

- Sybesma, Rint. 1999. *The Mandarin VP*. Dordrecht: Kluwer.
- Takano, Yuji. 1998. Object shift and scrambling. *Natural Language and Linguistic Theory* 17:817–889.
- Toivonen, Ida. 2003. *Non-projecting words: A case study of Swedish particles*. Dordrecht: Kluwer.
- Trinh, Tue. 2009. A constraint on copy deletion. *Theoretical Linguistics* 35:183–227.
- Trinh, Tue. 2010. Edges and linearization: A reply. *Theoretical Linguistics* 36:93–110.
- Wexler, Kenneth, and Peter W. Culicover. 1980. *Formal principles of language acquisition*. Cambridge, MA: MIT Press.
- Williams, Alexander. 2015. *Arguments in syntax and semantics*. Cambridge: Cambridge University Press.
- Williams, Alexander. To appear. Objects in resultatives. *Natural Language and Linguistic Theory*.
- Williams, Edwin. 1980. Predication. *Linguistic Inquiry* 11:203–238.
- Zeller, Jochen. 2001. *Particle verbs and local domains*. Amsterdam: John Benjamins.

Department of Linguistics and Cognitive Science
University of Delaware
Newark, DE 19716
bruening@udel.edu

Some People Think There Is Neg Raising, and Some Don't: Neg Raising Meets Ellipsis

Pauline Jacobson

The interaction of Neg raising (NR) with VP-ellipsis (VPE) shows that if NR is a rule of grammar, then the conditions on VPE must be exact syntactic identity and must be insensitive to major semantic differences between the so-called antecedent and the meaning understood at the ellipsis site. In particular, the conditions on ellipsis must be so blind to the semantics that they allow a polarity reversal between the antecedent and the understanding at the ellipsis site. But the behavior of indexicals shows quite clearly that meaning is what counts for the understanding of VPE, not form. This in turn provides new evidence against a syntactic process of NR.

Keywords: Neg raising, VP-ellipsis, indexicals, negative polarity items, Horn clauses

This remark makes two (related) points. Section 1 shows that if Neg raising is a rule of grammar (in any of its versions), then the conditions on VP-ellipsis must be impervious to any kind of identity of meaning between the so-called antecedent and the material that is silenced—or understood—at the ellipsis site. (The remarks extend to focus-based accounts.) Section 2 argues that

An earlier version of this remark was presented at the November 2015 Hornucopia Workshop Honoring Laurence Horn at Yale University. I am grateful to the audience there for useful comments, as well as the Brown LingLangLunch group. For additional discussion and/or very helpful comments on an earlier version, I would like to thank two anonymous *Linguistic Inquiry* reviewers, Larry Horn, Paul Postal, and Scott AnderBois. I am especially grateful to Chris Collins, whose comments have greatly helped sharpen this remark.

it is highly unlikely that the conditions on ellipsis so completely disregard meaning. Hence, if it is true that VP-ellipsis does not ignore the relationship between the *meaning* of the silenced/understood material and the antecedent material that licenses/supplies this meaning, then the interactions discussed in section 1.3 provide new evidence against Neg raising.¹

1 Why Neg Raising Requires (Only) Formal (Syntactic) Identity in Ellipsis

1.1 Background: Neg Raising

The results here concerning the connection between ellipsis and Neg raising (hereafter, NR) are rather theory-neutral, and so I phrase the remarks throughout in fairly broad terms, abstracting away as much as possible from the details of any particular theory. Hence, by *NR hypothesis* I mean only the hypothesis that a sentence like (1) has a representation along the lines of (2a) or (2b) and that this is the one relevant for the determination of meaning (the details of where the negation is located in the lower clause vary).

- (1) Joe does not believe that Donald will be president.
- (2) a. Joe believes [NEG [Donald be president]]
 b. Joe believes Donald [NEG [be president]]

This result could, for example, be accomplished by an account that includes a rule raising NEG in the syntax and in which the semantics interprets a pre-NR level (as was the original idea). Or it could be that NEG leaves a trace that is interpreted and the higher NEG is somehow invisible to the semantics. Or there could be a rule lowering NEG at LF.

The idea that (1) has one representation in which the semantics actually “sees” (2) was first proposed by Fillmore (1963) and has been argued for by many authors since. For some classic

¹ An anonymous reviewer points out that Lindholm (1969) explored the interaction of Neg raising (NR) and “pronominalization” (often noted to be closely related to ellipsis) and came to the opposite conclusion from that reached here. In particular, Lindholm was concerned with sentences like (i), where *it* is understood as the proposition ‘Donald didn’t pay his taxes’.

- (i) I don’t think that Donald paid his taxes, and Curt is quite sure of it.

Assuming (as Lindholm did) that *it* is derived from a full sentence (formally) identical to the antecedent, the conclusion here is that there must be a syntactic rule of NR that applied in the antecedent clause. However, it is well-known (since at least as early as Hankamer and Sag 1976) that *it* is a “deep anaphor” and picks up contextually salient individuals, properties, or propositions, and so there is no reason to assume that it has “replaced” a full sentence.

The same reviewer notes that a similar point holds for sluicing. Here, I will argue that if there is NR, then VP-ellipsis must be sensitive to formal and insensitive to semantic identity (and I will argue that since insensitivity to semantic identity is implausible, NR cannot exist). But data from sluicing can be used to show the reverse: if formal identity is required in sluicing, then there must be a rule of NR. The key datum (pointed out by the reviewer, but see also Kroll 2016) is (ii), where the understood material in the sluice is *he is not going to release his tax returns*.

- (ii) I don’t think he is going to release his tax returns, and I’m pretty sure I know why.

Again, though, there is independent evidence that sluicing cannot just require exact formal identity. Consider, for instance, the following striking example from AnderBois 2014:902:

- (iii) (Either) Freddie is baking a cake again or something is on fire, but I can’t tell which (one).

A full discussion of sluicing is beyond the scope of this remark, but see Kroll 2016 for relevant discussion not making use of NR.

arguments to this effect, see for example Lakoff 1969 and Horn 1971; a recent extensive defense of this hypothesis is put forth in Collins and Postal 2014a. Of course, it is well-known that (1) also has a representation in which the negation is in the upper clause for the purposes of the semantics (since nothing could block this), and the meaning of this representation is quite close to the NR reading in (1). In view of this, many have argued that NR is an illusion: (1) has only a literal meaning in which *not* is in the upper clause and the NR “reading” is due to pragmatic strengthening. A particularly influential proposal to this effect is given by Bartsch (1973), who points out that if we simply assume that for any proposition *p*, Joe either believes *p* or believes not *p*, then it follows that if Joe does not believe *p*, he also believes not *p*. Arguments back and forth about whether NR is an actual rule have gone on for decades; in addition to Bartsch 1973, see Horn 1978, Horn and Bayer 1984, Gajewski 2007, and many others for proposals that derive the NR reading without a literal NR rule (in any of its guises).

A review of the arguments on both sides is beyond the scope of this remark, but I mention two that have been taken to motivate NR, as these will be useful below. The first centers on the distribution of certain strong negative polarity items (NPIs) such as *until* (Lakoff 1969) and *in ages*. The argument assumes that *until* (on the relevant reading) must be within the local scope of negation; note the contrast between (3a) and (3b).

- (3) a. The weather forecaster said that it won't start snowing until midnight.
- b. The weather forecaster didn't say that it will start snowing until midnight.

There is an irrelevant reading in which the saying did not happen until midnight. But the reading of interest here is the one in which *until midnight* is linked to the time of the onset of the snow. This is available for (3a) but not (3b).

Contrast this to corresponding cases with NR verbs such as *think*.

- (4) a. The weather forecaster thinks that it won't start snowing until midnight.
- b. The weather forecaster doesn't think that it will start snowing until midnight.

(4b) has a natural reading in which *until midnight* is associated with the onset of the snowing. If it is correct that punctual *until* must be within the local scope of a negation, the goodness of (4b) follows under NR. But the issue remains unsettled because the assumption that punctual *until* requires a local negation has been challenged: Horn (1978), Gajewski (2011), and others argue that its distribution can be accounted for in part by pragmatic inferences and/or presuppositions. Nonetheless, defenders of NR have generally taken the distribution of punctual *until* to be subject to the local negation restriction; and given that assumption, its presence is useful as a diagnostic.

Similar remarks hold for the distribution of *in ages/in years*, and so on. The parallel facts are given in (5) and (6), and the reasoning and conclusions are the same.

- (5) a. Sarah said that the lawn hadn't been mowed in ages.
- b. *Sarah didn't say that the lawn had been mowed in ages.
- (6) a. Sarah thinks that the lawn hasn't been mowed in ages.
- b. Sarah doesn't think that the lawn has been mowed in ages.

A second argument, heavily featured in Collins and Postal 2014a, centers on what the authors dub *Horn clauses* (the key facts were originally presented in Horn 1975). Assume that the subject-aux inversion found in (7a) is licensed only when the preposed adverbial contains a negative. The full paradigm in (7) is designed to show that a negative is indeed required and that, moreover, the inversion is obligatory in such a case.

- (7) a. At no point in my life will I acquire a taste for cilantro.
 b. *At no point in my life I will acquire a taste for cilantro.
 c. *At some point in my life will I finally acquire a taste for cilantro.
 d. At some point in my life I will finally acquire a taste for cilantro.

But (8a) is also good, and (8b) shows that here the inversion is obligatory.

- (8) a. I don't think that ever again will I see so many beautiful oaks in October.
 b. *I don't think that ever again I will see so many beautiful oaks in October.

I will not give Collins and Postal's full analysis here, but the basic point should be clear: they assume that in (8a–b) NEG + *ever* is in the lower clause at the level of representation relevant both for the semantics and for licensing the inversion, and NEG raises into the *think* clause.

Again this is not the end of the story. Horn (2014) points out that similar facts obtain with the non-NR verb *know* (though see also Collins and Postal 2014b). Note too that certain affective predicates that license NPIs like *ever* also allow the inversion.²

- (9) I'll be surprised if ever (again) will/do I see such beautiful oak leaves.

This suggests that perhaps it is the presence of the NPI itself that licenses the inversion (although one might also try to decompose *surprise* into *expect not*). So once again the issue remains open, and it is not my purpose here to try to resolve this. The reason for bringing Horn clauses up is that once again they provide a useful diagnostic: if NR exists, then the Horn clauses provide a diagnostic for its presence.

1.2 Background: VP-Ellipsis

So many different accounts of VP-ellipsis (hereafter, VPE) have been proposed that there is no way to catalogue them all here, and so I paint the possibilities in broad strokes. The first major cut centers on the question of whether there is any linguistic material in the position of the ellipsis site at any level, or whether instead an auxiliary (or the ellipsis site itself) is a kind of "free" pro-form such that the grammar itself provides neither any material nor any meaning there. Under the latter view (detailed a bit more below), the listener simply supplies a contextually salient property. Under the former view, that there is indeed linguistic material (at some level) in the ellipsis site, it is always assumed that this material is deleted or silenced in virtue of some kind

² One might speculate that the inversion is actually just triggered by the presence of *if* since this does, after all, license NPIs. But *if* in its normal use does not trigger the inversion.

- (i) a. She'll be elected if ever again she gives such a good speech.
 b. *?She'll be elected if ever again does she give such a good speech.

of identity with some other overt VP in the discourse context. (The rubric of “silent linguistic material” is intended here to extend to views such as Williams’s (1977) in which material is copied in.) After all, if there were no identity condition, this would be a completely nonrecoverable deletion/silencing, and the ellipsis site could be interpreted as anything at all, which is obviously not the case. The main line of disagreement for the silent-material view, then, is whether the identity condition cares about identity of form or identity of meaning (see, e.g., Ross 1967 for the first and Merchant 2001 for the second (for the case of sluicing)). There are also many versions of the identity condition (whether it be formal or semantic) that allow for a certain amount of mismatch; I do return to this condition, though I gloss over some of these details whenever possible to keep the discussion manageable.

Thus, I initially consider three main positions:

(a) *Formal identity*. A VP may be silenced/deleted under exactly formal identity (at some level) with another overt VP in the discourse context. Again, this proposal has many variants. For example, the requisite condition could be formal identity on the surface or some similar level, or identity at LF.³ The latter is sufficiently similar to an “identity of meaning” condition (in terms of the predictions it makes) that, from now on, by the term *formal identity* I do not mean LF identity; rather, I mean formal identity at some other level of representation that—quite crucially—is not “seen” by the semantics.

(b) *Semantic identity*. A VP may be silenced/deleted provided its meaning is such that there is another, overt VP in the discourse context with the same meaning.

(c) *Deep anaphora*. There is no material (beyond possibly a silent pro-form) in the position of the ellipsis site. Either the ellipsis site is a pro-form (Hardt 1993) or the auxiliary is, or the auxiliary contains something akin to a pro-form as part of its meaning (see, e.g., Jacobson 2003). This means that the “missing” material in VPE is recovered much as in the case of free pronouns (under many accounts of free pronouns). The grammar provides neither any linguistic material nor any actual meaning in the ellipsis site; rather, some contextually salient property is supplied by the listener. Of course, Hankamer and Sag (1976) famously argued precisely against this position on the basis of the fact that one cannot infer the relevant property from context alone; they claimed that it had to actually be the meaning of an overt VP. But since then, numerous cases with no overt linguistic antecedent have been observed (for recent discussion with corpus data, see Miller and Pullum 2013). That said, it is certainly the case that pure pragmatic control is difficult and is not the common situation, and so the deep anaphora theory needs to say why this is so (and why pragmatic control is so much harder in this instance than it is in the case of ordinary pronouns). I will not discuss this in detail here; in Jacobson 2003, I speculate that “properties” (functions of type $\langle e, t \rangle$) are simply harder to access than individuals and thus are likely to be contextually salient by having been named. Thus, it is an illusion that there is an actual “antecedent”; what we tend to think of as the antecedent for the understood material at the ellipsis site is merely the linguistic expression that names the relevant property and hence makes it contextually salient. For convenience, most of the discussion can collapse (b) (semantic

³ This was basically Sag’s (1976) proposal, although he allowed for mismatch in the indices “up to alphabetic variance.”

identity) and (c) (deep anaphora, in which what we think of as the “antecedent” is the material that names the supplied property). Both have the consequence that what we call the “antecedent” must have a meaning identical to the meaning that is understood at the ellipsis site.

There are additional possibilities that do not fall under any of these positions. One has both a formal and a semantic identity condition. In a similar vein, many researchers have adopted Rooth’s (1992) theory or some variant thereof, which has both a formal identity and a focus condition (stated on a slightly larger domain). And there are accounts relying only on the focus condition (e.g., Fox 1999). I will initially set the focus condition aside and return to it in section 1.3; it too will be incompatible with NR. There are also claims to the effect that the grammar contains some kind of “hybrid” of a formal and a semantic condition. However, this is viable only if there is a general way to determine when identity must be semantic and when it must be syntactic. Lacking full detail, I will say little about this here (but see section 3).

1.3 *Neg Raising: The Consequences for VPE*

At this point we can note the key facts, which are presented in (10) and (11). (In Jacobson 2006, I note facts parallel to (11) for the narrower case of the *can’t seem to* construction; I use the ellipsis behavior to argue for a compositional analysis of that construction.)

(10) Chris thought that Donald would win, but Joe didn’t.

(11) Joe didn’t think that Donald would win, but Chris did.

Notice that (11) can be contrasted with a parallel case in which the negation in the first sentence is on the surface in the lower clause; see (12). While the judgment here is not completely clear (I return to this directly), the contrast is nonetheless quite strong (the ?* is a rough approximation of my own judgment; speakers vary).

(12) ?*Joe thought that Donald wouldn’t win, but Chris did.

One can, in fact, construct cases along the lines of (12) that are (for some speakers) surprisingly good, such as (13a).

(13) a. ??I think that a Socialist can’t possibly win in 2016. But Denise does.

b. I don’t think that a Socialist could possibly win in 2016. But Denise does.

Nonetheless, there is a clear difference between (11) and (12); and even for most of the speakers who find (13a) surprisingly “okay,” it contrasts strongly with (13b).⁴ I will put cases like (12)

⁴ During the course of three talks on this material (as well as comments on a handout), I can count about ten speakers who voiced firm opinions on (12) and (13a) (some of these speakers were naive informants who were not versed in syntactic theory and had never heard of NR). Four speakers found both sentences awful, although to be fair three were less than naive informants (but one was a completely naive informant). Of the remainder (some naive, some not), almost all found that (12) contrasted robustly with (11), where the negation in the first conjunct is in the upper clause (generally, speakers just didn’t like (12)). As to (13a), a common reaction was that speakers knew what its intent was, and some (like myself) found it not perfect but surprisingly good—but again, it contrasted with (11). Only one person (not a naive informant) reported finding (12) perfect. One anonymous reviewer also brought up cases like this, for a different reason. This reviewer also found that one can construct surprisingly good cases along these lines, but reported that not all such cases are good; see the discussion below.

and (13a) aside for the moment; and in the ensuing discussion, I will not construct examples along the lines of (12) (with the negation in the lower clause in the first conjunct) each time I introduce an example like (11) (with the negation in the upper clause). The interested reader can easily construct the minimal pairs for all the relevant examples and will probably find that the basic pattern remains: sometimes the negation in the lower clause is a bit better than one might expect, but there is a robust contrast with negation in the upper clause.

To simplify the exposition, I first develop the basic point using these sentences rather than ones containing strong NPIs and/or Horn clauses. This means that initially the point is not airtight, for it is uncontroversial that, for example, the first clause of (11) has a non-NR reading/representation under which the negation is always in the upper clause. But for the sake of exposition, assume for now that the presence of the “strong” reading is a diagnostic for NR (on theories in which that reading is due to NR and not to pragmatic strengthening). By the strong reading, I mean one where the first clause of (11) is understood as making a positive statement about Joe’s opinion: he thinks Donald will lose. After this initial discussion, I will strengthen the point by relying not on the strong reading but on the other diagnostics for NR (according to theories maintaining NR).

Consider first (11). What is understood after *do* in the second conjunct is the property [[think Donald will win]], not [[think Donald won’t win]]. But if the first conjunct (on the easy-to-get and relevant strong reading) involves NR, then the meaning of the matrix VP (i.e., the “antecedent”) is [[think that Donald won’t win]], and so this is what should be the silent or supplied material in the second conjunct. Incidentally, one cannot appeal to the fact that if the interpretation of the second conjunct were *Chris thinks Donald won’t win*, that conjunct should be followed by *too* or *also*. This is undoubtedly true, but inserting *too* in (11) just makes this quite strange.

(14) ?*Joe doesn’t think that Donald will win, and Chris does too.

In fact, the badness of (14) is not predicted if NR exists and if ellipsis has any sensitivity to meaning identity. As pointed out to me by Chris Collins, we need not even depend here on the assumption (made for temporary expository convenience) that the strong reading in (11) involves NR. For the very claim that there *can* be an NR derivation for, say, the first clause in (14)—combined with the assumption that the elided material will have the same meaning as the antecedent VP—would lead us to expect a reading for the *Chris* clause in which Chris thinks that Donald won’t win (as this is one possible meaning for the matrix VP in the *Joe* clause). Hence, (14) is predicted to be good: Joe and Chris think the same thing. This prediction is clearly incorrect.

Furthermore, consider (10). There is an understanding of the second conjunct in (10) in which the strong reading emerges. (This is brought out even more clearly by changing the second conjunct to *Joe most definitely doesn’t*.) If that reading comes from a grammatical NR process (rather than via pragmatic strengthening), then, first of all, there must be actual linguistic material in the ellipsis site out of which NEG has raised. Moreover, if what is needed is semantic identity between the antecedent and the silent material, then the silent material would be [[think that Donald will win]], not [[think that Donald won’t win]]. But in that case, the negation is always in the upper clause—and the strong reading cannot be due to NR. In other words, if NR is responsible for the strong reading in the second conjunct in (10) and the first in (11), then the “antecedent” and

the ellipsis site do not have identical meanings. Moreover, we are not talking here about a minor difference in meaning. The difference is quite radical: the polarity is reversed! And it is not even the polarity of the highest VP; these differences reverse the polarity of an embedded VP.

But a formal identity condition (supplemented with the assumption that a raised NEG does not leave a trace) is perfectly compatible with the NR hypothesis. In (11), the silenced material is *think that Donald will win*—at least at some level of representation. (Under the NR analysis, this is not the representation that the semantics “sees”; but we are assuming that formal identity checks identity at some surface or other non-input-to-semantics level.) In terms of the interpretation of (11) (under NR), the meaning of the relevant VP in the first conjunct is \llbracket thinks that Donald will not win \rrbracket , but that is irrelevant. Thus, the silenced/deleted material in the second conjunct is *think that Donald will win*; this is interpreted, and we get the right reading. (10) is similarly unproblematic. The silenced or deleted material in the second conjunct is *think that Donald will win*, but this means that the full second conjunct is as in (15) (with the overstrikes, as usual, representing the silenced/deleted material).

(15) . . . Joe doesn't ~~think that Donald will win~~

This meets the formal identity condition. But the semantics does not “see” this representation; rather, it sees one in which NEG is in the lower clause.

So far, the argument that the available readings require NR is not airtight; it assumes that the strong reading only emerges with NR. But this reading is not necessarily the best diagnostic for NR. After all, it is perfectly uncontroversial that there is one representation for, say, (11) in which the first conjunct is such that NEG is always in the upper clause, so that the meaning of the complement of *don't* is indeed \llbracket think that Donald will win \rrbracket . And so semantic identity does allow for the right reading (provided that the strong reading in the first conjunct is accounted for in some other way). I now tighten this argument by turning to those tests that are generally agreed to be diagnostics for NR—if it exists—and adding them into the mix. (Note too that the flip side of the argument—that NR predicts readings that are not available—was shown simply by (14).)

First, consider (16) in the context shown (I switch to the NR verb *want*).

(16) (Context: Thelma and Louise are the only two people in the department who can possibly be the Linguistics Graduate advisor next year. Both are eligible for semester-long sabbaticals, so there could be a problem. But there isn't, because . . .)

Thelma wants to take her sabbatical in the fall, but (fortunately) Louise doesn't until next spring.

This sentence may be a bit awkward, but only because it might be a case of so-called pseudogapping, which (except in comparatives) generally has an awkward feel. Note that it is no more awkward than similar cases with no NPI such as (17).

(17) Thelma will take her sabbatical in the fall, and Louise will in the spring.

Of course, one might object here that we have suddenly changed the scope of the discussion from VPE to pseudogapping. Given the need for the addition of *until* in order to make the argument

here, there is no way to avoid that. But—whether the two constructions are subsumed under the same mechanism (as some have argued) or involve different mechanisms—it is reasonable to hope that both are subject to the same basic identity conditions, so this move does not seem problematic. Besides, even if pseudogapping is quite separate, the entire discussion in this article could have been phrased in the context of pseudogapping instead of VPE. (17) is sufficient to show that the conditions on pseudogapping must be formal identity if NR exists, but the arguments in section 2 to the effect that this kind of formal identity condition is not plausible hold just as well for pseudogapping cases. I leave this to the interested reader to verify.⁵

Similar facts hold for *in years* (and here things conspire in such a way that we can construct cases that uncontroversially involve VPE and not pseudogapping).

- (18) ?Sally doesn't think that it has been this cold in April in ten years, but I do—in fact, I'm pretty sure that it was even colder than this just last April.

Again, while not perfect, this is not crashingly bad. But if the first clause must involve NR in order to license *in ten years*, then the semantic identity (b) and deep anaphora (c) theories of VPE would have to supply [[has not been this cold in ten years]] as the meaning in the ellipsis site, which is clearly not the intended meaning.

One can make the same point with Horn clauses.

- (19) Linda didn't used to think that ever in a million years would an African American be elected president. ?But Michele always did.

The interpretation of the second clause with ellipsis might strip out the meaning of 'ever in a million years', leaving just the meaning [[an African American would be elected president]]. Since [[ever in a million years]] itself is an NPI, this is not surprising (it would presumably not make sense to have this phrase as part of the meaning, although it is well-known that some NPIs—especially weak ones like *any*—are fine in the antecedent).⁶

⁵ Chris Collins points out to me that there are other infinitives with non-NR verbs that also appear to allow *until* with what seems at first glance to be a nonlocal negation.

- (i) You don't need to teach three classes until your third year.

A number of verbs of obligation exhibit this behavior.

- (ii) They won't require/order/ask you to teach three classes until your third year.

However, in all of these cases one can plausibly take *until your third year* to also be taking wide scope, where the semantics is such that the obligation does not take effect until the third year. This is not the case with *want*; the *until* clause in (16) does not modify the wanting time. Note, then, that the behavior of *need*, *require*, and so on, diverges from that of *want* when the *until* clause is preposed.

- (iii) a. Until your third year you don't need to teach three classes.
 b. Until your third year they won't require you to teach three classes.
 c. *Until spring semester, she does not want to take her sabbatical.

This is unsurprising if the *until* clause in (iiia) and (iiib) has wide scope, but the one in (iiic) does not.

⁶ Under the view of NPIs in Collins and Postal 2014a, one need not even look to cases of "classical" NR (i.e., interclausal NR) to make the point that VPE would have to be insensitive to meaning identity. This is because under that account, all NPIs underlyingly form a constituent with NEG that "raises" out by a much more local raising. Thus, (ia) is underlyingly (ib).

Let us take stock. If NR is a rule of grammar, and if the diagnostics above show that it is operative in these cases, then ellipsis ignores meaning quite dramatically. Consider again the VPE theories sketched above. Theory (b), semantic identity, says that there is actual linguistic material in the ellipsis site, and its meaning must be identical to the meaning of some other VP that is overt in the discourse context. Obviously, this cannot be right if NR does exist. Incidentally, there are versions of (b) that allow for slight differences in the meaning of the two. In the case at hand, though, the difference is hardly slight—it is a complete reversal of the polarity (moreover, of the polarity of embedded material). And we can easily show that—unsurprisingly—such polarity reversal is not possible. So, consider the Nader Trader scenario of the 2000 US presidential election. The scenario was this: An Al Gore supporter in a state that was “safe” for Gore (e.g., Massachusetts) was matched (via the Internet) with a Ralph Nader supporter in a swing state (e.g., Florida). The former agreed to vote for Nader if the latter would vote for Gore. Now, consider (20) in this context.

- (20) Nadine in Florida is willing to promise not to vote for Nader if Goran in Massachusetts is.

The second VP here cannot be understood as *is willing to promise to vote for Nader* even though this makes perfect sense in the Nader Trader context.

What are the consequences for theory (c), the deep anaphora view—that is, for the claim that the grammar supplies no meaning, but that the listener supplies a contextually salient meaning, which often—though not always—is made salient by having been named? Here it is a bit harder to pinpoint the predictions without a full theory of what counts as being contextually salient (note that it is equally difficult to pinpoint the predictions of theory (b) under the idea that “slight” meaning differences are allowed). But again, (20) clearly shows that embedded polarity reversal of the type needed here is generally not available.

Finally, we have yet to consider what happens under the proposal that ellipsis is subject to Rooth’s (1992) focus condition, with this possibly replacing (rather than just supplementing) an identity condition (as in, e.g., Fox 1999). Because the focus condition consults a larger domain, things are a bit different, but the end conclusion is the same. The focus condition assumes silent linguistic material and requires that an elided VP be within some constituent C such that there is another constituent C’, where $\llbracket C' \rrbracket$ is a member of the focus value of C. Put differently, C and

-
- (i) a. John didn’t vote for anyone.
b. John voted for [NEG someone]

But then consider (ii).

- (ii) John didn’t vote for anyone, but Bill did.

Of course, all theories need to account for the apparent “mismatch” in the ellipsis clause (since *anyone* would not be licensed there), but here there is the additional fact that the ellipsis allows not only a mismatch in form, but also a glaring mismatch in meaning. Thus, the elided VP is understood not as $\llbracket \text{vote for NEG-someone} \rrbracket$ (which is $\llbracket \text{vote for no one} \rrbracket$) but as—roughly— $\llbracket \text{vote for someone} \rrbracket$.

C' have the same meanings up to substitution of F(ocus)-marked material. Interestingly, if this is the only condition on ellipsis (i.e., if there is not also an identity condition), then (10) is actually compatible with the NR hypothesis. I leave this to the interested reader to verify.

Crucially, however, the reverse cases like (11) and (18) do not meet the focus condition if the first conjunct contains NR. Consider (18). Prosodic prominence on *do* in the second clause is optional, so assume that there is one version where this is not F-marked (*I* will be F-marked in any case). Then the meaning of the first conjunct (ignoring the contribution of *in years*) is the proposition 'Sally thinks [NEG [this cold in April]]', which is not a member of the focus value of I_F *think this cold in April*. So instead, assume that *do* is F-marked (with or without prosodic prominence). But the above proposition is still not a member of the focus value of I_F *do_F think this cold in April*—for those focus alternatives vary on whether or not the subject thinks it was cold; they are not about thinking that it was or wasn't cold. One might try to rescue this by positing some silent F-marked material in the elided portion of the second clause. But then we have no explanation for the failure of (20) in the Nader Trader scenario (note that [[Nadine is willing to promise not to vote for Nader]] is a member of the focus value of $Goran_F$ *is willing to promise to_F vote for Nader*). Hence, the NR hypothesis is incompatible not only with semantic identity and with deep anaphora, but also with Rooth's focus condition (whether or not it is supplemented with an identity condition).

There is one more tack we might try. An anonymous reviewer suggests that perhaps the focus account of ellipsis could accommodate the good cases like (11) and (18) if there is also a process raising a focused verum operator (VER) that can apply in just those cases where NR is also allowed. I illustrate with (11). The relevant representation for (11) (with NR) would be (21), where VER is not a silent operator in the embedded clause but instead has raised on the surface to the matrix clause.

(21) Joe think that Donald [NEG] would win, but Chris think that Donald [VER would win]

That the Nader Trader case in (20) shows that there can't be a lower VER is irrelevant here because the idea is that verum raising is allowed just where NR also is.

And in fact there are some NR cases where this is good; the reviewer constructs the following case using the Nader Trader scenario:

(22) ??Nadine in Florida is expected not to vote for Nader, but Goran in Mass. is.

But note that this is akin to cases like (12) and (13a), discussed earlier. While one can construct cases (e.g., (22) and perhaps (13a)) that are surprisingly good, this is not generally the case. If focused verum raising existed, then all of these cases should be just as good as the corresponding cases with the negation in the upper clause (such as (11) and all similar cases discussed in this article). But the contrast between the general situation like (11) (with NEG in the upper clause) and cases like (12) is robust. (See footnote 4 for informal informant work; the reviewer also mentions not liking all of these kinds of cases.)

Of course, we ultimately need to explain why these types of cases are sometimes improved from what one might expect under any identity condition (be it semantic or syntactic). One might

take such cases as evidence for the deep anaphora theory; there are cases where enough information and background context is available for the property of “not thinking p” to make salient the property of “thinking p” (and clearly it is the semantics of the NR predicates that helps this along: to not think p makes quite accessible the notion of “thinking not p,” whether one believes in NR or not). But this is not the end of the story, since one needs to explain why this sort of inference happens only sometimes, and why it happens when it does. This is tied in to the general fact that the deep anaphora theory by itself is quite weak and needs to be supplemented with a much richer set of predictions about what can be made salient without having been explicitly named. But providing that explanation is of course beyond the scope of this article.⁷

The upshot, though, is that a purely focus-based account of ellipsis, combined with the hypothesis that there is focused verum raising, will not account for the full range of facts, leaving us with a purely formal identity condition (and no semantic identity) as the only way to salvage NR. In other words, if NR is a rule of grammar, ellipsis must care only about formal identity, with no attention to meaning (and no trace of NEG in the lower clause). I thus now question whether that is plausible. If it is not, then ellipsis provides new evidence against NR.

2 Is the Requisite Account of VPE Plausible?

In section 1.2, I sketched various accounts of VPE; the only one that is obviously compatible with the above polarity mismatches under an NR account is account (a): formal identity only. The goal of this section is to show that this account is not very plausible for other reasons, casting doubt on the existence of NR.

There are many well-known cases of formal mismatch, for which various solutions have been proposed. For example, Fiengo and May’s (1994) vehicle change proposal is designed to account for mismatches of features relevant to pronouns, reflexives, and names; and there are other cases involving featural mismatches (I return to these below). Other kinds of formal mismatches (e.g., active/passive mismatches) have sometimes been dealt with by assuming that the grammar itself does not actually allow such mismatches, but speakers accommodate so as to get the obvious understanding (see Arregui et al. 2006). Incidentally, depending on one’s theory, some of these cases are also a challenge to complete meaning identity. If, for example, active VPs and passive VPs have different meanings (which is true in some but not all accounts), then a strict meaning-identity condition also needs to be slightly relaxed. Alternatively, such mismatches might provide evidence for the deep anaphora theory (where the relevant property is not named, but is easily inferable from something that is).

⁷ Another possibility makes use of the account of voice mismatches in Arregui et al. 2006. This account assumes a formal identity condition between the antecedent and the ellipsis site, such that voice mismatch cases (and cases like these) would be disallowed by the grammar. However, Arregui et al. also posit that listeners invoke a “repair strategy” and change the antecedent (in ways sanctioned by rules of the grammar) in such a way as to find an antecedent that *could* have been in the ellipsis site. Note that using this account for these cases (combined with Arregui et al.’s assumption that the repair strategy follows rules of the grammar) commits one to the view that NR exists, for there would have to be a grammatical rule allowing the repair to move NEG into the upper clause.

A thorough discussion of all of these cases and their implications is beyond the scope of this article. However, there is one case that has been discussed from time to time whose consequences have arguably been underappreciated: namely, mismatches of indexical pronouns and other indexicals. These are not like cases of passive/active mismatch, for here the relevant readings are completely smooth and do not in any obvious way involve some kind of accommodation, inferencing, or repair strategy.

Thus, as discussed at least as early as Hankamer and Sag 1984, (23) and (24) appear to show formal but not meaning mismatches.

(23) A: John is going to vote for me.

B: No, he isn't.

(24) A: Is John going to vote for you?

B: He is.

B's utterance in (23) is interpreted in the same way that the fuller VP *vote for you* is; it cannot have the form *vote for me*. The reverse is the case in (24); B's utterance is understood as 'vote for B', not 'vote for A'.

Before continuing, I should clarify the assumptions here regarding the meaning of indexical pronouns such as *I* and *you*. I embed this in Kaplan's (1989) theory whereby a sentence has both a character and a content. The character of (23A) is a function from speech contexts to propositions; it maps each speech context into the proposition that John will vote for the speaker in that context. The content is the proposition arrived at by applying the character function to the relevant context. In (23B), this is the proposition that John will vote for A. Under either the semantic identity theory or the deep anaphora theory, the understanding of B's utterance in each case will be the correct one assuming that what matters is content and not character.

As to the meaning of the indexicals themselves (as opposed to just the meaning of the entire VP), as a first pass I treat their character as functions from contexts to individuals. The character of *I* is a function from each speech context *C* to a function mapping each world to the individual that is the speaker in *C*. Because these pronouns are rigid designators (once context is fixed, they pick out the same individual in all worlds), I will henceforth skip the world step in the discussion.

There is, however, more to say. Sentences like (25) (attributed originally to Barbara Partee) show that—like other pronouns—these indexicals can function as bound variables.

(25) Only I love my mother.

The relevant reading here is one in which I am the only *x* who loves *x*'s mother. In view of this, one can take a slightly different perspective on the semantic contribution of the person features. In a theory with variables, all pronouns—including the indexical ones—have an index. For all assignments *g*, the character of I_i is a function mapping each context *C* to $g(i)$ provided that $g(i)$ is the speaker in *C*; it is undefined for *g* in all other contexts. This is recast in a variable-free semantics in Jacobson 2012, where all pronouns—including the indexicals—denote the identity function on individuals. The character of *I* is a function from any speech context *C* to the identity function whose domain is only the speaker in *C*. (Note that again I ignore worlds, and similar remarks hold for *you*.) Either of these options is consistent with the remarks here.

Moreover, the contribution of the first person features on the second pronoun *my* in (25) is also subject to debate. Von Stechow (2003), Kratzer (2009), and others have argued that these features are invisible to the semantics and are pure syntactic agreement features. Others (see Spathas 2009, Jacobson 2012) have argued that they have their usual semantic effect; they simply do not have an effect in the focus value (under the focus semantics of Rooth 1984). I return to this below. Thus, it is undoubtedly oversimplified to treat the character of, say, *I* as a function from speech contexts to individuals. But doing so temporarily is a harmless expository convenience; I return below to the more sophisticated accounts.

With that in mind, I turn again to the apparent formal mismatches in (23) and (24). In the discussion above, I used these to suggest that meaning (i.e., content) and not form is relevant for ellipsis, but perhaps that was too hasty. Is there a way to account for these mismatches while still maintaining that the identity condition requires only formal identity and is impervious to rather large meaning changes (such as those found in the NR cases)?

First, one might posit that certain features can simply be ignored or “switched” in the antecedent and the ellipsis site. (This is similar to, although not exactly like, Fiengo and May’s (1994) vehicle change proposal.) But the feature switch is not optional: B’s utterance in (23) can only be understood as ‘vote for A’, not ‘vote for B’. Moreover, the situation is much more complex: switches are (and must be) possible from 1st person to 3rd, 2nd to 3rd, exclusive ‘we’ to 2nd person plural or 3rd person plural—in fact, any combination one can imagine. The first of these switches is illustrated in (26), and the interested reader can construct additional ones.

(26) (Context: You and I are overhearing a conversation between a reporter and Hillary Clinton.)

Reporter to Clinton: “Do you think the convention will nominate you on the first ballot?”
I turn to you and say, “Oh, I’m sure it will.”

Hence, there appears to be no simple recipe for effecting the switch: it must switch to whatever is the appropriate form to preserve the meaning (“content”).⁸

But what about an attempt along the lines of the full proposal in Fiengo and May 1994: namely, that any of the features on a pronoun in the ellipsis site can be different *provided that the referential index stays the same*. Fiengo and May do not discuss indexicals in this context (but see Kitagawa 1991), but their proposal could extend to these. After all, we know from (25) that indexicals function like other pronouns, and so in a theory with indices, they must have an index. Under such a view, for any assignment *g* and speech context *C*, $\llbracket I_i \rrbracket$ in *C* is *g*(*i*) provided

⁸ Jason Merchant (pers. comm.) suggests that another possible strategy is to treat the indexical pronouns in, say, (23) and (24) as just bound pronouns by adopting a performative analysis so that there is indeed a higher *I*, *you*, and so on, to bind them. In that case, the features on these pronouns would be “ignored” (either in the focus value or by the semantics in general), just as they are in (25). But for this proposal to work, the higher hidden-performative clause would have to be extremely rich. Since all indexicals show the relevant behavior, and since we also see switches involving 3rd person (see (26)), the higher performative clause would minimally need to be something like *I say to you here and now about X-and-so . . .* And even this does not extend to indexicals like *today* and *tomorrow*, which also show the relevant pattern. Moreover, as shown below, indexicals such as *tomorrow* and *right here* resist a bound interpretation, rendering the analysis in which higher material in a performative clause “binds” the indexicals quite dubious.

that $g(i)$ is the speaker in C . Now suppose that the condition on ellipsis is formal identity, including the indices on pronouns, and that the other features are either stripped away or discounted. Indeed, Kitagawa (1991) puts forth this basic idea (in slightly different terms); I will not go into his particular proposal in detail as he was concerned with LF identity (via copying at LF) and of course here we are focusing on a theory with formal identity at a level not relevant to the semantics.

Still, the basic idea would be that a formal mismatch (at the relevant level) is allowed, as long as the indices are kept constant. The indices are used to ensure that the referent of the object in (23a) is the same as the understood object of the understood vote-for relation in (23b). In other words, identity of meaning is not directly imposed on ellipsis but is a by-product of a requirement on identical indices; the indices “track” the referents. Hence, (23) would have to have the representation in (27); we strip away the person features on the pronoun in B’s utterance.

- (27) A: John is going to vote for me_i .
 B: No, he isn’t ~~going to vote for~~ pro_i .

For any local speech context C and any assignment g , the value of $\llbracket me_i \rrbracket$ in A’s utterance is defined only if $g(i) = A$ (i.e., the speaker in C). The requirement that the index be held constant ensures that the silent pronoun in B’s utterance—while in a different local speech context—has the same index, and so $\llbracket pro_i \rrbracket$ is also A . There are no person features on this pro , so the fact that the local speech context in B’s utterance has a different speaker has no effect on the interpretation. (I assume that the assignment is held constant throughout the discourse.)

There are two difficulties with this. First, the existence of sloppy identity means that things are more complex; we cannot require full identity of indices and still allow both a strict and a sloppy reading for, say, (28).

- (28) A: I want Sally to write to me.
 B: I do too.

In fact, two of the well-known ways to account for sloppy identity in general are ruled out under the hypothesis here that formal identity is required in ellipsis and that, moreover, the requisite identity condition is on some representation that the semantics does not “see.” One simple account of sloppy identity is that the sloppy cases involve identity of meaning (see, e.g., Keenan 1971). Under that view, the meaning of the silent VP in B’s utterance (or, the understood property there) is the property $\lambda x[x \text{ want Sally write to } x]$, and this is the same as the meaning of the matrix VP (or of some LF constituent). (The simplicity of this account can itself be seen as suggesting that VPE actually cares about meaning, but here I continue to pursue the implications for a formal identity account.) Another account allows for differences in indices provided that the two are “alphabetic variants,” but the notion of two formal expressions being alphabetic variants is usually defined at LF (see, e.g., Sag 1976), and once again LF identity is not compatible with NR. Thus, in order to allow for a sloppy reading in (28B) (or, for that matter, in other cases involving sloppy identity) while using the index on the subject (and the deleted pronoun) as a way to track the referents in the strict cases, the ellipsis condition will have to be stated as “identical indices up to alphabetic variance,” where the definition of two things being alphabetic

variants must be stated at some non-LF level. I leave it to defenders of NR to formulate the appropriate definition of “alphabetic variance.”

But leaving sloppy identity aside, the attempt to bypass meaning by requiring identical indices is unlikely to extend to the full range of cases. For it is not just indexical pronouns that show the relevant behavior; other deictic expressions show it as well (Hankamer and Sag 1984). Thus, we find the same behavior with *now*, *then*, (*right*) *here*, *there*, *yesterday*, *today*, and so on. Consider for example (29) and (30).

(29) (Context: Sarah is on the phone, asking her husband to pull the alligator stew out of the freezer to defrost.)

He says, “Okay, I’m standing by the sink, do you want me to leave it right here?” She answers, “Yes, please do.”

(30) (Context: Sarah suddenly remembers that her friend Cathy is due to arrive the next day, although she isn’t sure.)

In the evening, Sarah dashes off an e-mail: “Oh gosh, I just realized—are you arriving tomorrow?” Cathy doesn’t read the e-mail until the next morning at the airport, and dashes back, “Yes, I am—but don’t worry, my plane doesn’t get in until five.”

There is, however, no independent motivation for thinking that the expressions *right here* and *tomorrow* (and many others that show this behavior) contain indices that would track their identity under ellipsis. In fact, while it is well-known that certain temporal expressions do behave like bound pronouns (and so—in a theory where indices are used in the account of binding—would have such indices), these particular expressions resist bound interpretations. Thus, compare (31), which shows the “bound” behavior of *I*, with (32); and compare (33) with (34). (Note that (31) and (33) are both good with either *I* or *they* in the *when(ever)* clause on the bound reading. Interestingly, all the bound pronouns must match; I will not pursue this point here.)

(31) I’m the only person who finally finished a paper when they/I said they/I would.

(32) Tomorrow is the only day that I’ll finally finish a paper due that day/*tomorrow.

(33) I’m the one person who always loses their/my keys whenever they/I put them in their/my desk drawer.

(34) Right here is the one place that I always lose my keys whenever I put them there/*right here.

Many other indexicals show the same pattern under VPE but resist a bound interpretation. For example, one can construct parallel data using *next Friday* instead of *tomorrow*.

Thus, the fact that *right here*, *tomorrow*, and other indexicals show the effect of preserving meaning (content) under ellipsis but resist bound interpretations indicates that the former effect cannot be reduced to a principle on preserving the referential indices. There is no motivation for providing these indexicals with such indices; in fact, were they to be treated this way, one would expect the illicit bound patterns shown above. Arguably, the indexicals are trying to tell us that meaning (content) and not form is what matters in ellipsis.

3 Conclusion

The upshot of these remarks is as follows. If NR is a rule of grammar, then VP-ellipsis can involve neither meaning identity nor deep anaphora, nor can it be subject to Rooth's (1992) focus condition. Large differences in meaning are ignored: an embedded VP can be understood with a polarity the reverse of the "antecedent's." And so, the grammar must care only about formal identity.

But the behavior of indexicals seems to be telling us just the opposite: any attempt to ignore the value (the content) of the indexicals in ellipsis in favor of some sort of formal identity seems problematic. The full picture, then, casts new doubt on the existence of NR. Of course, one might try to come up with some sort of "hybrid" theory, but I leave that to proponents of such a view. I will, however, end by noting that one can mix the indexicals and the NR facts into a single example, as follows:

- (35) A: Microsoft wants to hire me.
 B: Yeah, but Google doesn't until you finish your PhD.

References

- AnderBois, Scott. 2014. The semantics of sluicing: Beyond truth conditions. *Language* 90:887–926.
- Arregui, Ana, Charles Clifton, Jr., Lyn Frazier, and Keir Moulton. 2006. Processing elided VPs with flawed antecedents: The recycling hypothesis. *Journal of Memory and Language* 55:232–246.
- Bartsch, Renate. 1973. 'Negative transportation' gibt es nicht. *Linguistische Berichte* 27:1–7.
- Collins, Chris, and Paul Postal. 2014a. *Classical NEG raising: An essay on the syntax of negation*. Cambridge, MA: MIT Press.
- Collins, Chris, and Paul Postal. 2014b. Dispelling the cloud of unknowing. Ms., New York University. <http://ling.auf.net/lingbuzz/002269>.
- Fiengo, Robert, and Robert May. 1994. *Indices and identity*. Cambridge, MA: MIT Press.
- Fillmore, Charles J. 1963. The position of embedding transformations in a grammar. *Word* 19:208–231.
- Fox, Danny. 1999. Focus, parallelism and accommodation. In *Proceedings of SALT 9*, ed. by Tanya Matthews and Devon Strolovitch, 52–69. <http://journals.linguisticsociety.org/proceedings/index.php/SALT/issue/view/100>.
- Gajewski, Jon. 2007. Neg-raising and polarity. *Linguistics and Philosophy* 30:289–328.
- Gajewski, Jon. 2011. Licensing strong NPIs. *Natural Language Semantics* 19:109–148.
- Hankamer, Jorge, and Ivan Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry* 7:391–426.
- Hankamer, Jorge, and Ivan Sag. 1984. Towards a theory of anaphoric processing. *Linguistics and Philosophy* 7:325–345.
- Hardt, Daniel. 1993. Verb phrase ellipsis: Form, meaning and processing. Doctoral dissertation, University of Pennsylvania, Philadelphia. Distributed as IRCS report 93-23.
- Horn, Laurence. 1971. Negative transportation: Unsafe at any speed? In *Papers from the 7th Regional Meeting of the Chicago Linguistic Society*, 120–133. Chicago: University of Chicago, Chicago Linguistic Society.
- Horn, Laurence. 1975. Neg-raising predicates: Toward an explanation. In *Papers from the 11th Regional Meeting of the Chicago Linguistic Society*, ed. by Robin E. Grossman, L. James San, and Timothy J. Vance, 280–294. Chicago: University of Chicago, Chicago Linguistic Society.
- Horn, Laurence. 1978. Remarks on neg-raising. In *Syntax and Semantics 9: Pragmatics*, ed. by Peter Cole, 129–229. New York: Academic Press.
- Horn, Laurence. 2014. The cloud of unknowing. In *Black book: A festschrift in honor of Frans Zwartz*, ed. by Jack Hoeksema and Dicky Gilbers, 178–196. Groningen: University of Groningen.

- Horn, Laurence, and Samuel Bayer. 1984. Short-circuited implicature: A negative contribution. *Linguistics and Philosophy* 7:397–414.
- Jacobson, Pauline. 2003. Binding without pronouns (and pronouns without binding). In *Resource sensitivity, binding, and anaphora*, ed. by Geert-Jan Kruijff and Richard Oehrle, 57–96. Dordrecht: Kluwer.
- Jacobson, Pauline. 2006. I can't seem to figure this out. In *Drawing the boundaries of meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn*, ed. by Betty Birner and Gregory Ward, 157–175. Amsterdam: John Benjamins.
- Jacobson, Pauline. 2012. Direct compositionality and 'uninterpretability': The case of (sometimes) uninterpretable features on pronouns. *Journal of Semantics* 29:305–343.
- Kaplan, David. 1989. Demonstratives. In *Themes from Kaplan*, ed. by Joseph Almog, John Perry, and Howard K. Wettstein, 481–563. Oxford: Oxford University Press.
- Keenan, Edward L. 1971. Names, quantifiers, and a solution to the sloppy identity problem. *Papers in Linguistics* 4:211–232.
- Kitagawa, Yoshihisa. 1991. Copying identity. *Natural Language and Linguistic Theory* 9:497–536.
- Kratzer, Angelika. 2009. Making a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40:187–237.
- Kroll, Margaret. 2016. Polarity reversals under sluicing. Paper presented at Berkeley Syntax and Semantics Circle (April 2016). Handout available at https://people.ucsc.edu/~makroll/uploads/2/0/8/0/20801778/kroll_berkeley_handout.pdf.
- Lakoff, Robin. 1969. A syntactic argument for negative transportation. In *Papers from the 5th Regional Meeting of the Chicago Linguistic Society*, ed. by Robert I. Binnick, Alice Davison, Georgia M. Green, and Jerry L. Morgan, 140–147. Chicago: University of Chicago, Chicago Linguistic Society.
- Lindholm, James. 1969. Negative-raising and sentence pronominalization. In *Papers from the 5th Regional Meeting of the Chicago Linguistic Society*, ed. by Robert I. Binnick, Alice Davison, Georgia M. Green, and Jerry L. Morgan, 148–158. Chicago: University of Chicago, Chicago Linguistic Society.
- Merchant, Jason. 2001. *The syntax of silence: Sluicing, islands, and the theory of ellipsis*. Oxford: Oxford University Press.
- Miller, Philip, and Geoffrey K. Pullum. 2013. Exophoric VP ellipsis. In *The core and the periphery: Data-driven perspectives on syntax inspired by Ivan A. Sag*, ed. by Philip Hofmeister and Elisabeth Norcliffe, 5–32. Stanford, CA: CSLI Publications.
- Rooth, Mats. 1984. Association with focus. Doctoral dissertation, University of Massachusetts, Amherst.
- Rooth, Mats. 1992. Ellipsis redundancy and reduction redundancy. In *Proceedings of the Stuttgart Ellipsis Workshop*, ed. by Steven Berman and Arild Hestvig. Stuttgart: University of Stuttgart.
- Ross, John Robert. 1967. Constraints on variables in syntax. Doctoral dissertation, MIT, Cambridge, MA. Distributed by Indiana University Linguistics Club.
- Sag, Ivan. 1976. Deletion and Logical Form. Doctoral dissertation, MIT, Cambridge, MA. Distributed by Indiana University Linguistics Club.
- Spathas, George. 2009. Interpreting gender features on bound pronouns. In *Proceedings of the 2007 Workshop in Greek Syntax and Semantics*, ed. by Claire Halpert, Jeremy Hartman, and David Hill. MIT Working Papers in Linguistics 57. Cambridge, MA: MIT, MIT Working Papers in Linguistics.
- Stechow, Arnim von. 2003. Feature deletion under semantic identity. In *NELS 33*, ed. by Makoto Kadowaki and Shigeto Kawahara, 133–157. Amherst: University of Massachusetts, Graduate Linguistic Student Association.
- Williams, Edwin. 1977. Discourse and Logical Form. *Linguistic Inquiry* 8:101–139.

Department of Cognitive, Linguistic and Psychological Sciences
 Box 1821
 Brown University
 Providence, RI 02912
 Pauline_jacobson@brown.edu