

Ellipsis and Higher-Order Unification

Mary Dalrymple

Xerox PARC

Palo Alto, CA 94304, USA

and

Center for the Study of Language and Information

Stanford University

Stuart M. Shieber

Division of Applied Sciences

Harvard University

Cambridge, MA 02138, USA

Fernando C. N. Pereira

AT&T Bell Laboratories

Murray Hill, NJ 07974, USA

Abstract

We present a new method for characterizing the interpretive possibilities generated by elliptical constructions in natural language. Unlike previous analyses, which postulate ambiguity of interpretation or derivation in the full clause source of the ellipsis, our analysis requires no such hidden ambiguity. Further, the analysis follows relatively directly from an abstract statement of the ellipsis interpretation problem. It predicts correctly a wide range of interactions between ellipsis and other semantic phenomena such as quantifier scope and bound anaphora. Finally, although the analysis itself is stated nonprocedurally, it admits of a direct computational method for generating interpretations.

Contents

1	Introduction	2
2	Basics	2
2.1	An Abstract Statement of the Ellipsis Problem	2
2.2	Previous Analyses of Ellipsis	3
2.3	The New Analysis and an Example	5
2.4	Strict and Sloppy Readings	7
2.5	Constraints on Relation Formation	8
2.6	Constraints on Parallelism	9
2.7	Formal Semantic Background	10
2.7.1	Interpretation of Ellipsis Resolution Equations	11
2.7.2	Higher-order unification	13
2.7.3	Interpretation of quantification and long-distance dependencies	14
3	Interactions with Other Phenomena	17
3.1	Non-Constituent Abstractions	18
3.1.1	Sloppy readings with embedded antecedents	19
3.1.2	Non-subject abstraction	21
3.2	Multiple Property Abstraction	22
3.3	Cascaded Ellipsis	23
3.4	Interactions with Quantifier Scope	25
3.4.1	Quantification and antecedent-contained ellipsis	26
3.4.2	Quantification parallelism	29
3.4.3	Quantification and type raising	29
3.5	Other Phenomena	31
4	A Comparison of Approaches	32
4.1	Zero-Reading Analyses	33
4.2	Two-Reading Analyses	34
4.3	Three-Reading Analyses	35
4.4	Six-Reading Analyses	36
4.5	Five-Reading Analyses	37
4.6	Four-Reading Analyses	37
4.7	Summary	38
5	Problematic Cases	39
5.1	Non-Syntactic Parallelism	39
5.1.1	Semantic parallelism	39
5.1.2	Pragmatic parallelism	41
5.2	Further Constraints on Relation Formation	41
5.2.1	Obligatory sloppy readings	42
5.2.2	Antecedent-anaphor constraints	45
6	Conclusion	48

1 Introduction

In this paper, we present a new method for characterizing the interpretive possibilities generated by elliptical constructions in natural language. Unlike previous analyses, which postulate ambiguity of interpretation or derivation in the full clause source of the ellipsis, our analysis requires no such hidden ambiguity. For example, the ambiguity typically characterized as enabling “strict” versus “sloppy” readings of elliptical constructions does not arise from a corresponding ambiguity as to whether the pronoun in the antecedent clause is given a strict or sloppy interpretation; instead, the ambiguity follows from the process of interpreting the elided phrase on the basis of its unambiguous antecedent. Further, the analysis follows relatively directly from an abstract statement of the ellipsis interpretation problem and applies to the interpretation of a wide variety of elliptical constructions, including VP ellipsis, “do so” and “do it” anaphora, gapping, stripping, and related constructions involving recovery of implicit relations such as “only” modification and cleft constructions. It predicts correctly a wide range of interactions between ellipsis and other semantic phenomena such as quantifier scope and bound anaphora. Finally, although the analysis itself is stated nonprocedurally, it admits of a direct computational method for generating interpretations.

The analysis we present is intended to characterize the semantics of constructions involving ellipsis. Many interesting issues arise regarding the syntax of ellipsis, and we will touch on some of these issues; our main goal, though, is to characterize a method for ellipsis interpretation.

2 Basics

2.1 An Abstract Statement of the Ellipsis Problem

We can provide an abstract and reasonably theory-neutral characterization of ellipsis phenomena and their interpretation as follows. An elliptical construction involves two phrases (usually clauses) that are parallel in structure in some sense. The antecedent or *source* clause is complete, whereas the *target* clause is missing (or contains only vestiges of) material found overtly in the source. As a concrete example, which we will use as our primary source of data in the paper, consider the verb phrase (VP) ellipsis phenomenon, as in (1).

- (1) Dan likes golf, and George does too.

The sentence is interpreted as meaning that Dan and George both like golf. The source clause, ‘Dan likes golf’, parallels the target ‘George does too’, with the subjects ‘Dan’ and ‘George’ being parallel elements, and the VP of the target sentence being vestigially represented by the target phrase ‘does too’.¹

¹We emphasize that although the bulk of the examples in this paper involve VP ellipsis, the techniques can be applied to the semantic interpretation of many other elliptical constructions. This is in part because we do not restrict the notion of parallel elements to the subjects of the source and target clauses (see

Given this abstract view of ellipsis, the problem of ellipsis interpretation is just to recover a property of (or relation over) the parallel element (respectively, elements) in the target that the missing or vestigial material stands proxy for. Of course, this property is not arbitrary. We know that the application of the property or relation to the parallel elements in the source constitutes the interpretation of the source clause. In example (1) above, we know then that the property P being predicated of George in the second sentence is such that when it is predicated of Dan, it means that Dan likes golf. We might state this equationally as follows:

$$P(dan) = like(dan, golf)$$

A possible value for P in this equation is the property represented by the lambda term $\lambda x.like(x, golf)$. Predicating this property of George, we have $[\lambda x.like(x, golf)](george)$ which reduces to $like(george, golf)$.

In general, then, the abstract problem of ellipsis can be stated as the problem of recovering solutions to the equation

$$(2) P(s_1, s_2, \dots, s_n) = s$$

where s_1 through s_n are the interpretations of the parallel elements of the source, and s is the interpretation of the source itself. (The determination of the parallelism itself is a separate problem, about which more later.) Once P is determined, $P(t_1, t_2, \dots, t_n)$ serves as the interpretation of the target, where t_1 through t_n are the interpretations of the corresponding parallel elements of the target.

Not only is this an abstract characterization of the ellipsis problem, it is essentially the entire analysis proposed in this paper. It constitutes an analysis because the equational statement of the problem, together with some reasonable assumptions, determines rigorously the sets of interpretations for target clauses, which interpretations, we will see, correspond to the actual possible interpretations of the target.

2.2 Previous Analyses of Ellipsis

It is important that ellipsis analyses (including the equational one outlined above) allow for ambiguity in the target clause, that is, for a set of relations to be made available by the source clause. The availability of multiple relations is attested in various phenomena in which the target clause has multiple readings; it can be seen most clearly in the distinction between strict and sloppy readings. In the sentence

$$(3) \text{ Dan likes his wife, and George does too.}$$

Section 3.1). Indeed, the parallelism between the clauses need not even be purely syntactic (see Section 5.1). The generality of this ellipsis resolution method is one of its primary advantages.

(under the reading in which the pronoun ‘his’ refers to Dan) the property predicated of George might be the property of liking Dan’s wife or the property of liking one’s own wife. In lambda notation, these two properties are given by

$$\lambda x. \text{likes}(x, \text{wife-of}(\text{dan}))$$

and

$$\lambda x. \text{likes}(x, \text{wife-of}(x)) \quad ,$$

corresponding to the strict and sloppy readings of the sentence, respectively.² As we will see, the possibility of several available properties arises in other cases as well.

Most previous analyses of ellipsis have allowed for the possibility of multiple available properties by arranging for the source clause to be ambiguous as to what property it makes available. That is, in an individual instance, the source clause is interpreted in one of several possible ways, leading to the use of a particular property in the interpretation; the property used in the target clause is identical to the corresponding property in the source clause. Dahl (1972) was the first to draw a distinction between approaches that place the ambiguity in the source clause (which he called *strict identity* approaches) and those that place the ambiguity in the process of recovering a property or relation for the target (which he called *sloppy* or *non-strict identity* approaches). So as not to confuse the terminology here with that of strict and sloppy readings, we will call the former kind of analysis an *identity-of-relations* analysis, as opposed to a non-identity analysis.

Under an identity-of-relations analysis, ambiguity of interpretation in a target clause comes about because the source clause is ambiguous. However, only a single relation is available from the source clause in any given instance. The relation that the source clause makes available is, in most previous work, that associated with its VP, though we emphasize that this is not a necessary condition for an identity-of-relations analysis, nor do we believe it is a tenable stance.

The multiplicity of interpretations for the elided phrase in an identity-of-relations analysis may arise in various ways. Purely interpretive analyses (Gawron and Peters, 1990; Roberts, 1987) allow for multiple semantic interpretations arising from an unambiguous syntactic analysis of the source clause. Partially interpretive analyses involve either copying the syntactic tree from the source clause to the target but requiring identical semantic interpretations for the two VPs (Williams, 1977), or deleting the phrase structure of the

²For brevity, we represent the semantic interpretation of ‘Dan’s wife’ in this case as *wife-of(dan)*; the important thing to note about this representation is that the semantics of the pronoun ‘his’ is identical to that of its antecedent, ‘Dan’. Any other notation that has this property would do as well here. In particular, a treatment of possessives as introducing bound variables along the lines of the quantifier assumptions described in Section 2.7.3 is possible and perhaps preferable. Notations such as *wife-of(dan)* can be thought of as abbreviatory of such analyses. We only use such notations for cases in which the space of readings generated is unaffected.

second tree under the constraint of identical interpretation (Sag, 1976). Finally, syntactic analyses may also allow the semantic ambiguity to arise from underspecification in the copied constituent (Hellan, 1988).

The solutions therefore form a continuum, the ambiguity arising at more or less superficial levels. All of the analyses, however, share a reliance on semantic ambiguity in the source clause.

Our solution to the question of what properties are made available can be seen as lying at the far end of this continuum. We eschew not only syntactic ambiguity in the source clause, but semantic ambiguity arising from any source, as a generator of the multiple readings of elliptical constructions. Instead, multiple solutions come about as a natural result of directly stating the definition of the relation to be applied in the target clause.

2.3 The New Analysis and an Example

As described earlier, the problem of extracting a relation from the source clause can be stated equationally as

$$P(s_1, s_2, \dots, s_n) = s \quad .$$

In cases of VP ellipsis where the subjects of the source and target are parallel, the equation is simply

$$P(s_1) = s$$

where s_1 is the interpretation of the subject of the source clause. By solving this equation for the unknown, P , we generate the relation (or relations, if multiple solutions exist) that the resolution of the ellipsis requires.

Huet's *higher-order unification* algorithm (Huet, 1975) provides a means of completely enumerating representations of the solutions of such equations, under assumptions whose detailed discussion we defer to Section 2.7.³

As an example of the use of equations to state a problem in ellipsis resolution, consider (1) above, repeated here:

- (4) Dan likes golf, and George does too.

Recall that 'Dan' and 'George' are the parallel elements in this example, and the semantic interpretation of 'Dan likes golf' is

³The use of higher-order unification to resolve elliptical constructions has been independently noted by other researchers. Pulman and Milward implemented a prototype system that handled simple cases of VP ellipsis and gapping along these lines (Pulman, 1988). Pareschi and Steedman's "left conjunct decomposition" operation, which is used for the parsing of gapping constructions, bears a certain resemblance to higher-order unification as well (Steedman, 1990).

$$(5) \textit{like}(\underline{\textit{dan}}, \textit{golf})$$

We have underlined what we will call *primary occurrences* of the parallel element's interpretation, for reasons to be clarified later. In this case, the single occurrence of *dan* is primary. Any other occurrences will be referred to as *secondary*.

To form an interpretation for the second conjunct, we require a property P that, when applied to the interpretation of the subject of the first conjunct, will yield the interpretation of the first conjunct as a whole. This property will serve to generate the interpretation of the target clause. It will be applied to the interpretation of the parallel element, that of the subject 'George', in the second clause.

We can state this requirement directly with the following equation, an instance of the more general equation (2):

$$(6) P(\textit{dan}) = \textit{like}(\underline{\textit{dan}}, \textit{golf})$$

The latter term is the interpretation for the source sentence; the equation requires P to be a property that, when predicated of the subject interpretation *dan*, yields the first term.

A solution for an equation can be represented by a substitution of values for the free variables in the equation that makes both sides of the equation identical. For example, the following two alternative substitutions solve equation (6):

$$(7) \begin{aligned} P &\mapsto \lambda x. \textit{like}(\underline{\textit{dan}}, \textit{golf}) \\ P &\mapsto \lambda x. \textit{like}(x, \textit{golf}) \end{aligned}$$

The first substitution will be disregarded because it leaves a primary occurrence in the result. This constraint requiring *abstraction of primary occurrences* comes about because the parallel element in the target clause must play the primary role in the meaning of the target. We will have more to say about this in Section 2.5, especially as regards the distinction between primary and secondary occurrences. Given this constraint, the only remaining value for P is the property $\lambda x. \textit{like}(x, \textit{golf})$. We can now use this function as the interpretation of the elided VP in (1); with $P(\textit{george})$ as the interpretation of the target clause, this gives the following semantics for the sentence as a whole:⁴

$$(8) \textit{like}(\textit{dan}, \textit{golf}) \wedge \textit{like}(\textit{george}, \textit{golf})$$

In summary, our analysis of the abstract problem of ellipsis resolution, that is, generating appropriate properties to be used in interpreting the target clause, is to state the problem equationally based on the parallel structure in the two clauses and to solve the equation using higher-order unification (under the constraint requiring abstraction of primary occurrences). The properties that are generated as solutions to the equation are predicated of the parallel elements in the target clause to generate the target clause interpretation.

⁴At this point, we can ignore the distinction between primary and secondary occurrences.

2.4 Strict and Sloppy Readings

As seen in the preceding example, the equation stating an ellipsis interpretation problem may have several alternative solutions, which the higher-order unification algorithm will generate. Specifically, when there are multiple occurrences of some subterm in a term, multiple alternative substitutions for the relation formed by abstracting out that subterm will exist.⁵ Consider sentence (3), repeated here:

(9) Dan likes his wife, and George does too.

Let us assume the following interpretation for the first conjunct:

(10) $likes(\underline{dan}, wife-of(dan))$

The first occurrence of *dan*, which arises directly from the parallel element, the subject ‘Dan’, is primary; the occurrence arising from the pronoun, which is not a parallel element, is secondary. Solution of the equation $P(dan) = likes(\underline{dan}, wife-of(dan))$ by higher-order unification yields four ways of forming a property by abstracting out the semantics of the subject. The possible values for P are

(11) $\lambda x. likes(\underline{dan}, wife-of(dan))$
 $\lambda x. likes(\underline{dan}, wife-of(x))$
 $\lambda x. likes(x, wife-of(dan))$
 $\lambda x. likes(x, wife-of(x))$

Again, the first two solutions fail the constraint on abstraction of primary occurrences. Either of the other two remaining properties yields a possible interpretation of the target clause. The first gives rise to what has been called the strict reading of the second conjunct, while the second gives rise to the sloppy reading:

(12) $\lambda x. likes(x, wife-of(dan))(george) = likes(george, wife-of(dan))$
 (George likes Dan’s wife.)
 $\lambda x. likes(x, wife-of(x))(george) = likes(george, wife-of(george))$
 (George likes George’s wife.)

⁵Since multiple occurrences of the same proper name do not necessarily co-denote, we will use different constants as the representations of the denotata of such occurrences. For instance, the interpretation of the sentence

(a) Dan likes Dan’s wife.

should be $like(dan_1, wife-of(dan_2))$ where dan_1 and dan_2 are separate constants that only contingently co-denote. As a consequence, an ellipsis like

(b) Dan likes Dan’s wife, and Bill does too.

has only a strict reading, which accords with conventional wisdom.

2.5 Constraints on Relation Formation

As illustrated in the foregoing examples, we use the distinction between primary and secondary occurrences of a term to constrain the acceptable solutions of an ellipsis interpretation equation. We define a primary occurrence as an occurrence of a subexpression in the semantic form directly associated with one of the parallel elements in the source clause; we assume that the notion of “directly associated” is sufficiently well-defined for the purposes at hand. We then require that the solution process preserve the primacy of occurrences, with the consequence that solutions must abstract over all primary occurrences in the source. In other words: *Solutions must not include primary occurrences.* In example (9), this constraint removes from consideration the first two putative solutions in (11).

If the constraint were not in force, the following readings would be produced for ‘... and George does too’:

- (13) * ... and Dan likes Dan’s wife.
 * ... and Dan likes George’s wife.

These are just the readings where the parallelism between the clauses has been disregarded. Thus, the constraint is a reflex of the inherent parallelism in elliptical constructions.

The existence of this constraint means, not surprisingly, that it is necessary to retain a connection between the syntactic and the semantic representation of the source sentence. By maintaining this connection, we can ensure that the solutions produced by higher-order unification satisfy the constraint that parallelism must be maintained by abstracting out of parallel positions.⁶

⁶Solutions involving vacuous abstraction, such as

$$\lambda x. \text{likes}(\underline{\text{dan}}, \text{wife-of}(\text{dan})) \quad ,$$

are ruled out where necessary as special cases of this more general constraint. A direct prohibition against vacuous abstraction might be too strong, since verb phrase ellipsis is possible even in cases where the subject of the source clause is pleonastic and makes no semantic contribution (examples due to Ivan Sag):

- (14) a. John said it would rain, and it did.
 b. John said there would be trouble, and there was.

Suppose the interpretation of the former example (ignoring tense and aspect as usual) were

$$\text{said}(\text{john}, \text{rain}) \quad .$$

Then the second-order matching problem induced by the ellipsis would be

$$P(\Delta) = \text{rain} \quad .$$

(We use the symbol Δ for specifying the interpretation of pleonastic elements, following Gazdar et al. (1985, page 221).) The requirement of abstraction of primary occurrences would still allow a binding for P involving vacuous abstraction, namely $\lambda x. \text{rain}$. The elliptical clause would therefore be interpreted as *rain*.

2.6 Constraints on Parallelism

One of the distinguishing features of our analysis is that the ellipsis resolution problem is separated into two subtasks: a prior determination of the parallel structure of source and target, and consequent formation of the implicit relation to be used in the target. We have been addressing the latter subtask primarily, and will continue to do so, but we digress to mention some perhaps obvious facts about the parallelism determination that might get lost in the sequel.

The task of determining the parallel structure of two clauses is far more subtle, and less syntactic, than a cursory examination exposes. (We discuss this issue further in Section 5.1.) For this reason, the division of the ellipsis problem into two parts—separating parallelism determination from relation formation—allows a simpler description of relation formation and a more appropriate characterization of the problem of determination of parallelism. Nonetheless, this paper does not provide a theory of parallelism; previous attempts have been far too restrictive, limiting themselves to purely syntactic criteria. The wide range of possibilities for parallelism described in Section 5.1 indicate that the process is not a purely linguistic one. As an extreme example, we describe in Section 5.1 cases of parallelism with no linguistic source whatsoever.

Our emphasis in this paper on issues in relation formation and our liberal view of parallelism determination should not, of course, be taken to imply that no constraints apply to the task of determining parallelism. For example, parallelism must respect stativeness of verbs (15a) and pleonasticity of noun phrases (15b).

- (15) a. * Dan likes golf and George is too.
 b. * It is raining and George is too.

Depth of embedding imposes constraints as well.

- (16) * The mayor of Washington left, and New York did too.

(We ignore the nonsensical reading in which the city of New York was the agent of the leaving action.) Similarly, the sentence

- (17) It is obvious that Dan is happy, and George is too.

(pointed out to us by Mats Rooth) can only be interpreted as meaning that George is happy, not that it is obvious that George is happy. If the obviousness is included, the parallelism would have to hold between ‘George’ and ‘it’, not ‘Dan’.

Such constraints hold not only in VP ellipsis, but also in gapping, stripping, comparative deletion and other elliptical constructions. Thus, not all constraints on readings of elliptical sentences follow from the relation formation issues that are the primary topic of this paper.

Furthermore, there are syntactic constraints that apply differentially to different ellipsis constructions. Even among constructions eliding VPs, such as the “do”, “do so”, and “do too” variants, syntactic distinctions can be found. For instance, as noted by Haik (1985, page 177), these variants differ in their grammaticality in antecedent-contained-ellipsis contexts:

- (18) a. John greeted every person that Bill did.
 b. * John greeted every person that Bill did so.
 c. ? John greeted every person that Bill did too.

These fine syntactic distinctions are not addressed by the present analysis, which attempts to make clear only the space of semantic interpretations.

Of course, not all elements in the target clause must be analyzed as parallel to some element in the source. For instance, adverbial phrases can be viewed as modifying the target directly. This possibility is exemplified by the following sentence:⁷

- (19) Jim couldn't open the door, but Polly did with her blowtorch.

No empty adverbial modifier need be posited in the source clause; the instrumental modifies the target sentence directly. (Such elements can be made parallel in other cases, however; see Section 3.1.2 for examples.)

2.7 Formal Semantic Background

We outline here the formal machinery underlying the semantic analyses used in this paper.

Meanings of phrases are to be represented by terms of a typed higher-order system with lambda abstraction. Since we are not concerned here with intensional phenomena, we will just need the basic types e (entities) and t (truth values). Two type constructors will be used: \rightarrow to form the type $T \rightarrow T'$ of functions mapping arguments of type T to results of type T' , and \times to form the type $T \times T'$ of pairs $\langle t, t' \rangle$ such that t has type T and t' has type T' .

We will use various elementary concepts from the lambda calculus, specifically the notions of *free* and *bound* occurrences of variables, and of *substitution* of a lambda term for a variable. We will notate the substitution of term N for all free occurrences of x in M by $M[x \mapsto N]$. We require, as is typical, that substitution rename bound variables in M appropriately to avoid capture of free variables in N . The reader is urged to refer to Hindley and Seldin (1986) for precise definitions of these notions.

We have proposed codifying the ellipsis interpretation problem using expressions equating terms that represent phrase meanings. What counts as a solution to such an equation depends crucially on what notion of equality between terms we are considering, or, in

⁷We are indebted to an anonymous reviewer for this example.

other words, on when we consider that two terms denote the same semantic object. One salient notion of equality is that of $\alpha\beta\eta$ *interconvertibility* (Hindley and Seldin, 1986), which captures formally the intuitive notion that two terms can represent the same “recipe” for calculating a function. Specifically, two terms are considered equal if one can be converted to the other by repeated application of the following rules and their inverses:

α conversion: convert $\lambda x. M$ to $\lambda y. (M[x \mapsto y])$, that is, the names of bound variables are immaterial to their meaning. The two terms are said to be *alphabetic variants*.

β conversion: convert $[\lambda x. M](N)$ to $M[x \mapsto N]$. This represents formally the operation of applying a function to an argument.

η conversion: convert $\lambda x. M(x)$ to M when x does not occur free in M .

2.7.1 Interpretation of Ellipsis Resolution Equations⁸

We have claimed that the meaning of

(20) Dan likes golf, and George does too.

is

(21) $like(dan, golf) \wedge P(george)$

where the equation

(22) $P(dan) = like(dan, golf)$

must be satisfied. It is not obvious that the equation in 22 is semantically interpretable, as opposed to being a recipe for invoking a formal procedure, higher-order unification, with no underlying meaning in and of itself. Clearly, the combined meaning of 21 and 22 is not equivalent to

$$\exists P. like(dan, golf) \wedge P(george) \wedge P(dan) = like(dan, golf) \quad .$$

This would merely require that P be such that it is true of Dan if Dan likes golf. Since the first conjunct states that Dan does in fact like golf, P need only be a true property of Dan. The entire formula then requires only that George possess some property, any property, that Dan possesses, which would give an incorrect interpretation for the target sentence. The equation is not to be interpreted, then, as codenotation in a model.

⁸We are indebted to Mark Johnson and an anonymous reviewer for crystallizing these issues in our minds and for organizing the structure of possible responses. The particular statements made here are our own, however, and should not be interpreted as indicative of their opinions.

Instead, we want ellipsis to be more content-independent, in that the property should be such that the equation holds whether or not Dan happens to like golf. It should be independent of the particulars of a given model, that is, it should hold in any model in a suitable class of models. But we have to be careful about the choice of model class. Even necessary truths should not codenote over the class of models in which the ellipsis equation is interpreted. Otherwise, the sentence

(23) Every square has four sides, and every rhombus does too.

would be subject to an interpretation where what is predicated of every rhombus is some property of squares that is true of them in whatever models assign squares four sides. But these, *ex hypothesi*, include all the models, so any property true of all squares (such as having four equal angles) would do. The sentence might mean, then, that every rhombus has four equal angles. Of course, it does not. Similarly, logical tautologies must not be valid; sentence (20) does not, for example, mean that George likes golf and either it is raining or it is not raining, even though this is logically equivalent to George liking golf.

The class of models, then, in which the ellipsis equation is held as valid is very weak. In fact, for higher-order unification to be an appropriate procedure for determining valid instances, the valid equations must be exactly those whose two sides are $\alpha\beta\eta$ -interconvertible. Higher-order unification finds the most general substitutions of terms for the free variables in an equation that make the equation valid.

Friedman (1975) demonstrates that such equalities are exactly those that hold in all extensional models for the typed lambda calculus (without interpreted constants), and also exactly those that hold in any model consisting of all the higher-order functions over some infinite base set, with application interpreted as function application. Although the logic that we presuppose is augmented with a full first-order quantificational logic (and presumably, for intensional phenomena, would be an intensional logic), recall that the tautologies of this logic are not required to hold for the purposes of interpreting the ellipsis equations; the symbols of the logic (\forall , \wedge , etc.) can be viewed as uninterpreted function symbols of the appropriate type. The only structure that the model manifests is, then, the structure arising from the categorial or type structure of the language, together with the reasonable requirement of extensionality.

Thus, the semantic invariants in ellipsis resolution are those that follow from the type structure—the function-argument relationships—of natural language, and not from any contingent or even necessary truths. This accords with intuition, in that the felicity of ellipsis does not depend on the meaning of the words in the source sentence (though the elided property does), but does depend on their type structure. An ellipsis equation is not merely a recipe for a syntactic process. It has a meaning, but the meaning must be taken in a different, and much more profligate, model than that of the interpretation of the sentence itself. The ellipsis equation reflects semantic facts about the sentence, but just at the gross level of function-argument structure.

There is one remaining problem, however, for this view of ellipsis equations—the issue of primary occurrences. The primary occurrence notation serves to couple the ellipsis equations to the choice of parallel elements, and provides a way of forcing abstraction over the

meanings of the parallel elements. Intuitively, the distinction between primary and secondary occurrences is clear: a primary occurrence corresponds to a distinguished semantic role in the situations described by the source and target clauses. At present, however, we have no way of making precise the intuitions that led to the primary occurrence notation; that is, we know of no semantic correlate to the equational system with primary occurrence notation under the related constraint on abstraction. It remains for future work to reconstruct the semantical foundations of this variant of higher-order unification. Although we have some ideas as to how such a reconstruction might proceed, it is premature to discuss them here.

2.7.2 Higher-order unification

The unification problem for terms in a logical system is the problem of finding substitutions for the free variables of two terms t and t' that make the terms equal. Such a substitution is called a *unifier* of t and t' . A unifier σ is *more general* than another unifier σ' if there is a nontrivial substitution τ such that $(t\sigma)\tau = t\sigma'$, where $t\sigma$ is the result of applying substitution σ to term t . Informally, a more general unifier will leave more free variables, or make fewer variable identifications, than a less general one. Unifiers that are most general represent the solutions of a unification problem in their simplest form, since any less general unifier (solution) can be obtained from the output of a most general one by additional substitutions of terms for free variables. As is well known, the unification problem for first-order terms, and the related unification problem for certain kinds of feature graphs, admit of unique most general unifiers (up to variable renaming). However, this is not the case for higher-order unification, in which variables can range over functions of arbitrary order rather than just over [first-order] individuals. This multiplicity of unifiers corresponds to the possibility of multiple alternative interpretations for elided material.

Huet's higher-order unification algorithm enumerates the unifiers of higher-order terms in a typed λ -calculus of the kind we are using. Because higher-order unification is in general undecidable, given two terms, the algorithm will either stop without finding any unifiers (the terms are not unifiable), generate successive [most general] unifiers (possibly without end), or run forever without producing any unifiers. This computational property of higher-order unification has not been problematic on the cases we have examined that are engendered by ellipsis resolution, however, and there are several reasons why this might be so.

First, many of the equations arising in ellipsis resolution fall under the subcase called *second-order matching*. In the second-order subcase of unification, variables range only over individuals and first-order functions. The simpler matching problem occurs when the substitution need be applied only to one of the terms, that is, $t_1\sigma = t_2$. Huet and Lang (1978) use second-order matching as a way of applying program transformations in a manner reminiscent of the method used for ellipsis interpretation in this paper. Second-order matching is, fortunately, decidable, and Huet and Lang provide an algorithm, which is an adaptation of Huet's more general algorithm for the subcase of interest.

Furthermore, many of the equations we will be interested in solving are of the schematic

form

$$P(s_1, s_2, \dots, s_n) = s_0$$

where the s_i are all ground, that is, contain no free variables. The special case of second-order matching engendered by instances of this schema is computationally even more tractable. There are only a finite number of unifiers and these can be simply constructed as follows: Construct a term s from s_0 by replacing zero or more instances of the s_i by x_i . For each such s , construct a possible binding for P given by $\lambda x_1. \dots \lambda x_n. s$. Clearly, there are at most 2^c such unifiers where c is the number of occurrences of the s_i in s_0 , and these can be enumerated efficiently (in time linear in the output length).

Finally, although certain phenomena require use of the more general higher-order unification (as the examples in Section 3.4.3), the bulk of the cases considered in this paper rely on the ground subcase of second-order matching, and are therefore less computationally problematic. Of course, even the higher-order cases we consider may turn out to fall into a computationally reasonable subclass; further inquiry in this area would be useful.

2.7.3 Interpretation of quantification and long-distance dependencies

Before characterizing the interaction between our analysis of ellipsis and various other semantic phenomena, we must first lay out an approach to semantic interpretation—quantifier scoping in particular—in which to couch the discussion. For the most part, the particulars of the method for characterizing quantifier scoping are relatively unimportant; the analysis could be stated in terms of Cooper storage, say, or even quantifier raising.⁹ We will use here a variation of a method for interpreting quantifier scoping and long-distance dependencies developed by Pereira (1990). For those readers unfamiliar with this method, we provide some examples later in this section.

In general, the interpretation of a phrase will have the form $\Gamma \vdash m$ where Γ is a set of *assumptions* analogous to a quantifier store in the Cooper storage method (Cooper, 1983) and m is a matrix term in which free variables introduced by the assumptions in Γ may occur.

The assumptions used for quantifier scoping are triples of the form $\langle q \ x \ p \rangle$ where q is a determiner meaning, x is a free variable, and p is a term of type t in which x is free.¹⁰ The assumption $\langle q \ x \ p \rangle$ is said to *introduce* variable x . A quantified noun phrase is interpreted as a variable introduced by an assumption whose first component is the meaning of the

⁹The incorporation of an ellipsis analysis such as ours into a transformational framework is quite conceivable, merely requiring the ability to form abstractions over the syntactic objects representing semantic construals of sentences, that is, LF trees. The intrinsic portion of the analysis is its use of an equational framework for declaratively characterizing ellipsis resolution, not its use of particular logics for the representation of meanings. Nonetheless, the use of typed lambda calculus allows us to directly state our analysis with a minimum of extraneous machinery.

¹⁰Formally similar store elements have been used in quantifier scoping systems such as those of Schubert and Pelletier (1982) and Hobbs and Shieber (1987).

noun phrase’s determiner and whose third component represents the meaning of the noun phrase’s nominal.

Instead of the usual generalized quantifier type $(e \rightarrow t) \rightarrow (e \rightarrow t) \rightarrow t$ for determiner meanings, we will use the *pair quantifier* type $(e \rightarrow t \times t) \rightarrow t$. In other words, the meaning of a determiner takes a function from entities to pairs of truth values and yields a truth value. For example, the meaning of *every* will be the function that assigns ‘true’ just to those functions that take each entity to a pair whose second component is true whenever its first component is.

Our somewhat unusual type for determiner meanings, which is needed in our analysis of antecedent-contained ellipsis (Section 3.4.1) can be understood as a lambda calculus implementation of some aspects of Discourse Representation Theory (DRT) (Kamp, 1981; Heim, 1982).¹¹ Specifically, the interpretation of

(24) John greeted every person.

will be for us

(25) $every(\lambda x. \langle person(x), greet(j, x) \rangle)$

which we will abbreviate as

$every(x, person(x), greet(j, x))$.

This interpretation can be directly related to the discourse representation structure (DRS) for the same sentence:

(26)

x
$person(x)$

 \Rightarrow

$greet(j, x)$

The discourse marker x corresponds to the variable to be abstracted, the left inner DRS to the first element of the pair in (25) (the quantifier restriction), the right inner DRS to the second element of the pair (the quantifier scope) and the arrow to the determiner meaning *every*. The referential connection established by the discourse marker x in the DRS is simulated in our analysis by the simultaneous abstraction of the variable x over both the restriction and the scope of the quantifier.

It is straightforward to show that there is a one-to-one correspondence between Barwise-Cooper generalized quantifiers and pair quantifiers. Indeed, such a correspondence is established by two functionals, \mathcal{G} , mapping pair quantifiers to generalized quantifiers, and \mathcal{P} , mapping generalized quantifiers to pair quantifiers, satisfying $\mathcal{G}(\mathcal{P}(Q)) = Q$ and $\mathcal{P}(\mathcal{G}(P)) = P$.

¹¹Dynamic Montague grammar (Groenendijk and Stockhof, 1987; Chierchia, 1988) provides a more complex and possibly more general approach to incorporating some of the aspects of DRT into a compositional framework.

\mathcal{G} and \mathcal{P} are defined as:

$$\begin{aligned}\mathcal{G}(P) &= \lambda r. \lambda s. P(\lambda x. \langle r(x), s(x) \rangle) \\ \mathcal{P}(Q) &= \lambda p. Q(\lambda u. fst(p(u)), \lambda v. snd(p(v)))\end{aligned}$$

where $fst(\langle x, y \rangle) = x$ and $snd(\langle x, y \rangle) = y$.

For a derivation to be considered complete, all assumptions must be *discharged*. We will exemplify this process with sentence (24).

The quantified noun phrase ‘every person’ is given the interpretation

$$\langle every\ x\ person(x) \rangle \vdash x \quad .$$

that is, the meaning of the noun phrase is x under the assumption to the left of the \vdash .

The verb meaning $\lambda o. \lambda s. greet(s, o)$ applied to the NP meaning yields a VP meaning, still under the above assumption:

$$\langle every\ x\ person(x) \rangle \vdash \lambda s. greet(s, x) \quad .$$

Application of this VP meaning to the subject meaning $\vdash j$ results in this sentence meaning:

$$\langle every\ x\ person(x) \rangle \vdash greet(j, x) \quad .$$

Discharging the quantifier assumption involves applying the quantifier *every* to the result $\lambda x. \langle (person(x), greet(j, x)) \rangle$ of abstracting x over the pair consisting of the restriction and the scope of the quantifier. The resulting interpretation is

$$\vdash every(x, person(x), greet(j, x))$$

that is, a sentence meaning free of undischarged assumptions. When several quantifier assumptions are introduced, there is the option of discharging them in several different orders, leading to alternative quantifier scopings for the sentence.

In Pereira’s original system, the treatment of quantifier assumptions is semantically justified by showing how such derivations are convenient shorthand for derivations in the Curry system of semantic combination containing only the operations of functional application and abstraction. In the present variant, we could carry out a similar justification in a system that would include pairing in addition to functional application and composition (Lambek, 1980; van Benthem, 1989).

The treatment of the semantics of long-distance dependencies is handled similarly by introducing and discharging assumptions. Again, the form of these introduction and discharge rules could be justified on the basis of functional application and abstraction. As an example of a derivation involving a long-distance dependency, we consider the example

(27) John greeted every person that arrived.

The trace in the subject position of the relative clause can be thought of as introducing a *bind* assumption for a new variable

$$\langle \textit{bind } x \rangle \vdash x$$

which serves as the argument to $\lambda s. \textit{arrive}(s)$ (the interpretation of ‘arrived’):

$$\langle \textit{bind } x \rangle \vdash \textit{arrive}(x) \quad .$$

The relative clause being completed, we can now discharge the bind assumption. We do so by forming a higher-order predicate by abstracting the matrix, conjoined with a place-holder for the modified nominal, and abstracted by the bound variable.

$$\vdash \lambda N. \lambda x. N(x) \wedge \textit{arrive}(x)$$

This relative clause meaning serves as a function over the nominal meaning $\lambda y. \textit{person}(y)$.

$$\vdash \lambda x. \textit{person}(x) \wedge \textit{arrive}(x)$$

Finally, the rule for combining a quantifier *every* with a predicate forms a quantifier assumption over a new variable.

$$\langle \textit{every } z \textit{ person}(z) \wedge \textit{arrive}(z) \rangle \vdash z$$

From here, the derivation continues as before, yielding the sentence meaning (before the quantifier is discharged)

$$\langle \textit{every } z \textit{ person}(z) \wedge \textit{arrive}(z) \rangle \vdash \textit{greet}(j, z)$$

and the scoped meaning is

$$\vdash \textit{every}(z, \textit{person}(z) \wedge \textit{arrive}(z), \textit{greet}(j, z)) \quad .$$

This completes the background information on the formal semantics we will presume in the remainder of the paper.

3 Interactions with Other Phenomena

The approach to ellipsis resolution that is advocated here displays differences from previous approaches in its handling of various phenomena. We will discuss how our analysis differs from identity-of-relations analyses in general, and certain particular instances thereof, by briefly examining the predictions of the analyses with respect to the following phenomena:

Non-constituent abstractions: There are many cases in which the relation constructed from the source clause does not correspond in any straightforward fashion to the interpretation of some syntactic constituent: for example, when it must take more than one argument. For instance, the tense and aspect as well as the subject are abstracted in the sentence

Dan is running for president, and George did last term.

Examples demonstrate other nonstandard abstractions as well. Such cases are especially problematic for identity-of-relations analyses in which the relation is necessarily associated with some constituent such as the VP in the source clause.

Multiple property extraction: In some cases, a single sentence serves as the antecedent for two subsequent instances of ellipsis involving different parallel elements:

John finished reading the poem before Bill did, and the short story too.

This sentence has a reading on which John finished reading both the poem and the short story before Bill finished reading the poem. This is problematic for identity-of-relations analyses in which only a single property is available in any given instance for the interpretation of subsequent elided phrases.

Cascaded ellipsis: Analyses differ as to what readings are predicted for sentences containing multiple elliptical clauses in which the interpretation of one elided constituent depends partially or entirely on the interpretation of another elided constituent. An example is:

John realizes that he is a fool, but Bill does not, even though his wife does.

Interaction with quantifier scoping: As is well known, the ambiguities following from varying quantifier scope possibilities interact with ellipsis resolution possibilities. For instance, in the sentence

John greeted every person when Bill did.

two readings are possible, depending on whether the universal quantifier has wide scope over both the main and subordinate clause, or quantifies separately in each clause. But in

John greeted every person that Bill did.

only a wide scope reading is available.

We discuss each of these phenomena below, and demonstrate that our approach constructs appropriate solutions. In the succeeding section, we discuss in detail an example sentence which illustrates differences among a number of analyses of ellipsis that have been proposed in the past. Finally, we turn to problematic cases for this and other analyses.

3.1 Non-Constituent Abstractions

There are many instances in which the interpretation of elided phrases does not correspond to the interpretation of a syntactic constituent in the source clause. The most obvious cases include the elliptical constructions of gapping or stripping. But VP ellipsis provides examples as well. For instance, there are cases in which a deeply embedded constituent induces a sloppy reading; in other cases, relations are formed with multiple parallel elements in the source and target clause; in still other cases, as discussed in Section 5.1 below, the parallelism between the elements in the source and target clause is not syntactic, but semantically or pragmatically induced. These cases are problematic for identity-of-relations approaches in general, since such approaches would have to make available a very large number of different semantic analyses for each source clause, some of them otherwise unmotivated, to allow for all of the possible interpretations that might need to be provided for subsequent ellipsis. They are particularly problematic for identity-of-relations analyses in which the interpretation provided by the source clause corresponds to the translation of a syntactic constituent in the source.

3.1.1 Sloppy readings with embedded antecedents

The primary argument given by Reinhart (1983) for the distinction between bound variable and referential pronouns is the requirement that bound variable pronouns must be c-commanded by their antecedents. She uses this requirement to predict that the following example has only a strict reading:

(28) People from LA adore it and so do people from NY. [Reinhart's (17a), page 150]

Reinhart proposes a requirement that a pronoun must be c-commanded by its antecedent if the antecedent is a quantifier; further, she claims that a pronoun giving rise to a sloppy interpretation must be c-commanded by its antecedent. For Reinhart, then, the availability of a sloppy reading correlates with the possibility of a bound-variable interpretation of a pronoun, and she requires a c-command relation for this interpretation to be possible. This restriction simplifies the task of an identity-of-relations analysis because it reduces the number of cases in which a sloppy reading is available. An analysis postulating ambiguity of pronoun interpretation for only this restricted set of cases seems methodologically more plausible.

However, Reinhart herself (1983, page 178) notes certain counterexamples to this correlation, cases where a sloppy reading is available even when c-command does not hold:

- (29) a. Felix_{*i*}'s mother thinks he_{*i*}'s a genius and so does Siegfried's mother. [Reinhart's (8a)]
- b. We'll discuss Rosa_{*i*}'s problems with her_{*i*} parents and Sonya's problems too. [Reinhart's (8b)]

Wescoat (1989) notes a number of more extreme cases of sloppy readings involving non-c-commanding, embedded constituent antecedents, such as:

- (30) a. The policeman who arrested John failed to read him his rights, and so did the one who arrested Bill.
- b. The person who introduced Mary to John would not give her his phone number, nor would the person who introduced Sue to Bill.

Wescoat claims, and we agree, that sloppy readings are possible with these sentences; that is, that the following readings are available—perhaps even preferred—for them:

- (31) a. The policeman who arrested John failed to read John John's rights, and the one who arrested Bill failed to read Bill Bill's rights.
- b. The person who introduced Mary to John would not give Mary John's phone number, and the person who introduced Sue to Bill would not give Sue Bill's phone number.

Hirschberg and Ward (1991) have obtained experimental evidence consistent with these examples that the c-command criterion posited by Reinhart and counterexemplified by Wescoat is not a general requirement for sloppy readings. For example, their data show that as many as one-third of a group of test subjects preferred the sloppy reading of

- (32) People from Los Angeles think it's a scary place to live, and so do people from New York. [Hirschberg and Ward's (31)]

which parallels sentence (28) closely.

On our analysis, barring any stipulated prohibition, there is no obstacle to forming relations abstracted over arbitrarily deeply embedded positions. In skeletal form, the analysis of such an example would proceed as follows. Suppose the source clause of (30b) is interpreted as¹²

$$\text{refuse}(\text{pwi}(m, j), \text{give}(m, \text{phone}(j)))$$

where $\text{pwi}(x, y)$ is the person who introduced x to y , and $\text{phone}(x)$ is x 's phone number. The parallel elements in the construction are, respectively, 'the person who introduced Mary

¹²See Footnote 2 for a discussion of the status of the function symbols pwi and phone .

to John’ and ‘the person who introduced Sue to Bill’; ‘Mary’ and ‘Sue’; ‘John’ and ‘Bill’. Thus, the appropriate equation to solve is

$$P(pwi(m, j), m, j) = refuse(\underline{pwi(m, j)}, give(m, phone(j))) \quad .$$

The sloppy reading is engendered by the following unifying substitution for P :

$$(33) \lambda x. \lambda y. \lambda z. refuse(x, give(y, phone(z))) \quad ,$$

which, when applied to the interpretations of the parallel elements in the target, yields the target interpretation

$$P(pwi(s, b), s, b) = refuse(pwi(s, b), give(s, phone(b))) \quad .$$

Note that the recovered relation (33) is not a relation corresponding to a conventional interpretation for the VP—or any other constituent—in the source clause. On an identity-of-relations analysis, such relations would be available only by virtue of their use in the derivation of the source clause interpretation. This would necessitate the postulation of wild ambiguity in the source clause, one derivation for each possible case of subsequent ellipsis.

3.1.2 Non-subject abstraction

There exist many cases of multiple parallel elements in the source and target clause; it is very common for ellipsis to involve relations formed by abstraction of elements other than the interpretation of the subject noun phrase.

For example, the tense and aspect of the target clause might differ from that in the source clause:

$$(34) \text{Dan is running for president, and George did last term.}$$

The mood can also differ:

$$(35) \text{a. "I want to leave." "Well, do."}$$

$$\text{b. "Eat your dinner." "I did."}$$

The two clauses may differ in polarity:

$$(36) \text{Dan didn't leave, but George did.}$$

These examples show that relations of varying arity must be available as interpretations for elided phrases. The consequence of this for theories where the relation available for interpretation of subsequent ellipsis must be available in the source sentence is, again, that every sentence which can be the antecedent for subsequent ellipsis must be many ways ambiguous; an interpretation must be available for each relation that might be needed to interpret ellipsis in subsequent discourse.

On the equational analysis, however, such ambiguity is not required. A single interpretation for the source clause can give rise to any required interpretation for the target, since there is no inherent restriction as to the number or nature of the parallel elements involved in the ellipsis.

Take, for instance, the sentence in (36). Here, the parallel elements are ‘Dan’ and ‘George’ (the subjects), and the positive and negative polarities, represented semantically as operators *pos* and *neg*.¹³ Thus, the equation

$$P(dan, neg) = \underline{neg}(\underline{left}(dan))$$

can be solved yielding

$$P \mapsto \lambda x. \lambda S. S(\underline{left}(x))$$

and applied to the target parallel elements:

$$P(george, pos) = pos(\underline{left}(george)) \quad .$$

Of course, gapping and stripping provide abundant examples of non-subject abstraction involving other arguments or modifiers; Reinhart (1983, page 152) provides this example, which involves non-subject parallel elements and has both a strict and a sloppy reading:

- (37) You can keep Rosa in her room for the whole afternoon, but not Zelda. [Reinhart’s (18c)]

Jackendoff (1972, page 275) also discusses examples involving both subject and non-subject parallelism:

- (38) Maxwell killed the judge with a silver hammer, and I’d like to do the same thing to that cop, with a cudgel. [Jackendoff’s (6.196)]
- (39) Fred hung Tessie up in a tree and poured paint on her, but I bet he wouldn’t do it to Sue with glue. [Jackendoff’s (6.197)]

¹³Other analyses of polarity, for instance as an implicit argument of the predicate, will yield similar results.

3.2 Multiple Property Abstraction

A difficulty in any identity-of-relations analysis, which makes available only one interpretation for subsequent clauses exhibiting ellipsis, is seen when a single sentence is the antecedent for the ellipsis of two different noun phrases. Consider the following:

(40) John finished reading the poem before Bill did, and the short story too.

This sentence has a reading on which John finished reading both the poem and the short story before Bill finished reading the poem. On this reading, the source for both elliptical clauses is the same clause, ‘John finished reading the poem.’ To produce a relation which can be the interpretation for the elided VP whose subject is Bill, the interpretation for the sentence ‘John finished reading the poem’ must be derived as:

$$[\lambda x. \textit{finish-reading}(x, \textit{the poem})](\textit{john})$$

so as to make available the property $\lambda x. \textit{finish-reading}(x, \textit{the poem})$. Similarly, an interpretation must also be produced for the second conjunct ‘and the short story too’. On the desired reading, the interpretation for ‘John finished reading the poem’ must be derived as:

$$[\lambda y. \textit{finish-reading}(\textit{john}, y)](\textit{the poem}) \quad .$$

Under an identity-of-relations analysis, the source clause is deemed ambiguous between the two derivations. They do not simultaneously exist in a given analysis; only one or the other may be chosen. Thus, the reading noted above would not be generable. On the other hand, an analysis such as ours allows for the formation of two different relations from the semantic representation of the first sentence; the representation of the first sentence does not constrain the possibilities for construction of such relations. The next section provides an example of a similar problem and its derivation in our framework.

3.3 Cascaded Ellipsis

We use the term “cascaded ellipsis” to refer to cases of multiple ellipsis in which one of the elided constituents depends on another elided constituent for its interpretation. Analyses dependent on an identity-of-relations approach generally make available fewer readings in cascaded ellipsis cases than the analysis presented here; we believe that the greater number of readings available with our analysis is in fact warranted.

Dahl (1972) provides the following example (Dahl’s (12), an English paraphrase of Scheibe’s (58a) (1973)):

(41) John realizes that he is a fool, but Bill does not, even though his wife does.

Dahl claims that this sentence has, among other readings, the following one:

- (42) John realizes that John is a fool but
 Bill does not realize that Bill is a fool, even though
 Bill's wife realizes that Bill is a fool

Sag (1976, page 135 ff.) discusses this example; his claim is that this reading is not available for this sentence. We disagree, and find the reading acceptable.

On our analysis, this reading is readily available. Assume that the interpretation for 'John realizes that he is a fool' on the reading under discussion is:

$$realize(john, fool(john))$$

This sentence serves as the antecedent for the elided phrase in the second conjunct, 'Bill does not'. 'Bill' and 'John' are parallel elements; for the reading under discussion, second-order matching solves the equation

$$P(john) = realize(john, fool(john))$$

producing, among others, the following property (corresponding to the sloppy option):

$$P \mapsto \lambda x. realize(x, fool(x))$$

which is applied to 'Bill'. The interpretation for the second conjunct as a whole is, then, the following:¹⁴

$$realize(bill, fool(bill))$$

We assume that the second clause may serve as an antecedent for the elided portion of the third conjunct. The parallel elements are 'Bill' and 'his wife'; the ellipsis equation is¹⁵

$$Q(bill) = realize(bill, fool(bill)) \quad .$$

On the reading under discussion, the strict option is chosen; the property Q applied to the interpretation of 'his wife' is:

$$\lambda x. realize(x, fool(bill))$$

The resulting interpretation for the third conjunct is:

$$realize(wife-of(bill), fool(bill))$$

¹⁴For simplicity of exposition, we ignore the fact that the polarities of the two sentences differ. See Section 3.1.2 for a discussion of this issue.

¹⁵As the semantic roles for 'Bill' are parasitic on those for 'John', we let the primary occurrences of 'Bill' be determined by those of 'John' in the source sentence.

Although we have described the derivation of the meaning for this example in terms of temporal ordering (we resolve the first ellipsis, using its result to resolve the second), it is important to note that the analysis is truly order-free. In essence, we merely set up two equations in two unknowns and solve them using unification. The result, as is typical with declarative, equational methods, does not depend on solving the equations in a particular order.

Under an identity-of-relations analysis, such as Sag’s, the existence of this reading is problematic, as he notes. The problem is that there are conflicting requirements on the form of the semantic representation of the second clause.

Sag obtains strict and sloppy readings under ellipsis by optionally applying a rule that replaces the interpretation of a pronoun (which has an invariant referent and induces a strict reading) by a lambda-bound variable (inducing a sloppy reading). The representation Sag provides for the first two conjuncts is:¹⁶

$$(43) \textit{John}, \lambda x. [x \textit{ realize } x \textit{ is a fool}] \textit{ but} \\ \neg \textit{Bill}, \lambda y. [y \textit{ realize } y \textit{ is a fool}]$$

Crucially, the interpretation for the pronoun ‘he’ which is reconstructed in the second conjunct is represented by a bound variable.

In contrast, for the reading under discussion, the representation for the second and third conjuncts is:

$$(44) \neg \textit{Bill}_i, \lambda x. [x \textit{ realize } he_i \textit{ is a fool}] \textit{ even though } his_i \textit{ wife}, \lambda y. [y \textit{ realize } he_i \textit{ is a fool}]$$

The strict reading is only available when the option of replacing the pronoun interpretation with a lambda-bound variable is not taken. These conflicting requirements make it impossible for Sag’s analysis—or any identity-of-relations analysis, where the difference between a strict and a sloppy reading corresponds to a difference in the form of the semantic representation—to obtain the reading for this sentence that we (and Dahl) assume exists.

3.4 Interactions with Quantifier Scope

As described in Section 2.7, quantifier scope is generated through a mechanism of discharging of assumptions introduced in the course of a derivation. The interaction between quantifier scoping and ellipsis, then, will simply involve the relative derivational order of discharging such quantifier assumptions and resolving elliptical constructs. That is, when we set up the appropriate instance of the schematic equation $P(s_1, s_2, \dots, s_n) = s$, s is the meaning of the source clause, *possibly under one or more undischarged assumptions*.

¹⁶Sag uses the notation $a, \lambda x. [P]$ to represent the application of a VP meaning $\lambda x. [P]$ to the subject meaning a . The scope of the negation operator is the entire predication.

3.4.1 Quantification and antecedent-contained ellipsis

As an example, we consider the two sentences given in (45).

- (45) a. John greeted every person when Bill did.
 b. John greeted every person that Bill did.

As noted by Sag (1976), the quantifier scope possibilities differ for these two sentences: whereas (45a) admits of two readings, (45b) allows only one. Both Sag and Williams (1977) provide analyses for these semantic intuitions.

We will take the source of the ellipsis in (45a) to be the clause ‘John greeted every person’, and the target to be ‘Bill did’. The derivation of interpretations for (45a) proceeds as follows. The source clause is interpreted under a quantifier assumption generated by the subject NP. (See Section 2.7.3 for the derivation.)

$$(46) \langle \text{every } x \text{ person}(x) \rangle \vdash \text{greet}(\text{john}, x)$$

We might discharge the assumption at this point, but we choose not to in this first scenario. Consequently, the interpretation of the full sentence is

$$\langle \text{every } x \text{ person}(x) \rangle \vdash \text{greet}(\text{john}, x) \text{ when } P(\text{bill})$$

where P is constrained equationally by virtue of the interpretation of the source clause:

$$P(\text{john}) = \text{greet}(\underline{\text{john}}, x) \quad .$$

This equation has a single (most general) solution

$$P \mapsto \lambda z. \text{greet}(z, x) \quad .$$

It is this value for P that we will apply to *bill*. Thus, the interpretation of the full sentence, with ellipsis resolved is

$$\langle \text{every } x \text{ person}(x) \rangle \vdash \text{greet}(\text{john}, x) \text{ when } \text{greet}(\text{bill}, x)$$

The assumption may now be discharged, yielding the full interpretation

$$\text{every}(x, \text{person}(x), \text{greet}(\text{john}, x) \text{ when } \text{greet}(\text{bill}, x)) \quad .$$

This interpretation corresponds to a necessarily distributive reading, the ‘individual’ reading, in which each person is simultaneously greeted by John and Bill.

Alternatively, the assumption can be discharged before the ellipsis is resolved. Under this scenario, the interpretation of the source clause is

$$\vdash \text{every}(x, \text{person}(x), \text{greet}(\text{john}, x)) \quad .$$

The full sentence, then, is interpreted as

$$\vdash \text{every}(x, \text{person}(x), \text{greet}(\text{john}, x)) \text{ when } P(\text{bill})$$

where, again, the interpretation of the source clause is used to constrain the property P :

$$P(\text{john}) = \text{every}(x, \text{person}(x), \text{greet}(\underline{\text{john}}, x)) \quad .$$

The single value for P is

$$\lambda z. \text{every}(x, \text{person}(x), \text{greet}(z, x))$$

leading to the final interpretation

$$\begin{array}{l} \text{every}(x, \text{person}(x), \text{greet}(\text{john}, x)) \text{ when} \\ \text{every}(x, \text{person}(x), \text{greet}(\text{bill}, x)) \end{array} \quad .$$

This interpretation yields a ‘group’ reading paraphrasable as ‘John greeted every person when Bill greeted every person.’ The two derivations, then, correspond to just the interpretations noted by Sag.¹⁷

Now, we turn to the superficially similar sentence (45b). We take the source clause to be the entire sentence, and the target to be, again, ‘Bill did’. The interpretation of the sentence before resolution of ellipsis and discharge of assumptions is

$$\langle \text{every } x \text{ person}(x) \wedge P(\text{bill}) \rangle \vdash \text{greet}(\text{john}, x) \quad .$$

(Again, Section 2.7.3 records a derivation for a similar clause.)

As before, we will resolve the ellipsis by finding solutions for the equation

$$P(\text{john}) = \text{greet}(\underline{\text{john}}, x)$$

¹⁷Lappin (1984) proposes a combination of Cooper storage and copying for interpreting sentences such as (45a). His scheme differs substantively from ours in that he copies both matrix and store from the storage counterpart to give the interpretation of the elided material. However, such a scheme fails to preserve the scope parallelism noted in Section 3.4.2.

whose solution assigns P the relation $\lambda z. \text{greet}(z, x)$. With this substitution, the interpretation is

$$\langle \text{every } x \text{ person}(x) \wedge \text{greet}(\text{bill}, x) \rangle \vdash \text{greet}(\text{john}, x)$$

which, discharging the assumption, reduces to

$$\text{every}(x, \text{person}(x) \wedge \text{greet}(\text{bill}, x), \text{greet}(\text{john}, x)) \quad .$$

Alternatively, we might attempt to discharge the assumption before resolving the ellipsis. Recall the starting point for the previous derivation:

$$\langle \text{every } x \text{ person}(x) \wedge P(\text{bill}) \rangle \vdash \text{greet}(\text{john}, x) \quad .$$

Discharging the assumption yields

$$\vdash \text{every}(x, \text{person}(x) \wedge P(\text{bill}), \text{greet}(\text{john}, x))$$

Resolving the ellipsis, then, involves finding solutions to the equation

$$P(\text{john}) = \text{every}(x, \text{person}(x) \wedge P(\text{bill}), \text{greet}(\underline{\text{john}}, x)) \quad .$$

Notice that the variable P appears on both sides of this equation. For this reason, the derivation runs into problems, for there simply are no solutions to this equation; no unifier exists for this pair of typed terms. Thus, the derivation fails at this point; the sentence has only the ‘individual’ reading, in agreement with our judgments.

The computational reflex of the above lack of a solution is a violation of the so-called “occurs check” in the unification algorithm. This check prevents the construction of unifiers in which the substitution for a variable contains that variable. In other words, the occurs check blocks the creation of a circularity (a value for P containing P itself). It is interesting to note that this is formally analogous to the syntactic argumentation which Williams (1977) uses to eliminate such readings.

In sum, the ellipsis characterization described here allows for antecedent-contained ellipsis, and predicts correctly the interactions with quantifier scope.

3.4.2 Quantification parallelism

Another implication of the account of quantifier scope given above is that in cases where the source clause exhibits quantifier scoping ambiguities, if the quantifiers are separately quantified in the two clauses (i.e., as in the ‘group’ reading for (45a)) the scopes in the two clauses must be the same. For instance, in the sentence

(47) John gave every student a test, and Bill did too.

we predict (correctly, we take it) no reading of the form

$$\begin{aligned} & \text{every}(x, \text{student}(x), \text{exists}(y, \text{test}(y), \text{give}(j, x, y))) \\ & \wedge \text{exists}(y, \text{test}(y), \text{every}(x, \text{student}(x), \text{give}(b, x, y))) \end{aligned}$$

where the two quantifiers take different scopes in the two clauses. Consider the possible orderings of ellipsis resolution and discharging of quantifier assumptions. If ellipsis resolution occurs before some of the quantifiers have been discharged, these quantifiers will scope over *both* clauses. Thus, the only way for both quantifiers to scope separately is for ellipsis resolution to occur after both quantifier assumptions are discharged. But in that case, the ellipsis equation will include the quantifier order manifest in the source clause interpretation, and this will be carried over to the target interpretation.

3.4.3 Quantification and type raising

In general, the semantic types of parallel elements must be identical. In the case of a quantified NP parallel to a non-quantified NP, this implies that the type of the non-quantified NP must be raised to that of the quantified type. As an example, we consider the sentence

(48) Every student revised his paper, and then Bill did.

This sentence is ambiguous, depending on whether Bill revises his own paper or each student’s paper.

As usual, the ellipsis resolution can occur before or after the quantifier assumption is discharged. If the quantifier is discharged first, we have the meaning given in 49 for the first clause.

$$(49) \text{every}(x, \text{student}(x), \text{revise}(x, \text{paper-of}(x)))$$

To resolve the ellipsis, we need to set up an equation involving the interpretation of the parallel element in the source, “every student”. However, for this purpose we cannot directly use the stored noun phrase interpretation given in 50.

$$(50) \langle \text{every } x \text{ student}(x) \rangle \vdash x$$

This interpretation only makes sense as part of the derivation of the meaning of the whole source clause. Instead, we calculate the contribution of “every student” to the meaning of the source clause by examining the effect of applying it to an arbitrary property S . Combining S with 50, we obtain first the interpretation in 51.

$$(51) \langle \text{every } x \text{ student}(x) \rangle \vdash S(x)$$

The assumption can then be discharged to yield the sentence meaning in 52.

$$(52) \text{every}(x, \text{student}(x), S(x))$$

The contribution of “every student” is thus to map an arbitrary property S to the term in (52), that is, the contribution is the function given in 53.

$$(53) \lambda S. \text{every}(x, \text{student}(x), S(x))$$

It is easy to show that this is equivalent to the usual generalized quantifier meaning of “every student”.

To resolve the ellipsis, then, we set up the equation in 54 involving the meaning (53) of the parallel element in the source, “every student”:

$$(54) \quad P(\lambda S. \text{every}(x, \text{student}(x), S(x))) \\ = \text{every}(x, \text{student}(x), \text{revise}(x, \text{paper-of}(x)))$$

This equation has the solution

$$P \mapsto \lambda Q. Q(\lambda x. \text{revise}(x, \text{paper-of}(x))) \quad .$$

The type of the noun phrase meaning $\lambda S. \text{every}(x, \text{student}(x), S(x))$ is $(e \rightarrow t) \rightarrow t$. The meaning for “Bill” must be type-raised to $\lambda R. R(b)$ for type consistency. The value for P , when applied to the type-raised meaning for “Bill”, yields the target meaning

$$\text{revise}(b, \text{paper-of}(b))$$

according to which Bill revises his own paper.

On the other hand, if ellipsis resolution occurs first, the derivation yields

$$\langle \text{every } x \text{ student}(x) \rangle \vdash \text{revise}(x, \text{paper-of}(x))$$

just before resolution. The equation

$$P(x) = \text{revise}(x, \text{paper-of}(x))$$

admits of a strict interpretation for P :

$$P \mapsto \lambda y. \text{revise}(y, \text{paper-of}(x)) \quad .$$

The meaning for the conjoined sentence before ellipsis resolution

$$\langle \text{every } x \text{ student}(x) \rangle \vdash \\ \text{revise}(x, \text{paper-of}(x)) \text{ and then } P(b)$$

reduces, after ellipsis resolution, to

$$\langle \text{every } x \text{ student}(x) \rangle \vdash \\ \text{revise}(x, \text{paper-of}(x)) \text{ and then } \text{revise}(b, \text{paper-of}(x)) \quad ,$$

which, following discharging of the quantifier assumption, becomes

$$\text{every}(x, \text{student}(x), \text{revise}(x, \text{paper-of}(x)) \text{ and then} \\ \text{revise}(b, \text{paper-of}(x)))$$

On this reading, Bill revises each student's paper after the student revises it. A sloppy reading, on which Bill revises his own paper, is generable in this way as well; in this particular case, though, it is logically equivalent to the reading described above, in which type-lifting of 'Bill' is involved.

3.5 Other Phenomena

We defer discussion of several other important interactions of ellipsis and scoping phenomena to a companion paper in preparation. In that paper we intend to discuss, in addition to a more detailed explication of the mainstream quantifier cases:

Scope ambiguities with indefinites: Indefinites give rise to several readings under ellipsis, depending on their scope and the relative order of ellipsis resolution and discharge of the indefinite. This sentence, for example, has three readings:

John lost a book he owned, and so did Bill.

On the first reading, John and Bill lost the same book; on the second reading, John and Bill each lost one of John's books, possibly distinct; and on the third reading Bill lost one of his own books.

De dicto/de re ambiguities: Similar ambiguities are found in sentences with opaque verbs. This sentence has three readings:

Bill wants to read a good book and John does too.

Where the first conjunct has a de dicto reading, the second conjunct does also; where the first conjunct has a de re reading, however, there are two readings for the sentence as a whole, depending on whether or not Bill and John want to read the same book.

‘**Canadian flag**’ examples: Hirshbühler (1982) discusses examples such as the following:

A Canadian flag was hanging in front of each window, and an American one was too.

Examples such as these illustrate that the subject of the source clause need not take widest scope in VP ellipsis. On our analysis, these examples, like those in Section 3.4.3, involve matching of higher than second order.

4 A Comparison of Approaches

In order to more fully explicate the differences between our approach to ellipsis resolution and other approaches, we analyze a single example in detail from the perspective of several previous proposals. Rather than make the comparison in the respective notations of the original proposals, we normalize those notations by using lambda terms uniformly. When the analyses are viewed in this way, several apparently different analyses are seen to generate the same set of readings in the same manner, despite their having originally been stated in differing notations.

We will use the following example, discussed at length by Gawron and Peters (1990):

(55) John revised his paper before the teacher did, and Bill did too.

We follow Gawron and Peters in directing attention to the reading of the first conjunct where ‘John’ and ‘his’ corefer, and of the second elliptical clause where its source is the entire previous sentence. The following six readings exhaust the set of readings generated by any of the analyses (including our own) that we discuss. We present them with paraphrases for reference.

(56) a. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(teacher))) and
before(revise(bill, paper-of(bill)),
revise(teacher, paper-of(teacher)))*

Each person revised his own paper.

b. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(john))) and
before(revise(bill, paper-of(john)),
revise(teacher, paper-of(john)))*

Each person revised John’s paper.

- c. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(john))) and
before(revise(bill, paper-of(bill)),
revise(teacher, paper-of(bill)))*

John and then the teacher revised John's paper; Bill and then the teacher revised Bill's paper.

- d. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(teacher))) and
before(revise(bill, paper-of(john)),
revise(teacher, paper-of(teacher)))*

John and Bill both revised John's paper before the teacher revised the teacher's paper.

- e. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(john))) and
before(revise(bill, paper-of(bill)),
revise(teacher, paper-of(john)))*

John and Bill revised their own papers before the teacher revised John's paper.

- f. *before(revise(john, paper-of(john)),
revise(teacher, paper-of(john))) and
before(revise(bill, paper-of(john)),
revise(teacher, paper-of(bill)))*

John and then the teacher revised John's paper; Bill revised John's paper before the teacher revised Bill's paper.

As we will see, the equational analysis that we propose in this paper is the most profligate of the analyses, potentially generating all six of these readings, though restrictions might eliminate certain of these.

4.1 Zero-Reading Analyses

The strictest version of an identity-of-relations analysis requires that a lambda term used in the derivation of the meaning of the source clause be used in the derivation of the target clause meaning (either by copying or deletion under identity). Under such an analysis, the pertinent level of semantic representation of the source clause to use in the target clause derivation is that before beta reduction has occurred, as beta reduction eliminates the function-typed lambda terms. For the sentence 'John revised his paper', the unreduced meaning representation is one of:

$$\lambda x. \text{revise}(x, \text{paper-of}(\text{john}))(\text{john})$$

or

$$\lambda x. \text{revise}(x, \text{paper-of}(x))(john) \quad .$$

corresponding to the strict and sloppy readings, respectively. In forming the meaning of the first target clause ‘before the teacher did’, we use whichever term P is made available by the first clause, generating $P(\text{teacher})$. An issue remains as to how the two clause meanings are then combined to form a single sentence meaning. The most natural method, direct coordination, would yield (for the sloppy reading):¹⁸

$$\begin{aligned} &\text{before}(\lambda x. \text{revise}(x, \text{paper-of}(x))(john), \\ &\quad \lambda y. \text{revise}(y, \text{paper-of}(y))(teacher)) \end{aligned}$$

corresponding to a syntactic analysis under which the adverbial clause is attached at the S level. However, this is not itself of the form appropriate for being the source of a later ellipsis, that is, a function applied to the subject meaning. Thus, under this analysis, the second ellipsis, ‘and Bill did too’, would be uninterpretable. (Recall that we are ignoring the readings in which the source for the second ellipsis is merely ‘John revised his paper’. Such a reading would be possible, although it would be strict or sloppy dependent on the interpretation of the other two clauses.)

4.2 Two-Reading Analyses

The second ellipsis is not, of course, uninterpretable, so we attempt to design a meaning representation for its source that is of the appropriate form. A first method is to place the meaning of the ‘before’ clause within the meaning of its source VP:

$$\begin{aligned} &[\lambda x. \text{before}(\text{revise}(x, \text{paper-of}(x)), \\ &\quad \lambda y. \text{revise}(y, \text{paper-of}(y))(teacher))](john) \end{aligned}$$

(We will discuss a second method in Section 4.3.) The exact mechanism for constructing such a reading is not specified here. It may seem especially problematic, as the source and target terms are not alphabetic variants in this notation. Nonetheless, if this were the meaning representation of the source of the second ellipsis, it would allow for a sloppy reading of the target; this would correspond to the sloppy reading (56a). The strict variant

$$\begin{aligned} &[\lambda x. \text{before}(\text{revise}(x, \text{paper-of}(john)), \\ &\quad \lambda y. \text{revise}(y, \text{paper-of}(john))(teacher))](john) \end{aligned}$$

would yield reading (56b).

An alternative method of subordinating the ‘before’ clause meaning maintains the alphabetic variance property of the two clauses. If we assume, counterintuitively, that the

¹⁸Here and elsewhere, we uniformly rename bound variables apart for clarity.

before operator takes a property and a truth value to a property (that is, it is of type $((e \rightarrow t) \times t) \rightarrow (e \rightarrow t)$), we can form the meaning representation

$$[before(\lambda x. revise(x, paper-of(x)), \\ \lambda y. revise(y, paper-of(y))(teacher))](john)$$

and similarly for the strict case. Again, the mechanism is a bit mysterious (though less so) and readings (56a) and (56b) would be generated for the sentence as a whole.

The analyses of Sag (1976) and Williams (1977), although they do not consider cases such as these, might be reasonably viewed as generating these readings in much this way. Similarly, the analysis presented by Roberts (1987) and phrased in terms of DRT generates these two readings (as discussed by Gawron and Peters).

4.3 Three-Reading Analyses

The representation for the source of the second ellipsis might be extrapolated along different lines. In particular, the interpretation for each verb phrase, including the compound one, might be given by an overt abstraction. This would correspond to a syntactic analysis under which the adverbial clause is adjoined to the main clause verb phrase, with one verb phrase meaning appearing as a subconstituent of the other. For the sloppy reading (56a), we would have

$$\lambda u. [before(\lambda x. revise(x, paper-of(x))(u), \\ \lambda z. revise(z, paper-of(z))(teacher))](john) \quad ,$$

and for the strict (56b)

$$\lambda u. [before(\lambda x. revise(x, paper-of(john))(u), \\ \lambda z. revise(z, paper-of(john))(teacher))](john) \quad .$$

This representation has the benefit of being uniform, preserving alphabetic variance, and assigning a more attractive type to *before*. It also allows for a third reading when used as the source for the second elliptical clause. The second argument to *revise* might be—in addition to *j* (the strict option) or *x* (the locally sloppy option)—the reabstracted subject meaning *u*. This leads to a globally sloppy meaning (56c).

$$\lambda u. [before(\lambda x. revise(x, paper-of(u))(u), \\ \lambda z. revise(z, paper-of(u))(teacher))](john)$$

This analysis, which generates three readings for the example sentence, is essentially the analysis developed by Gawron and Peters within a situation-theoretic framework. (The final reading corresponds to the second argument of *revise* being “absorbed” by the lambda operator.) Our transliteration makes clear that, for this case at least, situation-semantics machinery is not necessary to yield the readings in question; an extrapolation of Sag’s or Williams’s analyses might achieve the same result. Of course, other aspects of the Gawron and Peters analysis depend intrinsically on the situation-theoretic foundation.

4.4 Six-Reading Analyses

The three readings provided by the Gawron and Peters analysis seem to exhaust the possibilities for an identity-of-relations approach. Our analysis produces six readings for the example sentence.

Let us examine in detail how one of the readings for this sentence, (56c) above, is obtained under the equational analysis of ellipsis. Assume that the semantics for ‘John revised his paper’ is:

$$\mathit{revise}(\mathit{john}, \mathit{paper-of}(\mathit{john}))$$

The first conjoined sentence then will have the meaning

$$\mathit{before}(\mathit{revise}(\mathit{john}, \mathit{paper-of}(\mathit{john})), \mathit{P}(\mathit{teacher}))$$

under the constraint

$$\mathit{P}(\mathit{john}) = \mathit{revise}(\underline{\mathit{john}}, \mathit{paper-of}(\mathit{john})) \quad .$$

The second elliptical clause takes its source to be the whole first conjunction. Thus, its interpretation will be

$$\mathit{before}(\mathit{revise}(\mathit{john}, \mathit{paper-of}(\mathit{john})), \mathit{P}(\mathit{teacher})) \wedge \mathit{Q}(\mathit{bill})$$

under the constraint

$$\mathit{Q}(\mathit{john}) = \mathit{before}(\mathit{revise}(\underline{\mathit{john}}, \mathit{paper-of}(\mathit{john})), \mathit{P}(\mathit{teacher})) \quad .$$

These two equations in two unknowns (P and Q) are solved, as usual, by higher-order unification; we will take the equation for P first. One solution is to take the strict reading for P ,

$$P = \lambda x. \mathit{revise}(x, \mathit{paper-of}(\mathit{john}))$$

leading to the following interpretation for the second equation:

$$\mathit{Q}(\mathit{john}) = \mathit{before}(\mathit{revise}(\underline{\mathit{john}}, \mathit{paper-of}(\mathit{john})), \\ \mathit{revise}(\mathit{teacher}, \mathit{paper-of}(\mathit{john})))$$

This equation, in turn, has a solution

$$Q = \lambda x. \mathit{before}(\mathit{revise}(x, \mathit{paper-of}(x)), \\ \mathit{revise}(\mathit{teacher}, \mathit{paper-of}(x)))$$

The semantics for this reading of the sentence as a whole is:

*before(revise(john, paper-of(john)),
revise(teacher, paper-of(john))) and
before(revise(bill, paper-of(bill)),
revise(teacher, paper-of(bill)))*

Our analysis allows for readings that are missing under the analyses discussed above because it is not an identity-of-relations analysis; interpretation of ellipsis does not involve copying the interpretation of a constituent in the source.

4.5 Five-Reading Analyses

In Section 5.2.2, we discuss a restriction on unifiers that uniformly eliminates certain readings of elliptical clauses. This restriction, when applied to the example at hand, eliminates reading (56f). Thus, our analysis, strictly speaking, generates only five of the six combinatorially possible readings of the sentence.

4.6 Four-Reading Analyses

An unpublished analysis attributed to Hans Kamp (personal communication to Mark Gawron and Stanley Peters, cited by Gawron and Peters (1990)) and couched in DRT assigns four readings to the sentence, and does so by eliminating the identity-of-relations assumption. In Kamp's analysis, as in our own, ambiguities between strict and sloppy readings do not arise from ambiguity in the source clause; the source has only a single interpretation. Essentially, Kamp makes a copy of the discourse representation structure of the source, and then imposes constraints identifying the participants in the source and target copies. These constraints must be applied in a symmetric manner. If a sloppy interpretation constraint applies to one copied discourse entity, it must apply to all; similarly for a strict interpretation constraint. Gawron and Peters mention a possible extension to Kamp's analysis that allows for the generation of all six of the readings listed above by relaxing the symmetry requirement. We refer the reader to the discussion by Gawron and Peters for a fuller description of Kamp's proposal.

Insofar as Kamp's analysis can be fleshed out, his analysis and ours make the same predictions as to the class of readings available in cases of cascaded ellipsis. Readings missing under other analyses are available for our analysis and his. The particular syntactic operations that Kamp (under Gawron and Peters's reconstruction) presupposes have no particular foundation other than efficacy. Our analysis can be seen as providing an argument for the operational view implicit in Kamp's analysis based on the underlying equational characterization of elliptical constructions. This equational foundation, as we have seen, articulates with other semantic phenomena in ways not appreciated in the previous research.

4.7 Summary

In summary, each analysis differs in the number of analyses that are predicted for the given sentence. Here is a scorecard.

<i>Method</i>	<i># readings</i>
Sag, Roberts	2
Gawron and Peters	3
Kamp	4 (or 6)
equational	5

The single reading that our method derives that remains underived by others is that in (56e). Clearly, this reading is difficult, if not impossible, to dig out (although plausibility considerations play a large role here). However, we have seen examples that demonstrate that the reason for its absence in the other analyses is faulty. For instance, its elimination on the basis of an identity-of-relations analysis, as Roberts or Gawron and Peters would have it, has repeatedly been seen to be too strong.

Rather, the pertinent distinction in differentiating the first four readings from the last two is that the resolution of the second elliptical construction in the last two readings must treat the parallel structures that ellipsis applies to in a non-parallel fashion. We conjecture that such non-parallel cases are highly dispreferred, if not disallowed entirely.

In order to test this hypothesis, we can try to construct an example which pragmatically favors this dispreferred reading to see whether it is obtainable. The following sentence is parallel in structure to sentence (55):

(57) Dewey announced his victory after the newspapers did, but so did Truman.

A reasonable, historically accurate reading for this sentence may be represented as:

(58) *after(announce(Dewey, Dewey's victory),
announce(newspaper, Dewey's victory)) but
after(announce(Truman, Truman's victory),
announce(newspaper, Dewey's victory))*

That is, Dewey claimed victory for himself after the newspapers announced Dewey's victory, but Truman also claimed victory after Dewey was announced the winner by the newspapers. This reading is parallel to reading (56e) described above.

Opinions differ as to the acceptability of this reading. One's opinion in this case can be seen as a litmus determining whether parallelism of the sort violated here is required, or merely preferred.

5 Problematic Cases

The following issues are problematic for most analyses of ellipsis interpretation. We present them, along with our conjectured solutions, to codify the range of phenomena that analyses of ellipsis might account for and to provide a preliminary guess as to their possible solutions in our framework.

5.1 Non-Syntactic Parallelism

Our analysis of ellipsis resolution presupposes identification of the source of the ellipsis and the parallel structuring of the source and target. This division of labor between identification of parallelism and resolution of ellipsis is purposeful, as the factors involved in the solution of the two problems are quite different. Although determining the parallelism may seem to be a purely syntactic operation, much like the matching that goes on at the semantic level, this similarity is illusory. Cases of semantic or pragmatic parallelism also exist. These cases are particularly problematic for theories of ellipsis in which the interpretation of an elided phrase is presumed to correspond to the interpretation of some syntactic constituent in the source clause, as is the case in most identity-of-relations analyses.

5.1.1 Semantic parallelism

Examples of ellipsis exist in which the parallelism is between the “logical subject” (i.e., passive agent) in the source clause and the surface subject in the target clause:

- (59) a. A lot of this material can be presented in a fairly informal and accessible fashion, and often I do. (Chomsky, 1982, page 41)
- b. It should be noted, as Dummett does, . . . (example due to Ivan Sag)

Similar examples involving “so-anaphora” are also found:

- (60) a. The formalisms are thus more aptly referred to as information- or constraint-based rather than unification-based, and we will do so here. (Shieber, 1989, page 2)
- b. It is possible that this result can be derived from some independent principle, but I know of no theory that does so. (Mohan, 1983, page 664)

Examples of this type are ubiquitous, but seem to be confined pragmatically to cases where the source clause states a general fact or rule, and the target clause provides a specific instance of this fact or rule.

Examples of passive/active parallelism are not confined to those in which the source and target appear in the main clause. In the following example (due to Wescoat (1989)), parallelism of the heads of the main clause subjects forces a parallelism of arguments of the modifying relative clauses; *John*, the object of the relative clause in the source sentence, is parallel to *Bill*, the subject of the relative clause in the target sentence:

- (61) The policeman who arrested John failed to read him his rights, and so, for that matter, did the one Bill got collared by.

Our analysis does not require that the property provided as the interpretation for the elided portion of the target clause in examples like those above correspond to the interpretation of any constituent in the source clause. It is not clear that there is any analysis available for examples of this sort within a theory in which the interpretation for elided phrases must be that of some constituent in the source clause, as is the case in most identity-of-relations analyses.

Other cases of semantic/thematic parallelism can also arise; sentence 62 is from instructions on a bottle of Agree shampoo:

- (62) Avoid getting shampoo in eyes—if it does, flush thoroughly with water.

Syntactically, the parallel elements are the object of the source clause and the subject of the target; thematically, these elements are the “theme” arguments of the intransitive/causative *get* verb pair.

Other combinations of logical-subject/surface-subject parallelism do not seem to arise. The following examples, where the parallel elements are intended to be the surface subject in the source clause and the logical subject in the target, are ungrammatical:

- (63) a. *Dummett notes, as it should be, . . .
 b. *We will refer to the formalisms as information- or constraint-based, as they more aptly are.

However, the following sentence (due to Peter Sells) has a similar structure, yet seems to be more acceptable:

- (64) John completed the assignment faster than it ever had been in the history of the school.

To the extent that this sentence is grammatical, it illustrates that either the source or target clause can contain a logical subject which is parallel to a surface subject in the other clause. Although examples such as these are often restricted in their distribution, they demonstrate that the parallelism between elements in the source and target clause need not be confined to surface syntactic parallelism.

5.1.2 Pragmatic parallelism

More extreme cases exist in which arbitrary information may need to be brought to bear to determine the appropriate parallelism between source and target. Webber (1978) cites some such cases:

- (65) a. Irv and Mary want to dance together but Mary can't since her husband is here.
b. Mary wants to go to Spain and Fred wants to go to Peru, but because of limited resources, only one of them can.

Recovery of the pertinent properties in these sentences requires nonlinguistic knowledge concerning social norms and economic processes. In example (65b), for instance, the implicit property that only one of Mary and Fred can have is *to do what he or she wants to do*. Other attested examples include:¹⁹

- (66) Fortunately, the first person to die in 1990 and the first couple to file for divorce in 1990 were allowed to do so anonymously. (Roeper, 1990)
- (67) Amid applause at the Congress of the Russian Federation (RSFSR), Mr. Yeltsin put forward a bill setting Russian law above the law of the Soviet Union – something Mr. Gorbachev, as Soviet president, declared unconstitutional when Estonia, Latvia and Lithuania did it last year. (Rettie, 1990)

Sentences of this sort illustrate that, to a greater or lesser extent, relations involved in the resolution of anaphoric processes such as ellipsis can be made available contextually. Identity-of-relations analyses allow for only the simplest cases of resolution of elided constituents, since the only mechanism that is available to provide an interpretation for the target is that of copying an interpretation from the source. Our approach goes beyond identity-of-relations analyses by allowing for the construction of new relations on the basis of old ones; the use of unification to construct relations is, as we have seen, more powerful and more flexible than copying.

Even for cases of what Hankamer and Sag (1976) call “surface anaphora”, such as verb phrase ellipsis, resolution is possible to relations that are pragmatically determined or influenced, as the examples in (65) show. To interpret these examples, not only the apparatus we introduce for constructing new relations but also pragmatic knowledge must be brought to bear.

¹⁹We are indebted to James McCawley and Bonnie Webber for bringing these examples to our attention.

5.2 Further Constraints on Relation Formation

Cases in which sloppy but not strict readings are available might seem to be problematic for an analysis like the one presented here. Below we will examine cases of control and reflexivization in which exclusively sloppy readings are available. Solutions will be proposed which do not involve constraining the process by which relations are formed as interpretations for elided constituents.

However, there are other cases in which readings are unexpectedly unavailable; these cases generally involve multiple occurrences of pronouns whose antecedent is a parallel element in the source clause. It seems that a constraint is necessary on the possibilities for forming relations in cases such as these.

5.2.1 Obligatory sloppy readings

Control In general, only sloppy readings are available for sentences involving control. The following sentence is not ambiguous:

(68) John tried to run, and Bill did too.

There is no reading according to which Bill tries to bring it about that John runs.

Chierchia (1983; 1984), noting facts of this type, proposes that the semantic type of a controlled verb phrase is a property rather than a proposition. That is, for a sentence like “John tried to run”, the correct semantic representation would be (a) and not (b):

- (69) a. $try(john, \lambda x. run(x))$
 b. $try(john, run(john))$

If Chierchia’s hypothesis is correct, the lack of a strict reading is predicted; in the representation in (a), there is only a single occurrence of “*john*”, and a strict reading is impossible to produce.

However, there are reasons to doubt the adequacy of an analysis like this one. First, anaphors whose antecedent is the subject of a controlled VP can give rise to both a sloppy and a strict reading under ellipsis. Consider this sentence:

(70) John tried to kill himself before Bill did.

Since the reflexive “himself” cannot be bound to a higher clause subject, its antecedent must be the subject of “kill”. On the property analysis, the interpretation for the first conjunct would then be:

- (71) $try(john, \lambda x. kill(x, x))$

However, we find this sentence to have two readings, corresponding to the following paraphrases:²⁰

- (72) a. John tried to kill himself before Bill tried to kill himself.
 b. John tried to kill himself before Bill tried to kill John.

The second of these readings would not be obtainable given (71) as the source clause meaning.

Zec (1987) demonstrates another problem for Chierchia’s property analysis. The data she cites bear on Chierchia’s claim that there is a correlation between the syntactic form and the semantic type of complements: that complements of type VP are properties, while complements of type S’ are propositions. Chierchia predicts that there would be no case in which an S’ complement gives rise to only a sloppy reading.

Zec discusses cases of obligatory control in Serbo-Croatian, showing that there are cases of obligatory control into verbal complements that are of the syntactic category S’ rather than VP. For example, the verb *pokušati* “try” takes an S’ complement, and yet only a sloppy reading is possible in the following sentence (Zec, 1987, page 143):

- (73) Petar je pokušao da postane predsednik a to je pokušala i Marija
 Petar Aux tried that become president and it Aux tried too Marija
 “Petar tried to become president and Marija tried it too.”

This sentence means that Marija tried to bring it about that Marija (not Petar) become president.

There seem to be two possibilities with regard to the Serbo-Croatian data. The first is to deny that there is a necessary correlation between the syntactic type of the complement and its semantic type. The claim would be in that case that, although the syntactic type of the complement of “try” is an S’ in Serbo-Croatian, semantically it is a property. In that case, the lack of a strict reading would be predicted for the Serbo-Croatian case, just as it is for the English case; the problem illustrated by example (70) would remain, however.

Another option, and the one that Zec takes, is to posit an obligatory coreference relation between the subject of “try” and the subject of its complement clause; this relation would presumably be induced by the control verb. If this option is taken, it would presumably force the abstraction of both arguments at the same time under second-order matching.

²⁰ Additional readings arise if the source clause for the ellipsis is taken to be not the matrix “John tried to kill himself”, but the VP complement “to kill himself”:

- (a) John tried to kill himself before Bill killed himself.
 (b) John tried to kill himself before Bill killed John.

Similar comments apply to these examples.

In our terms, the obligatory coreference relation and the obligatory sloppy readings Zec discusses would mean that controlled occurrences are primary in both Serbo-Croatian and English. Thus, in (70), the interpretation would be

$$try(\underline{john}, kill(\underline{john}, john)) \quad ,$$

which, when taken to be a source for ellipsis, would generate only two solutions to the equation

$$P(john) = try(\underline{john}, kill(\underline{john}, john)) \quad ,$$

manifesting a sloppy reading for the controlled subject occurrence and either a strict or a sloppy reading for the reflexive occurrence, as required.

Reflexivization Sells et al. (1987) present a number of cases of reflexivization in which only sloppy readings are available. The Dutch reflexive *zich* is such a case (Sells et al., 1987, page 182):

- (74) Zij verdedigde zich beter dan Peter
 She defended herself better than Peter
 “She defended herself better than Peter.”

Sells et al. characterize reflexive constructions involving only sloppy readings as “closed predicate” constructions. They discuss only examples in which the reflexive appears in object position, with its antecedent being the subject of the same clause.

We might assume, then, that for the examples they discuss, the presence of the reflexive correlates with the operation of a semantic relation-reducing rule, one which semantically “intransitivizes” the verb. In the Dutch case, then, the presence of *zich* signals a change in the meaning of the verb from the meaning in (a) to the one in (b):

- (75) a. $\lambda x. \lambda y. defend(x, y)$
 b. $\lambda x. defend-self(x)$

A solution of this type would work for all cases of obligatory sloppy readings for reflexives that are described by Sells et al., since they consider only cases where the reflexive and its antecedent are arguments of the same predicate, in which a relation-reducing operation can apply.

However, this solution would be inappropriate in cases where the reflexive and its antecedent are clearly arguments of different predicates, where a relation-reducing operation cannot apply. Although such cases are difficult to find, it may be that the Serbo-Croatian reflexive *sebe* (genitive *svoje*) is such a case.

The following sentence has only a sloppy reading (Draga Zec, personal communication):

- (76) Petar je sakrio sto hiljada dolara ispod svoje kuće a to je učinio
 Petar Aux hid one hundred dollars underneath self's house and that Aux did
 i Pavle
 also Pavle

“Petar hid one hundred dollars underneath self's house, and Paul did (that) too.”

The only reading available for this sentence is that Paul hid one hundred dollars under his own house.

We postulate that the reflexive *sebe* in Serbo-Croatian engenders primary as opposed to secondary occurrences, which would then be subject to the primary occurrence constraint. This would also be true of the English reflexive for those speakers who find that strict readings with reflexives are unacceptable.

In short, a variety of syntactic constructions give rise to multiple primary occurrences of parallel elements: control, both of the type seen in English and of the type seen in Serbo-Croatian, and reflexivization in some dialects of English and in Serbo-Croatian.

5.2.2 Antecedent-anaphor constraints

Missing readings with multiple occurrences of anaphora In cases where there are two pronouns coreferent with the parallel element in the source, one might expect that each pronoun would give rise to either a strict or a sloppy reading, giving a total of four readings for the target clause. This does not seem to be the case, however; one of the readings is systematically missing.

Dahl (1974) notes that the following sentence, with two occurrences of pronouns in the source clause, has only three and not four interpretations:

- (77) a. Bill believed that he loved his wife, and Harry did too.
 b. *Harry believed that Bill loved Bill's wife.*
 b''. *Harry believed that Harry loved Harry's wife.*
 b''. *Harry believed that Harry loved Bill's wife.*
 b'''. * *Harry believed that Bill loved Harry's wife.*

In this case, the missing reading corresponds to the following unifier for (78a):

- (78) a. $P(\text{bill}) = \text{believe}(\underline{\text{bill}}, \text{love}(\text{bill}, \text{wife-of}(\text{bill})))$
 b. $P \mapsto \lambda x. \text{believe}(x, \text{love}(\text{bill}, \text{wife-of}(x)))$

Other examples illustrate a similar phenomenon. Sag (1976, page 183) observes that this sentence has only three readings, not four:

- (79) a. Edith said that finding her husband nude had upset her, and Martha did too.
 b. *Martha said that finding Martha's husband nude had upset Martha.*
 b'. *Martha said that finding Edith's husband nude had upset Edith.*
 b''. *Martha said that finding Edith's husband nude had upset Martha.*
 b'''. * *Martha said that finding Martha's husband nude had upset Edith.*

The interpretation paraphrased in 79b''' is missing. The unifier for P for the missing reading is:

- (80) a. $P(edith) = say(\underline{edith}, upset(finding-nude(husband-of(edith)), edith))$
 b. $P \mapsto \lambda x. say(x, upset(finding-nude(husband-of(x)), edith))$

Examples (77) and (79) illustrate a constraint on relation formation. In example (77), the position corresponding to the pronoun “his” may not be abstracted unless the position corresponding to “he” is also abstracted. The constraint does not seem to correlate with linear order of the pronouns, though; in example (77) the position corresponding to the rightmost pronoun may not be abstracted unless the position corresponding to the leftmost pronoun is also abstracted. The reverse is true in example (79), however, where it is the position corresponding to the rightmost pronoun that must be extracted.

Although these examples show that the proper generalization about the ordering between pronominal positions does not have to do with linear order, they are consistent with the hypothesis that the ordering correlates with depth of syntactic embedding. In both example (77) and example (79), if the position corresponding to a more deeply embedded pronoun is abstracted over, the position corresponding to a less deeply embedded pronoun must also be abstracted over.

However, example (81) shows that the ordering between positions may not be dependent on syntactic facts at all. Recall that sentence (55), repeated here, was not associated with the reading in (81b), paraphrased in (81c):

- (81) a. John revised his paper before the teacher did, and Bill did too.
 b. *before(revise(john, paper-of(john)),
 revise(teacher, paper-of(john))) and
 before(revise(bill, paper-of(john)),
 revise(teacher, paper-of(bill)))*
 c. John and then the teacher revised John's paper; Bill revised John's paper before the teacher revised Bill's paper.

On the excluded reading, the source clause is “John revised his paper before the teacher did”, and the target clause is “Bill did too”. “John” and “Bill” are the parallel elements. The unifier for P for this reading is:

- (82) a. $P(\textit{john}) = \textit{before}(\textit{revise}(\textit{john}, \textit{paper-of}(\textit{john})),$
 $\textit{revise}(\textit{teacher}, \textit{paper-of}(\textit{john})))$
- b. $P \mapsto \lambda x. \textit{before}(\textit{revise}(x, \textit{paper-of}(\textit{john})),$
 $\textit{revise}(\textit{teacher}, \textit{paper-of}(x)))$

These various missing readings can be captured by positing a linking relationship between the semantics of pronouns and that of their antecedents, and generalizing it to include the relation between the semantics of terms induced by ellipsis and that of their source parallel element. Under a suitable definition of this generalized antecedent linking, all of the cases here can be captured by requiring that if an occurrence is abstracted over, so must its generalized antecedent.

Apparent Syntactic Constraints Finally, we turn to some simple examples that seem to lack any readings whatsoever. Consider the following examples, where “Mary” is taken to be the antecedent of “she”:²¹

- (83) a. * John gave Mary everything she did.
- b. * John likes Mary, and she does too.

These judgments would follow from an analysis on which syntactic structure is copied from the source to the target, as the sentences with copies in place violate constraints on binding.

- (84) a. * John gave Mary everything she_i gave Mary_i.
- b. * John likes Mary, and she_i likes Mary_i too.

However, a simple copying analysis faces problems in accounting for grammatical examples of similar structure:

- (85) a. John got to Sue_i's apartment before she_i did.
- b. John voted for Sue_i because she_i told him to.

On a copying analysis, these examples would be incorrectly predicted to be ungrammatical, just as their copied versions are:

- (86) a. *John got to Sue_i's apartment before she_i got to Sue_i's apartment.
- b. *John voted for Sue_i because she_i told him to vote for Sue_i.

²¹Examples of this type are due to Fiengo and May (1990).

The resolution of this puzzle remains an open question, as does its incorporation in the present analysis; for discussion of the problem, see Hellan (1988), Fiengo and May (1990), and Kitagawa (1991).

Another example that seems to argue for a quite superficial analysis of ellipsis is the following, due to Yoshihisa Kitagawa (personal communication to Peter Sells):

(87) John thinks that Mary will revise his paper before Bill will.

on the reading in which Mary revises John's paper and Bill revises his own paper. We find the intuition questionable, but it is clearly problematic for our (and many others') analysis if the reading is deemed to be available.

Finally, examples which seem to argue for the presence of a gap in the ellipsis site include the following:

(88) *John met everyone that Peter wondered when he could. (Haik, 1987, p. 511)

(89) *Tom visited everyone who told Sue where to. (Haik, 1985, p. 218)

On the assumption that long-distance dependencies are syntactically constrained and that subjacency violations involve an improper syntactic relation between a filler and a gap, these examples indicate that the ellipsis site contains a gap at syntactic structure.

6 Conclusion

The underlying idea in the analysis of ellipsis that we have presented here—namely, the construction of higher-order equations on the basis of parallel structures, and their solution by unification—has been exemplified primarily by the verb-phrase ellipsis construction. However, many other elliptical phenomena and related phenomena subject to multiple readings akin to the strict and sloppy readings discussed here may be analyzed using the same techniques. The ambiguities in cleft sentences such as

(90) It is Dan who loves his wife.

and interpretation of “only” with respect to its focus, as in

(91) Only Dan loves his wife.

as well as more standard elliptical phenomena such as stripping and comparative deletion can be analyzed in this way as well, making a broad range of predictions as to the space of possible readings and their interaction with other semantic phenomena. It remains for future work to test these potential applications more fully.

We adduce three advantages of the analysis of elliptical constructions presented here over previous alternatives. First, it is in certain respects simpler, in that it requires no postulation of otherwise unmotivated ambiguities in the source clause. Second, it is more accurate in its predictions, especially in allowing readings disallowed in identity-of-relations analyses. Third, it is methodologically preferable in that the analysis follows directly from a semantic statement of the ellipsis problem with little stipulation. The operation on which it relies, higher-order unification, is semantically sound in that the results it produces are determined by the meanings of phrases directly rather than by the form of the representations encoding those meanings, as operations of deletion or copying of portions of such representations are.

Acknowledgements

We would like to thank the following people for helpful discussion: Hiyan Alshawi, Sam Bayer, Joan Bresnan, Mark Gawron, Kris Halvorsen, Dan Hardt, Julia Hirschberg, David Israel, Mark Johnson, Ron Kaplan, Lauri Karttunen, Shalom Lappin, Richard Larson, Peter Ludlow, John Maxwell, Richard Oehrle, Stanley Peters, Hub Prust, Steve Pulman, Mats Rooth, Ivan Sag, Peter Sells, Gregory Ward, Michael Wescoat, Annie Zaenen, Draga Zec, and two anonymous reviewers. The written comments alone that we received prior to publication ran to well over 50 pages, longer, in fact, than the paper itself. We regret that we could not include discussion of all the important issues that they raised.

References

- Johan van Benthem. 1989. Categorical grammar and type theory. *Journal of Philosophical Logic*. To appear.
- Gennaro Chierchia. 1983. Outline of a semantic theory of (obligatory) control. In Michael Barlow, Daniel Flickinger, and Michael Wescoat, editors, *Proceedings of the West Coast Conference on Formal Linguistics 2*, pages 19–31. Stanford Linguistics Association, Stanford University.
- Gennaro Chierchia. 1984. Anaphoric properties of infinitives and gerunds. In Mark Cobler, Susannah MacKaye, and Michael Wescoat, editors, *Proceedings of the West Coast Conference on Formal Linguistics 3*, pages 28–39. Stanford Linguistics Association, Stanford University.
- Gennaro Chierchia. 1988. Dynamic generalized quantifiers and donkey anaphora. In M. Krifka, editor, *Proceedings of the 1988 Tübingen Conference*. Seminar für Natürliche Sprachliche Systeme der Universität Tübingen, November.
- Robin Cooper. 1983. *Quantification and Syntactic Theory*, volume 21 of *Synthese Language Library*. D. Reidel, Dordrecht.
- Östen Dahl. 1972. On so-called ‘sloppy identity’. In *Gothenburg Papers in Theoretical Linguistics*, volume 11. University of Göteborg.
- Östen Dahl. 1974. How to open a sentence: Abstraction in natural language. In *Logical Grammar Reports, No. 12*. University of Göteborg.
- Robert Fiengo and Robert May. 1990. Anaphora and ellipsis. MS, City University of New York and University of California, Irvine.
- Harvey Friedman. 1975. Equality between functionals. In R. Parikh, editor, *Lecture Notes in Mathematics 453*, pages 22–37. Springer-Verlag, Berlin, Germany.
- Mark Gawron and Stanley Peters. 1990. *Anaphora and quantification in Situation Semantics*. CSLI/University of Chicago Press, Stanford University. CSLI Lecture Notes, Number 19.
- Gerald Gazdar, Ewan Klein, Geoffrey K. Pullum, and Ivan A. Sag. 1985. *Generalized Phrase Structure Grammar*. Harvard University Press, Cambridge, MA.
- J. Groenendijk and M. Stockhof. 1987. Dynamic Montague Grammar. Paper presented at the Workshop on Discourse Representation Theory, Stuttgart, West Germany, December.
- Isabelle Haïk. 1985. *The Syntax of Operators*. Ph.D. thesis, MIT.
- Isabelle Haïk. 1987. Bound VPs that need to be. *Linguistics and Philosophy*, 10:503–530.

- Jorge Hankamer and Ivan A. Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry*, 7(3):391–428.
- Irene Heim. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts-Amherst.
- Lars Hellan. 1988. *Anaphora in Norwegian and the Theory of Grammar*. Foris Publications, Dordrecht.
- J. Roger Hindley and Jonathon P. Seldin. 1986. *Introduction to Combinators and λ -Calculus*. Cambridge University Press, Cambridge, England.
- Julia Hirschberg and Gregory Ward. 1991. Accent and bound anaphora. *Cognitive Linguistics*. To appear.
- Paul Hirshbühler. 1982. VP deletion and across-the-board quantifier scope. In James Pustejovsky and Peter Sells, editors, *Proceedings of NELS 12*. GLSA, University of Massachusetts-Amherst.
- Jerry R. Hobbs and Stuart M. Shieber. 1987. An algorithm for generating quantifier scopings. *Computational Linguistics*, 13:47–63.
- Gérard Huet and Bernard Lang. 1978. Proving and applying program transformations expressed with second-order patterns. *Acta Informatica*, 11(1):31–55.
- Gérard Huet. 1975. A unification algorithm for typed $\bar{\lambda}$ -calculus. *Theoretical Computer Science*, 1:27–57.
- Ray S. Jackendoff. 1972. *Semantic Interpretation in Generative Grammar*. MIT Press, Cambridge, MA.
- Hans Kamp. 1981. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*, pages 277–321, Amsterdam. Mathematical Centre.
- Yoshihisa Kitagawa. 1991. Copying identity. *Natural Language and Linguistic Theory*. To appear.
- Joachim Lambek. 1980. From λ -calculus to cartesian closed categories. In J.P. Seldin and J.R. Hindley, editors, *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, pages 375–402. Academic Press, London.
- Shalom Lappin. 1984. VP anaphora, quantifier scope, and logical form. *Linguistic Analysis*, 13(4):273–315.
- Fernando C. N. Pereira. 1990. Categorical semantics and scoping. *Computational Linguistics*, 16(1):1–10.
- S. G. Pulman. 1988. A contextual reasoning and cooperative response framework for the Core Language Engine. Internal report, SRI Cambridge, Cambridge, England.

- Tanya Reinhart. 1983. *Anaphora and Semantic Interpretation*. University of Chicago Press, Chicago.
- Craige Roberts. 1987. *Modal Subordination, Anaphora, and Distributivity*. Ph.D. thesis, University of Massachusetts-Amherst.
- Ivan A. Sag. 1976. *Deletion and Logical Form*. Ph.D. thesis, MIT.
- Traugott Scheibe. 1973. Zum Problem der grammatisch relevanten Identität. In F. Kiefer and N. Ruwet, editors, *Generative Grammar in Europe*, pages 482–527. D. Reidel Publishing Company, Dordrecht.
- Len Schubert and F J. Pelletier. 1982. From English to logic: Context-free computation of ‘conventional’ logical translations. *American Journal of Computational Linguistics*, 10:165–176. Reprinted in Grosz *et al.*, 1986.
- Peter Sells, Annie Zaenen, and Draga Zec. 1987. Reflexivization variation: Relations between syntax, semantics, and lexical structure. In Masayo Iida, Stephen Wechsler, and Draga Zec, editors, *Working Papers in Grammatical Theory and Discourse Structure*, pages 169–238. CSLI/University of Chicago Press, Stanford University. CSLI Lecture Notes, Number 11.
- Mark J. Steedman. 1990. Gapping as constituent coordination. *Linguistics and Philosophy*, 13(2):207–263.
- Bonnie Lynn Webber. 1978. *A Formal Approach to Discourse Anaphora*. Ph.D. thesis, Harvard University.
- Michael Wescoat. 1989. Sloppy readings with embedded antecedents. ms, Stanford University.
- Edwin Williams. 1977. Discourse and logical form. *Linguistic Inquiry*, 8(1):101–139.
- Draga Zec. 1987. On obligatory control in clausal complements. In Masayo Iida, Stephen Wechsler, and Draga Zec, editors, *Working Papers in Grammatical Theory and Discourse Structure*, pages 139–168. CSLI/University of Chicago Press, Stanford University. CSLI Lecture Notes, Number 11.

Sources of Attested Examples

Noam Chomsky. 1982. *Noam Chomsky on the Generative Enterprise*. Foris Publications, Dordrecht.

K. P. Mohanan. 1983. Functional and anaphoric control. *Linguistic Inquiry*, 14(4):641–674.

John Rettie. 1990. Yeltsin wants Russia to go it alone. *The Guardian*, 23 May.

R. Roeper. 1990. *Chicago Sun-Times*, 8 January. Cited by James McCawley, “1990 Linguistic Flea Circus”, unpublished manuscript.

Stuart M. Shieber. 1989. *Parsing and Type Inference for Natural and Computer Languages*. Ph.D. thesis, Stanford University.