain verb-noun combinations form impermissible incorporated structures, despite the fact there seems to be no discernible formal barrier. For instance, as Dayal points out for Hindi (assuming that the structure is indeed incorporation), one can "girl-see" but not "woman-see". Such restrictions could be encoded as the absence of any node in Σ in one instance, and the presence in another. It does leave the question of how "typical activities" are established (and not established) for a given language, but at least it situates the question into the broader question of lexicalization in general.

1.3.4 The S level

Interpretation at the V and V′ levels makes no reference to truth, context, times, worlds, or any of the familiar semantic notions aside from "denotation". However, at the S (or IP) level, all these elements come into play. Here, the representations are as commonly assumed. There is at least a set of individuals $\mathcal{A}$, a set of worlds $\mathcal{W}$, a set of events $\mathcal{E}$ (mnemonic for "actual events") with $e_1, \ldots$ as members, and a set of ordered times $\mathcal{T}$. I am going to talk about worlds as containing events as parts, but this is a matter of convenience. Verb meanings are (or rather, correspond to) functions of the sort nearly universally assumed in formal semantics, with argument slots in them, expressing relations among the denotations of its arguments to yield truth-values and define propositions via these truth-values (e.g., a proposition being a set of possible worlds, that set of world at which the evaluation of the relation and its arguments yields True). So what we need then is some means of mapping from the structures present at the V′ level into structures at the propositional level.

We need to ensure that the meanings expressed at the V′ level have the intuitively correct denotations in the propositional model. We do not wish for "dog" to denote a type of fish, or "run" to denote events of eating, for example. We accomplish this by mapping into $\mathcal{A}$ and the elements of $\mathcal{E}$ from the members of $\mathcal{N}$ and $\mathcal{E}$ respectively (i.e., the property-instantiations and the event-instantiations) *with respect to the parameters of evaluation* (e.g., times and worlds). So, for instance, we might say that a given individual, a, corresponds at world 7 at time 44 to *that* particular property instantiation in $\check{}[\![\text{student}]\!]$, etc. The event-instantiations in $\mathcal{E}$ are functionally mapped to events from $\mathcal{E}$ that appear in possible worlds. It is certainly the case that, different event- and property- instantiations may be to be mapped to the same individual or event. So, for instance, some snakes are also pets (but not all), and some instances of making a mark with a pencil are also cases of writing (but not all). This mapping from event-types and properties to elements of the world will (finally) ensure that "cat" at the lexical level will, ultimately, correspond to actual cat-creatures in the propositional model (or rather, any sentence "cat" appears in has truth-conditions that necessarily involve reference to some feline creature). The notion of "individual" finally appears at this level. An individual is something that exemplifies any number of property-instantiations at a given moment, and these properties may come and go, and the individual may remain the same individual.

Recall that the structure of the V′ consisted of relating argument NP's to an event-type denoted by a V via thematic roles. However, at the propositional level, the verbs will correspond to n-place functions which take arguments and yield truth-values. As mentioned above, however, not <u>all</u> arguments of

the verb appear at the V′ level—it is only the weak indefinite NP's that appear there; such things as referential definites, quantified NP's, pronouns, proper names, and so forth, do not. They are only interpretable at the propositional level. So, for instance, at the V′ level, a simple sentence like "John saw Mary" will only have the event-type of a "seeing" event (i.e., what is given to us in terms of extensions by the verb alone). Thus we must state three correspondences. (1) We need to make sure that John and Mary are placed in the propositional relation **see**(x,y) appropriately (i.e., in "John saw Mary" it would be **see**(j,m) and not **see**(m,j) (2) we need to make sure that they are interpreted as having or at least entailing the same thematic roles as weak indefinites would in the same argument positions (keeping the meaning of the verb itself constant). And lastly we need to go back and say something about what happens if there is a thematic role at the V′ level which is not linked to any NP or DP at that level—what is the interpretation then? We set out to do these things in several stages.

The task is not to give the lexical item, the verb, some kind of functional meaning. Its meaning was given at the lexical level. Rather, it is to give the event-type designated by the V′ a functional meaning. Let us begin with a case where the thematic roles associated with the V′ are "filled":

(12)

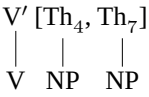

Let's suppose the V is "chase" and the NP's are "a cat' and "a mouse". This denotes the event-type of a cat chasing a mouse. This will intuitively correspond to a proposition, a function from worlds to truth values. If we apply the following, this achieves the desired result:

First, the propositional relation expressed by the verb needs to be defined. Suppose a V′ has the representation:

(13) If $[\![V']\!] = \{e_n \ldots e_q\}$,
then the interpretation of $[_S[\![V']\!]] = \lambda w\, \exists e \in [\![V']\!]$ & $f(e) \leq w$
(i.e., the propositional interpretation of an event-type is that set of worlds in which there occurs some event of that type; the '$\leq$' symbol here is used for convenience; f(e) is the mapping from event instantiations to members of Æ)

This will convert the event-type of cat-chasing-mouse to the proposition where in each world a cat chases a mouse, and does not in the remaining worlds. This proposition is equivalent to the one also defined by:

(14) $\exists x \exists y$ (cat(x) & mouse(y) & chase(x,y))

This is because each thing that is a cat and each thing that is a mouse finds

a corresponding property-instantiation in ˇ⟦cat⟧ and ˇ⟦mouse⟧ which "participates" in the chasing event. The same will hold for all extensional verbs. However, this is a different way of arriving at the same truth-conditionally defined proposition. I believe this (or its choice-function equivalent) is the structure of the specific reading of the indefinite. It involves assigning particular individuals as values of the individual variables x,y, whereas in the non-specific interpretation further above, there is no such operation—it only follows as a consequence. As the two are truth-functionally equivalent, it is easy to see why the specific/nonspecific distinction is intuitively such a difficult one.

Now let us consider a case where the argument noun phrases must appear at the S level, as would be the case with "The cat chased the mouse". This means that at the V′ level, the thematic roles would be "unfilled".

(15)
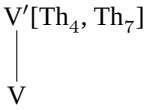

This would, quite obviously, correspond to a two-place relation in the propositional representation: V (,) (i.e., where the argument slots are not filled). But what relation is this? What we want to say is that this relation is one where, no matter how the arguments are filled in, all propositions so formed will be true only in (some subset) of those worlds in which there is an appearance of an event of the type designated by the verb. The functional representation is then defined as follows:

(16) $\forall x \forall y$ [ˆV(x,y) is interpreted as a subset of $\lambda w\ \exists e \in \llbracket V' \rrbracket$ & f(e) $\leq$ w]

This underspecifies the exact identity of the relation, but ensures that "the cat chase the mouse" will be true only in case there are chasing events.

But we also need to make sure that this relation requires both the cat and the mouse to be participants in the event-type as well, and, further, that they stand in a "chasing" relation, with the cat as the Agent and the mouse as Theme (let us suppose). But since they only appear outside the V′, and not within it, we have no means of making sure this occurs. In the case of the weak non-specific NP's there was something for the thematic roles to operate on, but in this instance there are no NP's in the V′ to be linked to the appropriate thematic roles.

Let us return then to the issue of morphology that we began the paper with. One way of looking at things is that it is the lack of morphology which has meaning—that of non-specificity. However, this turns out to be a side-effect of the fact that the unmarked NP's may occur within the V′. The thought here

is instead that it is not the lack of morphology that has any real function, but rather it is the presence of morphology that has it. And that function, I am hypothesizing, is to enable NP's that occur outside the domain of the thematic roles to be linked to them. Thus, the presence of morphology in those instances where there is a contrast enables the construction of a propositional representation where "higher" arguments nevertheless find their appropriate associations in the lower, V′ representations.

Let us illustrate with an example from English using definite articles, taking "the cat chase(d) the mouse". Since the definite article, at least in this instance, requires context in order to be interpretable, it must appear outside the V′ domain. There are various proposals for where, precisely, these "higher" arguments might appear in the syntax, and I not presume to choose among them. But schematically, the representation would be as follows:

(17)

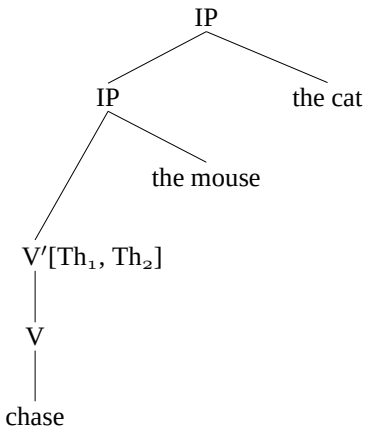

One function of the definite article is to simultaneously identify the appropriate thematic role (I am assuming this is accomplished via Dowty-type principles, but at this time I have no clear reason to prefer this over any other means of doing so), and it transmits the property-type defined by the N′ under the NP (or DP) to that role for the thematic role to operate on. So, in the syntactic representation, we coindex the roles with the definite articles:

(18)

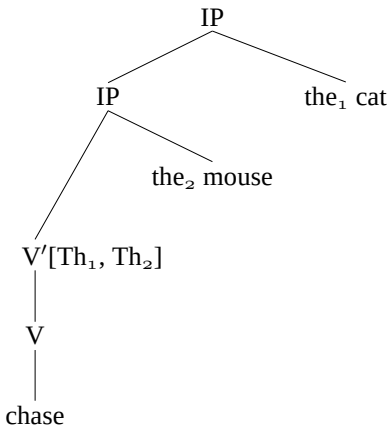

Having made this syntactic association we are now able to calculate the exact relation that the cat and the mouse stand in. Recall that nouns stand for properties, which have as extensions under V′ property-instantiations. Individuals have property-instantiations by virtue of being of certain types. In this instance, the interpretation of the N′ determines that the cat is a cat, and the mouse is a mouse, from the contents of the nouns themselves. Some of the property-instantiations are going to be mapped onto the particular individual cat that is talked about here, and also onto the mouse that is being talked about here. That is, each will find some extension in $\mathcal{N}$. Suppose these just happen to be i_1 and i_2, respectively. The event-type in this instance will be a very specific event-type defined by $[\![\text{chase}]\!]\text{Th}_1(i_1)\text{Th}_2(i_2)$. That is, where a "cat" property-instantiation of this particular individual that corresponds to i_1 is the Agent (or whatever the appropriate role name might be) and a property-instantiation of the mouse, i_2, is going to be Theme.

This enables us to define the **chase**(x,y) relation at the IP or S level, then. It is going to be guaranteed to be some sub-type of a chasing, and the particular sub-type of the chasing is going to be determined by the values assigned to x and y. These values are individuals which have property-instantiations in $\mathcal{N}$ by virtue of belonging to given classes. Thematic roles operate on property instantiations to derive, along with the verb, event-types. So, **chase**(a,b) is going to define the set of worlds where a has a property-instantiation, and b does too, and where they stand in the appropriate thematic role relations. This is exactly those worlds where **chase**(a,b) is true; the set of world so defined will be isomorphic to the event-type defined.

Assuming that proper names are syntactically in D (as Longobardi, 1992, has argued), there is no common noun expression, as is the case with pronouns and demonstratives. Yet there are no members of $\mathcal{N}$ that fail to fall under a

common noun meaning. One can deal with this by providing an interpretation of a very general sort for the empty N in these representations (e.g., "entity", or something of that character), but at this point it is only one among other possibilities which can be entertained with equal or perhaps greater convincingness.

1.4 Summary and conclusion

This paper began by considering a series of instances where there appears to be a contrastive lack of morphology associated primarily with noun phrases, but also with verbs, which collectively were associated with the phenomenon of indefiniteness, and in some very clear instances with non-specific indefiniteness. The question was posed as to whether we had an instance where a lack of morphology was associated with a function or a meaning, rather than its presence. It was certainly plausible to think of it that way. Three "stages" of semantic interpretation were introduced (when "normally" only one stage is presented), each of which with its own structures and elements. At the lowest level is lexical meaning, which is then projected homomorphically onto a higher domain (here, we focused on the V′ level). This domain of interpretation was projected homomorphically onto a still higher domain, the propositional representations that are most commonly used in semantics.

Within the V and V′ domains, due to their structure and elements, only weak indefinite noun phrases could receive any formal definition. It is these that may combine directly with the thematic roles associated with a verb. However, all other types of arguments fall outside this domain, and cannot be semantically defined except within the propositional representation. This necessitates having to associate the definites, quantified NP's, names, as well as specific readings of indefinites, with thematic roles that appear at lower positions in the syntactic representation. This, I am proposing, is the function of the extra morphology. It is not the absence that has a function. Rather, it is the presence.

Obviously, a lot of matters here have remained undiscussed, and many matters of important (and crucial) detail have been simply passed over. For example, what about languages which do not have any contrastive omission of morphology? Or, what about morphologically sparse languages, such as Mandarin? And, what about those instances where a given construction might be interpreted as either definite or indefinite, while a change or an addition of morphology yields only a definite? And, finally, what are the exact semantic facts for a much wider array of languages than those yet carefully studied? In a few instances they have been carefully researched, but in many other instances the semantics of the forms has not been examined with detailed care.

Many such questions remain open. My intent here, though, has been to provide a framework in which to pose questions and to further investigation. At this point the only secure prediction I can claim, though, is that one should not ever find a language where it is the morphologically more complex (or marked) forms that are non-specific indefinites, and the less complex forms are reserved for the specifics and definites. If this is confirmed by further research, I would feel that this paper has made at least one correct claim and has given some kind of answer as to why this might be so.

References

Aissen, J. (2003). Differential object marking: Iconicity vs. economy. *Natural Language and Linguistic Theory* 21, 435–83.

Asudeh, A. and L. Mikkelsen (2000). Incorporation in Danish: Implications for interfaces. In R. Cann, C. Grover, and P. Miller (eds.), ??A collection of Papers on HPSG. Stanford University.

Axelrod, M. (1990) Incorporation in Koyukon Athapaskan. *International Journal of American Linguistics* 56(2), 179–195.

Bach, E. (1986). The algebra of events. *Linguistics and Philosophy* 9, 5–16.

Baker, M. (1988). *Incorporation*. University of Chicago.

Basilico, D. (1998). Object position and predication forms. *Natural Language and Linguistic Theory* 16, 541–95.

Bealer, G. (1982). *Quality and Concept*. Oxford.

Bittner, M. (1994). *Case, scope, and binding*. Kluwer.

Borer, H. (in press). Exo-skeletal vs. endo-skeletal explanations: Syntactic projections and the lexicon. In J. Moore and M. Polinsky (eds.). *The Nature of Explanation in Linguistic Theory*. CSLI.

Borthen, K. (1998). *Bare singulars in Norwegian*. Cand. philol. thesis. NTNU, Department of Linguistics.

Borthen, K. (2003). *Norwegian Bare Singulars*. Doctoral dissertation, NTNU, Trondheim.

Carlson, G.N. (1984) Thematic Roles and their Role in Semantic Interpretation. *Linguistics* 22, 259–279.

Carlson, G. (2003a). Weak indefinites. In Martine Coene and Yves D'Hulst, eds. *From NP to DP: on the Syntax and Pragma-Semantics of Noun Phrases*. Benjamins Publishing. 195–210.

Carlson, G. (2003b). Interpretive asymmetries in major phrases. In A-M. Di Sciullo (ed.). *Asymmetry in Grammar, Vol. 1*. Benjamins. 301–14.

Cheng, Lisa, and Rint Sybesma (1999). Bare and not-so-bare nouns and the structure of NP. *Linguistic Inquiry* 30: 509–42.

Chierchia, Gennaro (1998). Reference to kinds across languages. *Natural Language and Linguistic Theory* 6, 339–405.

Chierchia, G. and S. McConnell-Ginet (2001). *Meaning and Grammar*. MIT press.

Chung, S. and W. Ladusaw (in press) *Restriction and Saturation*. MIT Press.

Cook, E-D. & A. Wilhelm (1998) Noun incorporation: New evidence from Athapaskan. *Studies in Language* 22(1), 49–82.

Dayal, V. (1999) Bare NP's, reference to kinds, and incorporation. *Proceedings of Salt IX*.

Diesing, M. (1992). *Indefinites*. Cambridge, Mass.: MIT Press.

Dowty, D. (1991). Thematic proto-roles and argument selection. *Language* 67, 547–619.

Dunn, M. (1999). *A Grammar of Chukchi*. Ph. D. Dissertation.

Enç, M. (1991). The semantics of specificity. *Linguistic Inquiry* 22:1, 1–26.

Farkas, D. and H. de Swart (to appear). *The semantics of incorporation: from syntax to discourse transparency*.

van Geenhoven, V. (1998). *Semantic incorporation and indefinite descriptions: semantic and syntactic aspects of noun incorporation in West Greenlandic*. Stanford: CSLI.

van Geenhoven, V. (2002). Noun incorporation. *GLOT International* 8, 261–71.

Gerdts, D. (1998). Incorporation. In A. Spencer and A. Zwicky (eds.) *The Handbook of Morphology*. Blackwell. 84–100.

Kallulli, D. (1999). *The comparative syntax of Albanian: On the contribution of syntactic types to propositional interpretation*. University of Durham dissertation.

Kamp, H. and U. Reyle (1993). *From discourse to logic*. Kluwer, Doredrecht.

Krifka, M. et al. (1995). Genericity: An introduction. In G. Carlson and F.J. Pelletier, (eds). *The Generic Book*. University of Chicago Press.

Kuroda, S.-Y (1972). The categorical judgement and the thetic judgement. *Foundations of Language* 9, 153–85.

Ladusaw, W. (1994). Thetic and categorical, stage and individual, weak and strong. In M. Harvey and L. Santelman (eds.), *Proceedings from Semantics and Linguistic Theory* IV. Cornell University, Ithaca. 220–29.

Longobardi, G. (1994). 1994. Reference and proper names: a theory of N-movement in syntax and logical form. *Linguistic Inquiry* 25, 609–669.

Longobardi, Giuseppe (2000). "Postverbal" subjects and the mapping hypothesis. *Linguistic Inquiry* 31: 691–702.

Marantz, A. (1997). No escape from syntax. In A. Dimitriadis, I. Siegeal et al (eds.), *University of Pennsylvania Working Papers in Linguistics* 4, 201–25.

Massam, D (2001). Pseudo noun incorporation in Niuean. *Natural Language and Linguistic Theory* 19, 153–197.

McNally, L. (1995). Bare plurals in Spanish are interpreted as properties. In Glyn Morrill and Richard Oehrle (eds.), *Formal Grammar*, Polytechnic University of Catalonia, Barcelona. 197–212.

Mithun, M. (1984). The evolution of noun incorporation. *Language* 62, 847–94.

Mohanan, T. (1995) Wordhood and Lexicality: Noun incorporation in Hindi. *Natural Language and Linguistic Theory* 13, 75–134.

Mosel, Ulrike and Even Hovdhaugen. (1992). *Samoan Reference Grammar*. Scandinavian University Press: Oslo, Norway.

Mulder, Jean Gail. (1994). *Ergativity in Coast Tsimshian (Smíalgyax)*. University of California Press: Berkeley.

Munn, Alan, and Christina Schmidt.(1999). Bare nouns and the morphosyntax of number. *Proceedings of the Linguistic Symposium on Romance Languages* 1999.

Myhill, J. (1988). Agent incorporation in Indonesian. *Journal of Linguistics* 24(1), 111–136.

Mohanan, T. (1995). Wordhood and lexicality: noun incorporation in Hindi. *Natural Language and Linguistic Theory* 13, 75–134.

Perrot, D. V. (1972). *Teach Yourself Swahili*. David McKay, NY.

Pesetsky, D. and E. Torrego (to appear). Tense, case, and the nature of syntactic categories. To appear in J. Guéron and J. Lecarme (eds.), *The Syntax of Time*. MIT Press.

de Reuse, W. (1994) Noun incorporation in Lakota (Siouan), *International Journal of American Linguistics*, 60(3), 199–260.

Schmidt, Christina and Alan Munn (1999). Against the nominal mapping parameter: Bare nouns in Brazilian Portuguese. *Proceedings of NELS 29*.

Spencer, Andrew. (1985). Incorporation in Chukchi. *Language* 71, 439–89.

Sullivan P. (1984) Noun incorporation in Yucatec Maya. *Anthropological Linguistics* 26, 138–160.

Suñer, M. (1988). The role of agreement in clitic-doubled constructions. *Natural Language and Linguistic Theory* 8, 391-434.

Woolford, E. (1996). Object agreement in Palauan: Specificity, Humanness, Economy, and Optimality. In J. Beckman, S. Urbanczyk, and L. Walsh

(eds), *University of Massachusetts Working Papers in Linguistics* 18. University of Massachusetts, Department of Linguistics.