Lecture Notes on Ellipsis

Irene Heim
heim@mit.edu
M.I.T. 24.956, Spring 1996

I. Sag's recoverability condition for VP ellipsis

1. Background: Deletion under identity

Under what conditions is a sentence with a deleted VP interpretable, and how is it interpreted? The simplest guess would be that a deleted VP is interpretable whenever it is possible for the listener to guess which property the speaker intends to refer to. Unfortunately, this does not appear to be sufficient.

(1) * The garbage can is full. I hope that YOU will [$_{VP}$], for a change.

Compare:

(2) The garbage can is full. I think it's YOUR turn.

(3) It's time to take out the garbage. I hope that YOU will [$_{VP}$], for a change.

It seems that a deleted VP can only be interpreted if it has an antecedent: another VP in the surrounding discourse from which it gets to pick up its content.

How exactly does this work? How does the antecedent VP contribute to and constrain the interpretation of the deleted VP?

When VP ellipsis began to be studied by generative grammarians, the first hypothesis they explored was that a VP was optionally deleted under identity with another VP. This deletion was assumed to take place after the stage in the derivation which feeds semantic interpretation. For instance, to obtain the surface form and meaning of (4),

(4) John didn't go out. He wasn't allowed to.

we can first generate the corresponding non-elliptical text (4') and its semantic interpretation:

(4') John didn't go out. He wasn't allowed to go out.

and then delete the second copy of *[$_{VP}$ go out]* in (4') under identity with the first.

This approach explains why the missing VP needs an antecedent, and it answers the question how the choice of antecedent determines the semantic interpretation of the elliptical sentence. Its limitations are not immediately obvious. Let's try to apply it to a series of increasingly challenging examples.

Assuming that the grammar is organized as in (5),

(5)

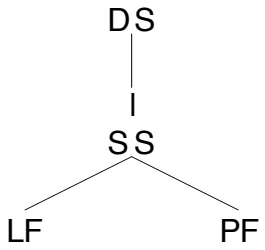

VP Deletion must take place on the SS-to-PF branch. That way, the VP will be present for semantic interpretation but absent for prosodic realization.

For a simple example like (6), the SS-to-PF derivation proceeds as in (6').

(6) John left. Bill didn't.

(6') SS: John PAST [$_{VP}$ leave]. Bill didn't [$_{VP}$ leave].
⇒ (VP Deletion) John PAST [$_{VP}$ leave]. Bill didn't [$_{VP}$ ∅].
⇒ (other PF-rules) John [$_{VP}$ left]. Bill didn't [$_{VP}$ ∅].

This example shows that the identity which licensed deletion need not obtain on the surface. The overt counterpart of (6) would be *John left. Bill didn't leave*, not *John left. *Bill didn't left.* Identity must obtain at the point where deletion applies, but may be obliterated by subsequent PF-operations like (in this case) affix movement and morphological spell-out.

Sag (1976) begins with a critcal discussion of the deletion-under-identity approach. He finds problems with the following types of examples. (Many of these had already been discussed in earlier literature -- refer to Sag for proper attributions.)

Polarity alternations

(7) You didn't eat anything, but I did.

If (7) is derived by deletion under identity, then either the deleted VP must be *eat anything*, or else the SS representation of the antecedent must not be *eat anything*. The first possiblity implies that it is okay for *any* to appear outside of the domain of an appropriate negative trigger at SS, as long as it doesn't surface at PF. The second possibility implies that the SS

representation of the antecedent VP contains either *something* or an abstract representation still unmarked for polarity, which then is changed to or completed to *anything* later on the PF branch. Both options thus lead to the conclusion that the distribution of negative polarity items is determined at PF. This is incompatible with evidence that NPI distribution depends on scope relations, and contradicts all current theories of NPI licensing.

Features of pronouns

(8) I turned in my homework, but most of the other students didn't.

The deletion-under-identity analysis seems to predict that the missing VP should be *turn in my homework*, and thus that the second sentence means that most other students didn't turn in my homework. This is one possible reading of (8), but not the only one. Another available interpretation is that most other students didn't turn in their homework. How is this reading generated? Is there a version or amendment of the deletion-under-identity approach that correctly accounts for it?

Should we fix up the semantic rules in such a way that *turn in my homework* can mean either 'turn in my homework' or 'turn in their homework'? Probably not, because we would have to stipulate then that the second interpretation arises only when the VP is not overt.

Should we relax the deletion rule in such a way that it doesn't require strict identity, but tolerates a limited number of differences between delendum and antecedent, for instance those pertaining to person, number and gender features of pronouns? This too seems problematic, because we can then no longer explain why (8) doesn't have many additional readings as well, for instance one where the second sentence is equivalent to 'most other students didn't turn in your homework'.

Here is a more promising line that might be taken by the defender of deletion-under-identity: At S-structure, anaphoric pronouns do not yet have features; they acquire them (by agreement with their antecedents) only later in the derivation towards PF. The basic idea is that this way, one possible SS representation for (8) could look as in (8') (where "PRONOUN" stands for a pronominal NP without feature specifications).

(8') I PAST [$_{VP}$ turn in PRONOUN's homework],
but most other students didn't [$_{VP}$ turn in PRONOUN's homework].

At this level, we still have the identical VPs required for deletion.

For this solution to function, we have to make sure that the representation derived from (8') on the LF branch receives the intended meaning, and that all undeleted instances of PRONOUN are

spelled out on the PF branch in their correct forms. A bit of reflection shows that the calculations on both branches must involve choosing an antecedent for each PRONOUN: the pronoun's semantic import depends on which NP it is bound by (or corefers with), and its surface form depends on which NP it agrees with. But if the pairing of pronouns with their antecedents proceeds independently on the two branches, we make inadequate predictions.

Consider, for instance, the following LF/PF derivation from the SS (8'): On the LF branch (or in the semantics), we make both PRONOUNs anaphoric to *I*. So the meaning obtained is that the speaker, but not most other students, turned in the speaker's homework. On the PF branch, on the other hand, we choose not to delete the VP and then we let each PRONOUN agree with the subject of its own clause. The surface form thus obtained is *I turned in my homework, but most other students didn't turn in their homework*. The problem is that this pair of meaning and form is not in fact grammatical.

The moral seems to be that we must fix the anaphoric relations before the derivations of LF and PF branch apart, i.e. by SS. So (8') cannot be the complete SS, we rather need something like (8'').

(8'') I**1** PAST [$_{VP}$ turn in PRONOUN**1**'s homework],
but most other students**2** didn't [$_{VP}$ turn in PRONOUN**2**'s homework].

But now we have destroyed the strict identity between the VPs: there is a mismatch of indices.

Let us break off the discussion of this example here. We will see shortly that it is not straightforward in other theories either. We have seen that its treatment in the deletion-under-identity theory raises non-trivial problems, but it remains to be seen whether this amounts to a decisive argument against such a theory in favor of another one.

Quantifiers and pronouns

(9) Sandy greeted everyone before Betsy did.

This sentence too has two readings of which only one is straightforwardly predicted by the deletion-under-identity analysis. If the deleted VP is *greet everyone*, (9) means that Sandy greeted everyone before Betsy greeted everyone. But it can also mean that, for each person x, Sandy greeted x before Betsy greeted x. What VP do we have to delete to obtain this reading, and under what assumptions would that VP meet the identity condition?

If we want to maintain the deletion-under-identity approach in the face of examples like (9), we are pushed towards a conception of SS that brings it closer to what is commonly considered LF.

Concretely, suppose that quantifiers were in a position c-commanding their semantic scope at SS. So the source of (9) could be an SS in which *everyone* c-commands the entire sentence:

(9') everyone$_1$ [Sandy PAST greet e$_1$ before Betsy did greet e$_1$]

(9') would be a suitable input for VP-deletion under identity, and also a suitable basis for an LF-derivation that yields the right meaning. (Assume, e.g., that nothing significant changes on the LF-branch, so that the LF looks just like (9'), and the semantic rules treat e_1 as a variable bound by the quantifier *everyone*$_1$.) But the PF-derivation from (9') would have to be non-trivial. Specifically, we need a PF-rule of Quantifier Lowering that substitutes *everyone* for the first occurrence of e_1. Also, if we allow any derivations from (9') in which the second VP is not deleted, we need rules that can spell out the second occurrence of e_1 as *them* .

"Antecedent-contained" deletion

(10) Alan will eat anything you want him to.

The deleted VP in (10) cannot possibly be identical to another surface VP in (10). The only other VP in (10) is *eat anything you want him to*. This in turn would have to result by deletion from *eat anything you want him to eat anything you want him to*, whose source in turn should have been ... . We are caught in infinite regress.

On the other hand, if what is deleted in (10) is what we see in its non-elliptical paraphrase, namely the VP *eat t* (with a wh-trace in object position), then deletion evidently does not require identity.

Again, the only straightforward way to maintain deletion under identity for this example is to posit an SS from which the surface form would be derived by Quantifier Lowering:

(10') [anything you want him to eat e$_1$] Alan will eat e$_1$.

Ambiguous antecedent VPs

What happens when the antecedent of a VP is ambiguous? Cases in point (involving a variety of different kinds of ambiguity) are (11) - (13).

(11) The children are ready to eat, but the chickens are not.

(12) John hit the wall yesterday, and Pete did today.

(13) Mary assigned too many problems more than once. Bill didn't.

Be ready to eat can mean 'be ready to eat something' or 'be ready to be eaten'. *Hit the wall* can have an idiomatic or a literal meaning. *Assign too many problems more than once* displays a quantifier-scope ambiguity: on one reading, it is true of x iff there are too many problems that x assigned more than once, and on another reading, it is true of x if there is more than one occasion on which x assigned too many problems.

In each of these cases, the deletion-under-identity analysis seems to predict that the sentence as a whole should be 4 ways ambiguous: the SSs of (11) - (13) each contain two ambiguous VPs. One observes, however, that these sentences are only 2 ways ambiguous: whichever reading is chosen for the antecedent VP must be chosen for the deleted VP as well. The "mixed" readings, in which the antecedent VP has one reading and the deleted VP the other, are unavailable.

Is this observation really problematic for the deletion-under-identity theory? That depends. To argue that it is, we must first be sure that the two readings are not only homophonous, but have identical representations at SS. In (11) and (12), this is far from evident[1], but it is arguably the case in (13). Still, the deletion-under-identity theorist may be able to rule out the mixed readings by an independent mechanism. Indeed, it seems that even when we don't delete anything, but instead have two overt identical VPs in short succession, such readings are hardly possible. Try to find a context in which (13') could naturally be understood as expressing one of the mixed readings!

(13') Mary assigned too many problems more than once. Bill didn't assign too many problems more than once.

Disambiguation effects in the antecedent sentence

A slightly different issue is raised by (14).

(14) Someone hit everyone, and then Bill did.

There is a potential scope ambiguity in the sentence containing the antecedent VP, and in isiolation it could indeed mean either $\exists x.\forall y.hit(x,y)$ or $\forall y.\exists x.hit(x,y)$. If the derivation of (14) proceeds by deleting *hit everyone* after the PF-derivation branches off from the LF-derivation, it is not apparent what principles should prevent both choices of scope-disambiguation for the first

---

[1]The infinitival complement of *ready* might be have the SS $[_{CP}\ Op_i\ [_{IP}\ PRO_{arb}$ *to eat* $t_i]\ ]$ on one reading, and the SS $[_{CP}\ [_{IP}\ PRO_i$ *to eat*$]\ ]$ on the other (where *i* is the index of *the chickens* in each case). Idiomatic *hit the wall* may simply be its own lexical item.

conjuncts on the LF-branch. The result could be a reading with the truthconditions ∀y.∃x.hit(x,y) & ∀y.hit(Bill,y), but this is not in fact attested for (14).

Again, it is difficult to be sure that this is really a defect specific to the deletion-under-identity hypothesis. Is it really easier to obtain the above-mentioned reading for the non-elliptical variant (14')?

(14') Someone hit everyone, and then Bill hit everyone.

If not, then the absence of such a reading in (14) might have an independent explanation compatible with the deletion-under-identity treatment of ellipsis. We can't decide before we have some idea of what that independent explanation might consist of.

Summary

The treatment of VP ellipsis by means of a rule that deletes VPs under identity has not been definitively refuted. We have learned, however, that such a treatment can only be maintained if the input to deletion (which by definition cannot be on the LF branch) is a representation that is semantically explicit in several respects, including dsiambiguation of anaphoric relations and quantifier scope. This observation has led many scholars, including Sag, to the suspicion that it might be misguided to impose a syntactic identity requirement on VP deletion at all. The natural alternative is that semantic equivalence with the antecedent might be the decisive precondition for deletion. This is the next hypothesis that we will spell out and explore.

2. Deletion under semantic equivalence

Informally, the hypothesis we will explore is that a VP may be deleted just in case it receives the same semantic interpretation as its antecedent. The idea is that semantic equivalence is the decisive factor, whereas syntactic identity per se is irrelevant. Of course, there will be many cases where a semantically equivalent VP is also syntactically identical, in which case the new hypothesis won't make different predictions from the old one. But we expect discrepancies in two kinds of cases: first, when two dissimilar VPs nevertheless mean the same thing, and second when two identical-looking VPs have different meanings. Indeed, many of the examples we were concerned with in the last section can be viewed as exemplifying one or the other of these two situations, and a quick review of them suggests that our new hypothesis may suit them well.

Re (7): The VPs *eat something* and *eat anything* both express the property of eating something. The polarity features that distinguish *something* from *anything* affect the distribution of these NPs in well-formed LFs, but they are not themselves semantically interpreted; both NPs denote the same quantifier (see e.g. Ladusaw, Linebarger). So our new hypothesis predicts that *eat*

something may be deleted under semantic identity with *eat anything*, which is exactly what seems to happen in (7).

Re (8): It is plausible that the VPs *turn in my homework* and *turn in their homework* can both express the property of turning in one's homework, i.e., that property which an arbitrary individual x has if and only if x turns in x's homework. If this is the correct semantic analysis of these VPs, then we expect one of these VPs to license deletion of the other, despite lack of complete syntactic identity. (We will have more to say about this sort of example.)

Re (9): Concentrate on the reading where (9) means that for every person x, Sandy greeted x before Betsy greeted x. On this disambiguation of the sentence, there is really no such thing as the "meaning of the VP *greet everyone*." The meanings of *greet* and *everyone* are never composed into a new meaning during the interpretation of this sentence (under this reading). It follows that under our new hypothesis, the surface VP *greet everyone* could not possibly be an antecedent for deletion of any other VP.

The picture changes if we consider the meanings of the VPs in the LF representation of (9). Suppose the LF of (9) (under the relevant reading) is the same as that of the non-elliptical sentence *Sandy greeted everyonewhen Betsy greeted them* on the same reading, viz. (9").

(9") everyone$_1$ [Sandy PAST [$_{VP}$ greet t_1] when Betsy did [$_{VP}$ greet them$_1$]]

Since both the QR-trace t_1 and the bound pronoun *them$_1$* are interpreted as variables, more precisely, as instances of the same variable (= the index 1), the two VPs in this LF receive identical meanings. So the first one licenses deletion of the second by our hypothesis.

There is something a little worrisome about this line of explanation: the deletion evidently applies on the PF branch, yet it is licensed by a relation between LF constituents. (Sag calls this the "correspondence problem".) We will return to this point when we make the analysis more precise.

Re (10): The situation here is similar to the one in (9); we postpone a more detailed analysis.

Re (11) - (13): In each of these examples, the surface form of the antecedent VP is ambiguous, and so is the surface VP that appears to have been deleted. Our new hypothesis says that deletion requires sameness of meaning and thus implies that the legitimacy of deletion depends on the choice of disambiguation. More specifically, it predicts deletion is licensed if and only if the antecedent and deleted VP are disambiguated in favor of the same reading. This is what we observe.

Re (14): This example cannot quite be explained by means of the current informal version of our hypothesis. We will return to it when we have a precise implementation.

Excursion: Sag's worries about indefinites

What I have presented so far is essentially Sag's reasoning. Yet he stopped short of endorsing the hypothesis that VP deletion is licensed by semantic equivalence. What deterred him were the following examples.

(15) They caned a student severely when I was a child, but not like Miss Grundy did yesterday. (credited to Nunberg)

(16) Jane ended up marrying a doctor, although she didn't want to. (credited to Kuno)

Sag thought that these were counterexamples, i.e., examples in which deletion was okay despite a lack of semantic equivalence between deleted VP and antecedent. He observed that in (15), the antecedent occurrence of *cane a student* contained a "generic" instance of *a student*, whereas the deleted VP contained a "specific" one. Similarly for (16), where he described the occurrence of *a doctor* in the antecedent as "specific," and the one in the deleted VP as "non-specific". He thought it was obvious that these differences in the readings of indefinites meant that the VPs containing them were not semantically equivalent.

However, this is not so obvious after all. Only a precise semantic analysis of the so-called "specific", "non-specific", and "generic" readings of indefinites will allow us to determine which VPs are semantically equivalent and which are not. For (16), such a precise analysis is by now widely available and generally accepted: Both occurrences of *a doctor* are existential quantifiers with scope in the minimal VP containing them. The occurrence in the antecedent is "specific" only in the sense that no other operator happens to have scope over it, whereas the occurrence in the deletion site is in the scope of *ddn't want* (a negation and a quantifier over possible worlds). On this analysis, (16) is a perfectly well-behaved case of deletion under semantic equivalence.

The case of (15) is more complicated. See Tomioka (LSA, Jan. 96) for a recent argument that (15) is not a genuine counterexample to semantic equivalence either. This is what I will assume henceforth. Moreover, I assume that Sag would have endorsed the semantic identity theory if it hadn't been for just these examples, and so I take the liberty of crediting it to him in the sequel.

Implementation of the "semantic equivalence" condition on VP deletion

The hypothesis to be implemented makes the applicability of a PF-operation (deletion) dependent upon a relation that is only defined for LF-constituents (semantic equivalence). To circumvent this difficulty, I will assume that the antecedency relation among VPs is represented before the PF/LF fork, by coindexing. We can thus state two separate principles:

(17) (a) A VP that is coindexed with another VP may delete.
(b) Coindexed VPs are semantically equivalent.

I will refer to (17) as "Sag's Recoverability Condition"[2]. (17b) can be derived from a general semantic rule for indexed LF-nodes.[3]

Apart from (17), we need a definition of "semantic equivalence" and some concrete assumptions about the derivation and interpretation of Logical Form. We assume that the semantic component of the grammar specifies an interpretation function [[ ]] which maps pairs of LF-constituents α and variable assignments g to intensions $[[\alpha]]^g$.

(18) Two LF-nodes α and β are <u>semantically equivalent</u> iff
for all variable assignments g, $[[\alpha]]^g = [[\beta]]^g$.

The assumptions about LFs and their interpretation will be mostly standard; we specify details as we get to examples.

The derivation and licensing of examples like *John left. Bill didn't* is obvious. It is also clear how the undesirable "mixed reading" in, e.g. *John hit the wall yesterday, and Pete did today* is excluded. That's precisely what the semantic equivalence requirement is designed to do.

<u>Quantifiers in VP</u>

Most of our other examples from the introductory discussion involve quantifiers somewhere in the VPs. To deal with these, we have to clarify some general points about the LF-representation and interpretation of quantification. Consider what should be an entirely unproblematic case of VP ellipsis:

(19) (a) John answered every question. (b) Bill did not.

(19b) means that Bill did not answer every question. The ellipsis is presumably licensed by the fact that both the deleted VP and its antecedent express the property of answering every question. There is a potential problem, however, with getting this prediction from the theory. Sag, like May (1977) and many other contemporaneous writers on LF syntax, assumed that quantifying NPs like *every question* (and also operators like negation) could not be interpreted in situ, but

---

[2]See below for some comments comparing Sag's original proposal with the present implementation.

[3]See e.g. Heim (1982), (1993).

had to adjoin to a sentential domain in the derivation of LF. This implies that the LF of (19) should look like this:

(20) (a) every question$_i$ [$_S$ John PAST [$_{VP}$ answer t_i]$_1$]
(b) not [$_S$ every question$_j$ [$_S$ Bill PAST [$_{VP}$ answer t_j]$_1$]

In this structure, deletion is predicted to be licensed only if i=j, i.e., if the two occurrences of the QRed quantifier happen to bind instances of the same variable. But it is problematic to allow such coindexing in principle. If it were always allowed to choose indices in applications of movement with complete freedom, in particular to use the same index in different applications, the current approach to ellipsis would allow far too many readings for ellided VPs. For instance, there would be no reason why the second VP in (21) shouldn't delete just as well as that in (20).

(21) (a) every question$_2$ [$_S$ John PAST [$_{VP}$ answer t_2]$_1$]
(b) not [$_S$ many phone calls$_2$ [$_S$ Bill PAST [$_{VP}$ answer t_2]$_1$]

But then the text in (19) ought to be able to mean that John answered every question and Bill didn't answer many phone calls. Clearly, the intuition that deletion is recoverable only under semantic equivalence is not adequately captured on these assumptions.

Sag argues, in effect, that to implement this intuition appropriately, one must disallow <u>any</u> use of coindexing in places where this does not express a semantic relation (binding or coreference). I state this convention explicitly for further reference:[4]

(25) <u>No Meaningless Coindexing</u>:
If an LF contains an occurrence of a variable **i** that is bound by a node α, then all occurrences of **i** in this LF must be bound by the same node α.

This applies not only to the LFs of individual sentences, but also to the LF-representations of multi-sentential texts. It prohibits two kinds of "spurious" (i.e., semantically inert) coindexing: cases where the same variable occurs both free and bound, and cases where bound occurrences of the same variable have different binders. Simple illustrations can be given if we apply (25) to a standard formal language such as predicate logic. It disallows formulas like (26a, b).

---

[4]It would not be necessary to stipulate (25) if we adopted a different notation for variables and variable-binding at LF, such as the one sketched in the logic text by Kalish, Montague & Marr, which employs blanks and connecting lines instead of the usual letters or numerical subscripts. As these authors and many others have pointed out, the customary notation is a historical compromise between conceptual appropriateness and typographical convenience.

(26) (a) * Fx & ∃x Gx

(b) * ∀x Fx & ∃x Gx

The expressive power of the language is unaffected by this convention, because every prohibited formula has semantically equivalent variants that are permissible, e.g. (26a',b') instead of (26a,b).

(26) (a') Fx & ∃y Gy

(b') ∀x Fx & ∃y Gy

Excursion on "alphabetic variance": Sag's definition differs from the standard definition in logic. [... Comment on the difference, and on the relations between alphabetic variance, semantic equivalence, and the "No Meaningless Indexing" convention.]

Assuming the No Meaningless Coindexing convention, an LF of the form of (20) must involve distinct variables i≠j, and thus does not license VP deletion under semantic equivalence. So how do we account for (19)? Two types of solutions suggest themselves. The first, familiar since Montague Grammar, is to introduce some semantic mechanism that allows in situ interpretation of quantifiers (for instance, assigning type <<et,t>,et> to transitive verbs, or shifting quantifying DPs to type <<e,et>,et>). Let us assume, for the time being, that we have something like this. Then we can assume an LF for (19) that is essentially like its SS:

(27) (a) John PAST $[_{VP}$ answer every question$]_1$

(b) Bill didn't $[_{VP}$ answer every question$]_1$

The second type of solution is to posit more internal LF-structure inside the VP, including a clausal node to attach quantifiers to. We will return to this option, and independent motivation for it, below.

If quantifiers can be interpreted in situ, or adjoined to non-clausal nodes like VP, then the present theory works not only for (19), but also for a number of other examples we encountered earlier:

(28) You didn't eat anything, but I did.
LF: you didn't $[_{VP}$ eat anything$]_1$, but I did $[_{VP}$ eat something$]_1$

Since *anything* and *something* are both existential quantifiers, *eat anything* and *eat something* express the same property, that of eating of something. Their differing shapes are no obstacle to deletion. We may assume, with most currents accounts of negative polarity, that polarity features are determined by scope relations at LF, so that only *something*, not *anything*, is well-formed in the LF of (28).

We also have an account of our earlier observation about (13) *Mary assigned too many problems more than once. Bill didn't.*. [...]

Let's look at (14) now.

(14) Someone hit everyone, and then Bill did.

One derivation is unproblematic: Generate SS (30) and leave everything in situ in the derivation of the LF.[5]

(30) someone PAST [hit everyone]$_1$, and then Bill did [hit everyone]$_1$

This derivation clearly meets Sag's recoverability condition and represents the grammatical reading of (14). The interesting question now is: Can we predict that there are no other readings, in particular none on which *everyone* scopes over *someone* in the antecedent clause? Such inverse scope would require *everyone* to raise above *someone* at LF, hence necessarily outside its VP. The LF of the antecedent would thus have to be essentially as in (31).

(31) everyone$_2$ [someone PAST [hit t_2]$_1$]

Given the semantic equivalence requirement, the deleted VP would have to have the same meaning as *hit* t_2, i.e., it would have to denote, for each variable assignment g, the property of hitting g(2). This is only possible if the deleted VP itself contains an occurrence of the variable 2. But then "No meaningless coindexing" implies that this variable-occurrence and t_2 in (31) must be either both free or both bound by the same binder. This is impossible, since the scope of *everyone*$_2$ in (31) does not extend over the elliptical clause.

The situation is quite different in an example like (9), where the elliptical clause is subordinate to the antecedent VP.

(9) Sandy greeted everyone before Betsy did.

Apart from the (by now familiar) derivation for the reading where the deleted VP means 'greet everyone', (9) allows another derivation:

(32) SS: Sandy PAST [greet everyone]$_1$ before Betsy did [greet him or her]$_1$
LF: everyone$_2$ [Sandy PAST [greet t_2]$_1$ before Betsy did [greet him-or-her$_2$]$_1$]

This derivation conforms to both "No meaningless coindexing" and Sag's recoverability, and it expresses the other reading that (9) was observed to have.

---

[5]We may also QR some or all of the quantifiers a little bit, as long as neither instance of *everyone* leaves its VP. This makes no difference.

My choice of the form *him or her* in this derivation was of course somewhat arbitrary. Sag's theory predicts only that there has to be an occurrence of the variable 2 in this LF-position. It says nothing about the morphological shape of this variable. *him or her* is a natural choice if we take it for granted that the deleted VP must meet whatever conditions would apply if it were to surface. But this assumption is by no means forced on us, and it is certainly not implied by Sag's theory of ellipsis. [More discussion of this later in the course ...]

Antecedent-contained deletion and the LF-representation of restricted quantifiers

The current proposal also works well for antecedent-contained deletion, provided we follow Sag in certain general assumptions about the LF-representation of restricted quantifiers and relative clauses. What we must not assume here is the standard Quine-Montague line of treating relative clauses as predicate abstracts and determiners as selecting arguments of the type of 1-place predicates. If we followed that line, then (33) would be the closest thing to a suitable LF for (21) that we would be able to generate. (*wh* is the non-overt relative pronoun.)

(21) Sue did everything she could.

(33) $\left[\text{every } [\text{thing } [\text{wh}_2 \text{ she could } [\text{do } t_2]_1]\right]_3 [\text{Sue PAST } [\text{do } t_3]_1]$

The problem here is that the variables t_2 and t_3 are distinct, and have to be, since each has its own binder. So semantic equivalence does not obtain and deletion should be blocked.

Sag does not have this problem, because he adopts essentially the following assumptions: (i) Determiners (such as *every*) are variable binders, and the arguments they select (restrictor and nuclear scope) are open sentences. (ii) Relative clauses are open sentences (i.e., relative pronouns are variables rather than variable-binders, and they reconstruct to their argument positions at LF). (iii) The head NP with which the relative clause combines is likewise an open sentence (it has an implicit subject position), and head NP and relative clause combine semantically by conjunction. So his LF for (21) was not like (33), but more like (22) below.[6]

---

[6]Actually, what Sag really assumed was that *every* is an unrestricted quantifier as in Predicate Logic, which combines with a single open sentence, formed by connecting restrictor and nuclear scope into a material conditional. The shortcomings of such an analysis are well-known by now. What I substituted for it here in the text avoids these problems, but is otherwise similar enough to preserve the assumptions that Sag really needed for the purposes of his ellipsis-theory. (See Heim (1982) for a more explicit rendition of the pertinent assumptions about LF-structure and semantic interpretation.)

(22)

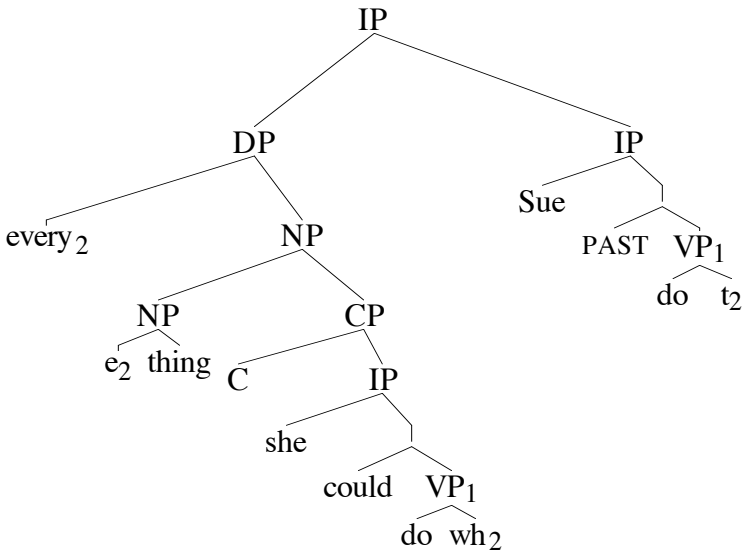

Notice that all three occurrences of the variable 2 are bound by the single occurrence of *every*$_2$.

The embedded VP is predicted to be deletable, as it should be.

Predictions about pronoun interpretation

So far, we have seen everything work well. The one phenomenon for which my current rendition of Sag's theory is blatantly inadequate is sloppy identity. In fact, it rules it out in a principled fashion. Let's see how.

To set the stage, we look at an uncontroversial case of a free (referential) pronoun.

(23) A: John$_1$ won't stay with Mary$_2$. He$_1$ doesn't get along with her$_2$ roommate.
B: Does Tom$_3$? (Or do we have to put him$_3$ up too?)

Imagine a conversational setting in which you would spontaneously take the *he* and *her* in the second clause to refer to John and Mary respectively (as indicated by the subscripts). It is easy, then, to understand the third, elliptical, clause as asking whether Tom gets along with Mary's roommate. This is the reading we predict -- in fact, the only one we predict for the last clause, given the indicated reference assignments in the text preceding. The pronoun *her*$_2$ in the LF of the overt VP *get along with her*$_2$ *roommate* is a free variable, which the contextually given assignment maps to Bill. To meet the semantic equivalence condition, the deleted VP must contain the same variable 2 in its LF, and hence the contextually given assignment will map it to the same referent, Mary. Even independently of the specific approach to referential pronouns (as

free variables) that I am assuming here, it is clear that a different referent for the *her* in the deleted VP would destroy semantic equivalence.

The prediction that referential pronouns in the antecedent VP always maintain their reference in the deleted VP appears to be generally correct. Consider, for instance, (24):

(24) John$_1$ won't stay with Mary$_2$. He$_1$ doesn't get along with her$_2$ roommate. Tom$_3$ won't mind staying with Susan$_4$, though. He$_3$ does.

It is very hard to read the last clause here as saying that Tom gets along with Susan's roommate. The fact that this would be pragmatically most plausible does not suffice.

Suppose now the antecedent VP for a deleted VP contains a bound variable pronoun. If the binder of that pronoun is also contained within this VP, as in (34), matters are straightforward.

(34) John made a boy$_1$ clean his$_1$ room, and Bill did too.

The antecedent and deleted VPs can have identical (and therefore equivalent) LFs *a boy$_1$ [make t$_1$ clean his$_1$ room]*. The text then means that John made some boy x clean x's room, and Bill made some (possibly, but not necessarily, different) boy y clean y's room. One might describe this as an instance of "sloppy identity" (on the grounds that different boys' rooms may be involved in the verifying scenario for this reading). But it is not the sort of case that people usually have in mind when they talk about "sloppy identity".

The standard cases of sloppy identity, in fact, are intractable for the present theory. Consider (35), on the reading which implies that John cleaned John's room and Bill cleaned Bill's room.

(35) John cleaned his room. Bill did too.

If the antecedent VP has an LF representation in which *his* is a free variable referring to John, then the deleted VP is licensed only on the reading according to which Bill cleaned John's room -- the "strict" reading, in other words. If, on the other hand, the *his* in the antecedent VP is a bound variable bound by *John*, as in (36),

(36) John$_2$ [t$_2$ PAST [clean his$_2$ room]$_1$]

then the deleted VP would have to contain another instance of this same variable, bound by *John* as well. This is not possible, since the scope of *John* is confined to its clause. (And even if *John* could be QRed to a position with scope over the whole text, the only reading we could obtain in this way would be the strict one.) The sloppy reading cannot be generated.

More generally, Sag's recoverability condition implies the following generalization: A bound variable pronoun may occur in a deleted VP only if either (a) its binder is also inside the deleted

VP, or (b) its binder has scope over both the deleted VP and its antecedent VP. Under our current background assumption, this amounts to the undesirable prediction that standard cases of sloppy identity (as in the most natural reading of (35)) are excluded.

3. Sloppy identity and the Derived VP rule

Sag's solution to this problem did not, as one might expect at this point, consist in a weakening of his semantic equivalence condition. Instead, he embraced the generalization about bound variables that we just stated and concluded that, contrary to appearances, a VP-internal pronoun that seemed bound by the subject was actually bound within the VP. (This had already been proposed by Keenan and elaborated further by Partee and others, so Sag did not take a view here that was particularly original or controversial at the time of his writing.)

How could the pronoun be bound within the VP, yet seem to be bound by the subject? Sag added to his rules for the derivation of LF from SS a version of Partee's so-called "Derived VP Rule":

(24) A VP α may be rewritten as 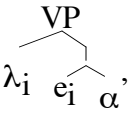

, for a freely chosen variable i.

e_i here is an empty category interpreted as a variable, which effectively saturates the external argument place of the verb. If only this 'VP-internal subject' were added, derived VPs would no longer be 1-place predicates, but open sentences. But this consequence is undone by the simultaneous insertion of a λ-abstractor which binds the VP-internal subject.

In fact, applications of the Derived VP Rule normally do not have any effect on either the semantic type or the specific meaning of the VP. Pairs like *[$_{VP}$ leave]* vs. *[$_{VP}$ λ_1 e_1 leave]* are semantically equivalent, of course. The only cases where the Derived VP Rule gives rise to new meanings are those in which additional occurrences of the variable i get bound along with the inserted subject e_i. An example is the following SS-to-LF derivation for the first clause in (35).

(37) SS: John PAST [clean his$_2$ room]$_1$
⇒ LF: John PAST [λ_2 e_2 clean his$_2$ room]$_1$

This derived VP denotes the property of cleaning one's own room. Hence when it is predicated of *John* in the containing clause, we obtain the proposition that John cleaned John's room. This creates the illusion of *his* being bound by *John*, when in fact it isn't.

By providing a VP-internal binder for pronouns apparently bound by the subject, we have removed the obstacle to generating sloppy identity readings for texts like (35). If the LF-derivation of the first conjunct in (35) is (37), and that of the second conjunct is (38):

(38) SS: Bill PAST [clean his$_3$ room]$_1$
⇒ LF: Bill PAST [λ_3 e$_3$ clean his$_3$ room]$_1$

then all conditions for deletion are met. (Notice that the indices on the antecedent and deleted occurrences of *his* are distinct, and they must be distinct to obey "No meaningless coindexing". But this doesn't hurt, because the complete λ-terms are equivalent, and it is those that constitute the relevant VPs at LF.)

We have seen that the Derived VP Rule is essential for the proper functioning of Sag's proposal. And it is not merely motivated in the context of his particular approach to VP ellipsis. It brings an additional, independent, benefit: It allows us to dispense, after all, with the special semantic mechanisms needed for in-situ interpretation of quantifiers. We now have a semantically sentential landing site for QR inside the VP and can exploit this, e.g., in the analysis of examples like (19) above. Still, the Derived VP Rule is unappealing, at least as it stands. We would like to reduce its effects to more elementary and generally accepted syntactic operations. Unfortunately, this does not seem to be entirely feasible. One part of it, namely the introduction of a VP-internal empty subject, looks quite unexceptional from the point of view of current syntactic theory. Many people nowadays assume that this position is projected to begin with, so it needn't be added by a special LF-rule. But the other part of the DVPR, the insertion of a λ-operator, seems no less ad hoc today than it did in the seventies.

II. Beyond Sag

1. A closer look at the distribution of sloppy identity readings

As we have seen, Sag's proposal predicts that sloppy identity interpretations are possible only for NPs that can be analyzed as bound variables with a binder inside the deleted VP. He offers quite a few examples in order to support the correctness of this generalization. The following observations about unavailable sloppy readings are all predicted by his analysis:

(1) *Norma$_1$ told Beth$_2$'s boyfriendto give her$_2$ a dime, and Judy$_3$ told Loi$_4$s's boyfriend to ~~give her$_3$ a dime~~.
(p. 123, attributed to Ross)

(2) *Sam$_1$ liked Janet's picture of Sam$_1$, and Bill$_2$ did ~~like Janet's picture of him$_2$~~ too.
(p. 129, attributed to Witten)

(3) *His$_1$ mother loves John$_1$, and Bill$_2$'s mother does ~~love him$_2$~~ too.
(p. 129)

(4) *John$_1$ scratched his$_1$ arm, and [the boy who knew Tom$_2$] did ~~scratch his$_2$ arm~~ too.
(129, Ross)

(5) *John$_1$'s sister scratched his$_1$ arm, and Bill$_2$'s sister did ~~scratch his$_2$ arm~~ too.
(129, Ross)

(6) *John$_1$ scratched his$_1$ arm, and Bill$_2$'s brother did ~~scratch his$_2$ arm~~ too.
(129, Ross)

(7) *Bill$_1$ and Mary are afraid that he$_1$ may be sent to India, and Tom$_2$ and Lucy are ~~afraid that he$_2$ may be sent to India~~ too.
(130, Witten)

(8) *John$_1$ said that Mary$_2$ hit him$_1$, and Bill$_3$ said she$_2$ did ~~hit him$_3$~~ too.
(131)

(9) *Mary$_1$ persuaded John$_2$ to say she$_1$ hit him$_2$, and she$_1$ persuaded Bill$_3$'s sister to ~~say she$_1$ hit him$_3$~~, also.
(132)

The following examples involve wh-traces (rather than sloppy pronouns), but the reason that Sag's theory excludes them is the same: The deleted VP contains a variable (in this case a wh-trace) whose binder (in this case *what*) is outside it.

(11) What did Harry take a picture of? ... *What did Bill? (p. 110)

(12) *What Sandy carried was the baseball bat, and what Betsy did was the catcher's mit. (p. 112)

(The following example is mentioned by Sag in the same context:

(10) We finally got in touch with John, who my brother Al tried to visit, but (*who he) couldn't. (p. 111)

It is unclear, however, that it illustrates the same point. As Sag himself notes elsewhere, if wh-pronouns introducing appositive relatives are referring expressions and can be reconstructed to their argument positions at LF, the asterisk in (10) is not actually predicted. This might be a good thing, after all, if (13), attributed by Sag to L. Horn, is okay.

(13) My brother Al, who John liked, but who Harry didn't, was a nice guy.

See Sag, ch. 1, fn. 26, p. 84.)

Elsewhere in the literature, however, we find quite a few different judgments on similar types of examples. Some of these predate Sag's writing, but there seems to have been an especially vigorous resurgence of published anti-Sag judgments since the late eighties. All of the following have been presented as counterexamples to Sag, and as examples which a better theory of VP-ellipsis ought to generate:

(14) Almost every boy$_1$ thinks that the teacher$_3$ likes him$_1$. [Only John]$_2$ thinks she$_3$ doesn't ~~like him$_2$~~. (Reinhart, 1986 LSA Institute presentation?)

(15) Tom$_1$ wanted Sue to water his$_1$ plants, while John$_2$ wanted Mary to ~~water his$_2$ plants~~. (Jabobson 1992, SALT 2)

(16) Bagels, I like. Donuts, I don't. (Jacobson, citing F. Evans)

(17) John$_1$'s coach thinks he$_1$ has a chance, and Bill$_2$'s coach does ~~think he$_2$ has a chance~~ too. (various places, including Rooth 1992)

(18) The policeman who arrested John$_1$ read him$_1$ his$_1$ rights, but the policeman who arrested Bill$_2$ didn't ~~read him$_2$ his$_2$ rights~~. (Wescoat 1989)

(19) If John$_1$ has trouble at school, I'll help him$_1$, but if Bill$_2$ does ~~have trouble at school~~, I won't ~~help him$_2$~~. (Hardt 1992)

2. Replacing VP-equivalence by a parallelism condition on larger domains

We saw that in at least some cases, Sag's recoverability condition was too restrictive. The following, for example, are grammatical, despite the fact that the requirement for semantically equivalent VPs cannot be met.

(1) Almost every boy$_1$ thinks that the teacher$_3$ likes him$_1$. [Only John]$_2$ thinks she$_3$ doesn't ~~like him$_2$~~.

(2) Tom$_1$ wanted Sue to water his$_1$ plants, while John$_2$ wanted Mary to ~~water his$_2$ plants~~.

(3) Bagels, I like. Donuts, I don't.

It is natural, therefore, to explore theories according to which deletion is licensed by a more permissive condition than semantic equivalence. There are many different plausible candidates for theories of this kind. The one we will explore here is presented (though not endorsed!) in Rooth (1992b).[7] Like other approaches (Chomsky, Tancredi), it incorporates a version of the intuitive idea that ellipsis is essentially an extreme form of deaccenting.

2.1 Background: Rooth's (1992a) analysis of contrastive focus

Judgments like those in (4) indicate that one of the semantic functions of prosodic prominence is the expression of contrast. (Capitalization represents main stress.)

(4) John saw Mary.
  (a) No, BILL saw Mary.
  (b) *No, Bill saw MARY.
  (c) No, John saw BILL.
  (d) *No, JOHN saw Bill.

Rooth (following Jackendoff) assumes that the relation between the prosody and its semantic effects is not direct, but mediated by a syntactic focus feature, which is subject to both phonological and semantic interpretation. As regards phonological interpretation, it is suffcient here to know that the stress pattern in (4a) is compatible only with the syntactic representation in (5a), and the stress pattern in (4b) at most with those in (5b), (5c), or (5d).[8]

(5) (a) Bill$_F$ saw Mary
  (b) Bill saw Mary$_F$
  (c) Bill [saw Mary]$_F$
  (d) [Bill saw Mary]$_F$

---

[7]The main point of Rooth's paper is to show that the condition I am about to formulate is not an adequate replacement for Sag's recoverability condition, at least not by itself. We will return to his arguments at a later point in the course, under the heading "VP-Ellipsis and Destressing".

[8]The proper formulation of the general principles which derive these predictions has been a matter of much debate. For a recent proposal, see Truckenbrodt 1995, ch. 2.

On the semantic side, F-marking determines the so-called focus (sermantic) value of a sentence. This is aset of propositions, intutively the set of evoked alternatives. Basically, the focus-value of ϕ is calculated by recalculating the ordinary semantic value of ϕ for each alternative to the value(s) of the F-marked constituent(s) in ϕ. For example, the ordinary and focus values of (5a,b) are as follows[9]:

(6) [[*Bill*$_F$ *saw Mary*]] = that Bill saw Mary
||*Bill*$_F$ *saw Mary*||
= {that Bill saw Mary, that John saw Mary, that Mary saw Mary, ...}
(= {that x saw Mary: x is an individual})[10]
[[*Bill saw Mary*$_F$]] = that Bill saw Mary
||*Bill saw Mary*$_F$||
= {that Bill saw Mary, that Bill saw John, that Bill saw Bill, ...}
(= {that Bill saw x: x is an individual})

See Rooth (1985) for a precise account of how such focus values are calculated compositionally.

Besides F-marking, the LFs of the second sentences in (4) contain a covert focus-interpretation operator ~ which is coindexed with the intended target of contrast, viz. the initial clause of (4). So we have LF representations like (7).

(7) [John saw Mary]$_1$ No, [Bill$_F$ saw Mary] ~$_1$

The operator ~$_1$ contributes an implicature to the semantics of (7), namely that [[*John saw Mary*]] ∈ ||*Bill*$_F$ *saw Mary*||. Since this is true (see (6) above), (7) is felicitous. By contrast, an analogous LF-representation (8) carries the manifestly false implicature [[*John saw Mary*]] ∈ ||*John*$_F$ *saw Bill*||.

(8) * [John saw Mary]$_1$ No, [John$_F$ saw Bill] ~$_1$

Thus Rooth accounts for the contrast between (4a) and (4d).
The general rule for the interpretation of ~ may be stated (for our purposes here) as in (9).

---

[9]I use double vertical lines for focus semantic values, and the usual double brackets for ordinary semantic values. (Rooth uses [[ ]]^f and [[ ]]^o.)

[10]Here we use an extension version of the standard set-abstraction notation, which is defined as follows: $\{\alpha : \phi\} := \{v : \exists v_1, ..., v_n [v = \alpha \,\&\, \phi]\}$, where $v_1, ..., v_n$ are all the variables that occur free in α, and v is a variable that does not occur free in α or ϕ.

(9) An LF containing constituents ϕ_i and $\psi \sim_i$ is felicitous only if $[[\phi]] \in \|\psi\|$. (More precisely: only if $[[\phi]]^g \in \|\psi\|^g$ for every variable assignment g.)

When the condition $[[\phi]] \in \|\psi\|$ is met, we also say that ψ contrasts appropriately with ϕ. Bear in mind below that this relation is not a symmetrical one: When ψ contrasts appropriately with ϕ, it may or may not be the case that ϕ also contrasts appropriately with ψ. The former is a matter of where focus is placed in ψ, the latter depends on where focus is placed in ϕ.

2.2 The Appropriate Contrast condition on VP ellipsis

In a nutshell, the revision of Sag's theory that we want to explore consists in dropping the requirement that the deleted VP be equivalent to its antecedent, and imposing instead a requirement that some constituent containing the deleted VP must contrast appropriately with some constituent containing its antecedent. Technically, this means we drop (10) (= (17a) on p. 10).

(10) A VP that is coindexed with another VP may delete.

In its place we adopt (11).

(11) A VP α may delete in the configuration

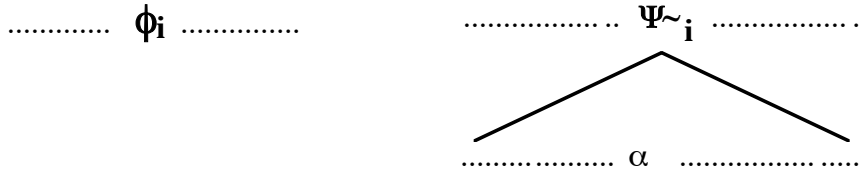

Let's call this the Appropriate Contrast Condition on VP deletion.

Here are two very simple illustrations of how (11) applies.

(12) You don't drink. You should.

(13) Mary drinks. And John does.

The natural pronunciation of these examples is consistent with positing F-subscripts on *should* in (12) and on *John* in (13). So we can have the following derivations (possibly among others).

(12') SS: $[\text{you}_1 \text{ don't } [_{VP}\text{t}_1 \text{ drink}]]_8 \; [\text{you}_2 \text{ should}_F [_{VP}\text{t}_2 \text{ drink}]]_{\sim 8}$

PF: you don't drink you SHOULD ∅

LF: $[\text{ don't } [_{VP}\text{you drink}]]_8 \; [\text{ should}_F [_{VP}\text{you drink}]]_{\sim 8}$

felicity condition: that you don't drink ∈

{f(that you drink): f is a function from propositions to propositions}

(13') SS: $[\text{Mary}_1$ PRES $[_{VP}t_1$ drink$]]_8$ and $[[\text{John}_F]_2$ does $[_{VP}t_2$ drink$]]_{\sim 8}$
PF: Mary drinks and JOHN does ∅
LF: $[$ PRES $[_{VP}$Mary drink$]]_8$ and $[$ does $[_{VP}$ John$_F$ drink$]]_{\sim 8}$
felicity condition: that Mary drinks ∈ {that x drinks: x is an individual}

The respective felicity conditions are met, and so these are good derivations.

3. The relation between Appropriate Contrast and Sag's Recoverability

The Appropriate Contrast condition was announced as a "more permissive" substitute for Sag's equivalence requirement. This ought to mean (a) that some cases predicted ungrammatical by Sag are predicted grammatical under Appropriate Contrast, and (b) that all cases predicted okay by Sag are still okay under Appropriate Contrast.

It is easy to substantiate (a). In fact, examples (1) - (3), which motivated the abandonment of Sag's proposal, are all cases in point. Let's see why they are now predicted grammatical.

(2) Tom$_1$ wanted Sue to water his$_1$ plants, while John$_2$ wanted Mary to ~~water his$_2$ plants~~.

Can we find a constituent containing the deleted VP that will contrast appropriately with something else in this text? Yes, we can, provided that we make two assumptions: First, that *John* and *Mary* are focussed; and second, that the two instances of *his* are bound variables, bound by *Tom* and *John* respectively. On these two assumptions, we can choose the following derivation. (I leave out the PF, and instead underline the VPs slated for deletion in the SS.)

(14) SS: $[\text{Tom}_1$ PAST t_1 want Sue$_3$ to t_3 water his$_1$ plants$]_8$,
while $[[\text{John}_F]_2$ PAST t_2 want $[\text{Mary}_F]_4$ to $\underline{t_4 \text{ water his}_2 \text{ plants}}]_{\sim 8}$
LF: $[\text{Tom}_1$ PAST t_1 want to Sue water his$_1$ plants$]_8$,
while $[[\text{John}_F]_2$ PAST t_2 want to Mary$_F$ water his$_2$ plants$]_{\sim 8}$
felicity condition: that Tom wanted Sue to water Tom's plants
∈ {that x wanted y to water x's plants: x and y are individuals}

How about (3)?

(3) Bagels, I like. Donuts, I don't.

Here we can give a licensing derivation if we posit foci on *donuts* and the negation:

(15) SS: $[\text{bagels}_1\ \text{I}_2\ \text{PRES}\ t_2\ \text{like}\ t_1]_8\ [[\text{donuts}_F]_3\ \text{I}_4\ \text{don't}_F\ \underline{t_4\ \text{like}\ t_3}]_{\sim 8}$

LF: $[\ \text{PRES I like bagels}]_8\ [\ \text{don't}_F\ \text{I like donuts}_F]_{\sim 8}$

fel. cond.: that I like bagels $\in \{f(\text{that I like x}): x \in D_e \text{ and } f \in D_{<st,st>}\}$[11]

Another example that the present analysis licenses is Hirschbühler's counterexample to Sag:

(16) An American flag is in front of every building, and a Canadian flag is too.

Hirschbühler pointed out that (16) was fine on a reading where the universal object in each clause scopes over the existential subject. We can derive this reading now if we posit focus on *Canadian* .

(17) SS: [[an American flag]$_1$ is t_1 in front of every building]$_8$,
and [[a Canadian$_F$ flag]$_2$ is $\underline{t_2\ \text{in front of every building}}$]$_{\sim 8}$ too

LF: [is [every building$_3$ [an American flag in front of t_3]]]$_8$,
and [is [every building$_4$ [a Canadian$_F$ flag in front of t_4]]]$_{\sim 8}$ too

fel. cond.: that for every building there is an American flag in front of it
$\in \{$that for every bldg there is an X-ian flag in front of it: X a nationality$\}$

It should be noted that in all these examples, the focus placement we have posited appears to be consistent with the phonetic evidence.

Having seen that Appropriate Contrast indeed licenses cases that Sag ruled out, can we show the converse, i.e., that Appropriate Contrast does not rule out anything that Sag allowed? This is not completely obvious. It depends on assumptions that we have not yet made explicit.

A question that arises in this connection is whether the containing constituent ψ in (11) can be the deleted VP α itself. For instance, could we also have licensed the deletion of *drink* in (12) as follows?

(12") SS: $\text{you}_1\ \text{don't}\ [_{VP}t_1\ \text{drink}]_8\ \ \text{you}_2\ \text{should}_F\ [_{VP}\underline{t_2\ \text{drink}}]_{\sim 8}$

LF: don't [you drink]$_8$ should$_F$ [you drink]$_{\sim 8}$

Notice that (12") has no F-marked constituent in the domian of ~. What is the focus-value of a phrase with no foci in it? It turns out to be the singleton set containing just its ordinary value: ||*you drink*|| = {that you drink}. This being so, the felicity condition associated with the LF in

---

[11] "$x \in D_e$ and $f \in D_{<st,st>}$" abbreviates "x is an individual and f is a function from propositions to propositions".

(12'') is 'that you drink ∈ {that you drink}', equivalent to "that you drink = that you drink', a truism.

If this sort of representation is allowed, we can give a straightforward proof that every case in which the deleted VP has a semantically equivalent antecedent VP is ipso facto a case which obeys the Appropriate Contrast Condition. All we need to do in each such case is choose the deleted VP itself as ψ and its antecedent as ϕ. Since it is an immediate consequence of the construction of focus values that $[[\alpha]] \in \|\alpha\|$, it follows that semantic equivalence of ψ and ϕ (i.e. $[[\phi]] = [[\psi]]$) will automatically imply $[[\phi]] \in \|\psi\|$.

4. Predictions about the distribution of sloppy identity readings

-- examples ruled out by lack of "parallelism"

(6) *John$_1$ scratched his$_1$ arm, and Bill$_2$'s brother did ~~scratch his$_2$ arm~~ too.

(20) *Tom$_1$ wanted John to water his$_1$ plants, and Fred wanted Bob$_2$ to ~~water his$_2$ plants~~.

-- examples ruled out by constraints on bound variable anaphora ("Weak Crossover")

(4) *John$_1$ scratched his$_1$ arm, and [the boy who knew Tom$_2$] did ~~scratch his$_2$ arm~~ too.

prediction: If sloppy identity is okay, then bound variable anaphora with a quantifying antecedent is also okay in the same configuration.

Tomioka (1995) refines this generalization: If sloppy identity is okay, then bound variable anaphora with a quantifying antecedent <u>or E-Type anaphora</u> is also okay in the same configuration.

5. Implications for LF-syntax

The Derived VP Rule becomes superfluous.

(21) John cleaned his room, and Bill did too.

SS = LF: $[$John$_1$ PAST $[_{VP}$t$_1$ clean his$_1$ room$]]_8$
and $[[$Bill$_F]_2$ did $[_{VP}$t$_2$ clean his$_2$ room$]]_{\sim 8}$ (too)
fel. cond.: that John cleaned John's room ∈ {that x cleaned x's room: x ∈ D}

The standard (Quine-Montague) analysis of relative clauses and QR-structures becomes tenable, provided that

either QR-derived λ-abstracts can qualify as constituents for the purpose of the Appropriate Contrast Condition,

or quantifier-meanings qualify as appropriate alternatives to variable-meanings, and the Appropriate Contrast Condition allows overlap at LF.

(22) every [thing [wh$_1$ could$_F$ she do t$_1$]~$_8$] [λ$_2$[did Sue do t$_2$]]$_8$

(23) [every [thing wh$_1$ [could$_F$ she do t$_1$]~$_8$] λ$_2$[did Sue do t$_2$]]$_8$

Main references

Sag, Ivan (1976) Deletion and Logical Form, MIT Ph.D. thesis

Rooth, Mats (1992) (1992b) "Ellipsis Redundancy and Reduction Redundancy," in Steve Berman and Arild Hestvik (eds.) Proceedings of the Stuttgart Ellipsis Workshop, report series Sprachtheoretische Grundlagen für die Computerlinguistik, no. 29, SFB 340, IBM Heidelberg.