

Semantics in Context

Forthcoming in Contextualism, Gerhard Preyer, ed. (OUP).

Jason Stanley

Consider an utterance of the sentence "Some philosophers are from New York". If no philosopher in the world comes from New York, competent speakers of English know that it is false. They also know that this utterance is true, if six philosophers in the world come from New York. In other words, competent English speakers have clear intuitions about the conditions under which what is said by an utterance of this sentence is true or false.

The apparent *source* of such intuitions is not difficult to locate. Competent English speakers know the meanings of the words in the sentence "Some philosophers are from New York." They also know how to combine the meanings of each of the words in this sentence to arrive at what is said by the utterance of the sentence, "Some philosophers are from New York." It is that linguistic competence that seems to be the source of their ability to report correctly about the truth of what is said by that sentence relative to different possible circumstances, e.g. the circumstance in which there are no philosophers from New York, or the circumstance in which six philosophers come from New York.

So, the explanation for our ability to report about the truth and falsity of what is said by an utterance of "Some philosophers are from New York" in various possible situations is as follows. Competent English speakers know the meanings of the words used, and understand how they are combined. Their grasp of the truth-conditions of the utterance of that sentence is due to their ability to combine the meanings of the words, relative to the context of utterance.

With this explanation in mind, consider an utterance of the sentence "Every philosopher is from New York", made at a small philosophy conference. It is natural to take this utterance to say something that is true if and only if every philosopher *at the conference* is from New York. If we cleave to the model of understanding just described, we will seek to explain our understanding of the truth-conditions of this utterance by appeal to a process of combining the elements of the sentence "Every philosopher is from New York", using our understanding of the words used in the sentence. But of course,

there appears to be no expression in the sentence "Every philosopher is from New York" that corresponds to the understood constituent expressed by "at this conference".

Similarly, suppose, pointing at a 5 foot tall seven year old child, I utter the sentence "He is tall." I am most naturally understood as saying something that is true if and only if the child in question is tall *for a seven year old child*. Preserving the model of understanding we began with, according to which our intuitions about the truth-conditions of an utterance are due to a process of combining meanings of the parts of the sentence uttered, would require us to find some constituent in the sentence that could be taken to supply the understood property of *being a seven year old child*. But again, it appears that the sentence "He is tall" contains no such constituent.

So, we have a predicament. If we look at certain sentences, there seems to be a clear and elegant explanation of why we have the intuitions we do about the truth conditions of utterances of those sentences. But if we consider utterances of other sentences, the explanation appears to break down. The first response to this predicament is to attempt to preserve the clear and elegant explanation in the face of the apparently recalcitrant data. The second is to abandon the clear and elegant explanation of the source of our truth-conditional intuitions in favor of a different one.

My concern with the second response to the predicament is that the suggestions I am aware of for dealing with the additional complexity essentially end up abandoning the project of giving a systematic explanation of the source of our intuitions. They invariably involve appeal to unconstrained and non-explanatory notions or processes (cf. Stanley (2002a)). I have therefore been inclined to pursue the first of these options (cf. Stanley (2000)).

My purpose in this paper is to continue the project of defending the clear and elegant explanation of the source of our intuitions about the truth-conditions of utterances. I will do so by considering some replies to previous arguments in favor of it. I will argue that proponents of abandoning the clear and elegant explanation have not yet made their case.

Section I. The Challenge from Context-Sensitivity

On the simple explanation of the source of our intuitions about the truth-conditions of utterances of sentences we understand, it is due primarily to a compositional process of interpretation. Our knowledge of meaning, together with our knowledge of relevant contextual facts, allows us to assign meanings to the parts of a sentence, and the intuitive truth-conditions of an utterance of that sentence are what results from combining these values. Somewhat tendentially, I will call proponents of the simple explanation, *semanticists*.

Innumerable researchers from pragmatics have challenged the semanticist's model. Here is the form of the standard challenge. First, a linguistic construction C is produced that appears intuitively to have a certain reading R. Secondly, the researcher claims that the readings cannot be due to the semantics of that construction. That is, the claim is that R cannot be due to the compositional semantic interpretation of C, relative to the envisaged context of use. The conclusion the researcher draws is that the assumption that the intuitive truth-conditions of a sentence relative to a context are due to semantics is incorrect.

A large number of researchers opposed to the semanticist employ arguments of this sort (a brief list of the most prominent exponents includes Kent Bach, Herman Cappelen, Robyn Carston, Ernie Lepore, Stephen Levinson, Francois Recanati, Dan Sperber, Charles Travis, and Deirdre Wilson; there are many more). Typically, such researchers do not just supply a single example, but a list of disparate examples. For example, the following is a representative list that could occur in any one of a hundred papers written in the past decade by researchers in this tradition:

- (1) John is tall. (for a fifth grader)
- (2) John is finished. (with grading)
- (3) Every boy (in the class) is seated.
- (4) John and Mary went to Paris (together/separately)
- (5) If Lincoln hadn't gone to the theater, he wouldn't have been assassinated. (fixing certain background assumptions)

- (6) John ate breakfast. (this morning)
- (7) John had breakfast this morning. (in the normal way, through his mouth)
- (8) John ate. (mushrooms)
- (9) The ham sandwich (person who ordered the ham sandwich) is getting annoyed.
- (10) The apple is green. (on the inside)

The bracketed material is intended to indicate the material that cannot be provided by semantics, but only by pragmatics.

As one may imagine, any such list will include cases that virtually all proponents of the simple model of interpretation believe are uncontroversially generated by the compositional interpretation of the sentence uttered. For example, there is much recent investigation into the syntax and semantics of gradable adjectives that generates the supposedly pragmatic material as the semantic value of some element in the syntax (either a comparison class variable, or a degree variable). Semantic treatments of plurals treat collective-distributive ambiguities in a variety of ways. For example, some treat collective-distributive ambiguities as structural ambiguities due to the relative scope of an event quantifier. On such a treatment, the two readings of (4) are due to scope facts in the syntax, rather than pragmatics. Since Partee (1973), most linguists have defended the view that verbs are associated with temporal variables, that have their references filled via deixis. On this view, generating the relevant reading of (6) simply involves speaker intentions determining the value of a temporal variable in the syntactic structure of the sentence uttered. Finally, few semanticists would balk at associating the provision of an accessibility relation for counterfactuals to some element in the syntax, either the conditionals words themselves, or some covert element.

On the other side of the spectrum, some of the examples that are provided in such lists seem to require pragmatic treatment, on the grounds that the alleged intuitive truth-conditions are richer than those delivered by tutored intuitions about truth-conditions. For example, as I will argue in the final section, with the help of recent work by Luisa Marti, it is clearly accessible to a native speaker of English that it is no part of the truth-conditions of an utterance of (7) that John ingested his breakfast through his mouth. Of course, when someone tells us that John ate this morning, we assume he did so in the

normal way. But no one would deem an utterance of (7) *false* if, contrary to default assumptions, they discovered that John ingested breakfast in some non-standard way, such as being spoon fed. So the manner of eating is no part of the intuitive truth-conditions of (7), but is rather pragmatically conveyed information. We also assume other things when we hear an utterance of (7), e.g. that John's breakfast wasn't prepared by a Martian. But none of this is information that is carried semantically, and, *pace* Carston (2002, p. 203) and Wilson and Sperber (2002)), it is odd to suppose that anyone has ever advanced a theoretical position that would commit them otherwise.

Nevertheless, between the two extremes I have just discussed, there are some examples that are genuinely worrying for the semanticist. For instance, it certainly appears that the intuitive truth-conditions of an utterance of (9) involve a person, rather than a ham sandwich. Yet it's not clear that a process that maps ham sandwiches onto persons counts as genuinely semantic. To take another example, the intuitive truth-conditions of (3) certainly involve reference to a domain of quantification. But if domain restriction is a matter of information being freely provided by context, that too does not seem to be a process that can be considered genuinely semantic.

Some have tried to respond to this predicament by arguing that the semantic content of a sentence, relative to a context, is only a minimal part of the intuitive truth-conditions of that utterance, a version of what King and Stanley (forthcoming) call *semantic modesty*. As King and Stanley emphasize, the worry with this response is that it is unclear what role the minimal semantic core ends up playing in an account of the intuitive truth-conditions, if one accepts that processes such as free enrichment account for much of our intuitions in examples such as (1)-(10). If free pragmatic enrichment is a process hearers regularly use to interpret utterances, and speakers are aware of this, then why can't speakers utter sentences whose semantic content is minimal or vacuous, and rely on such pragmatic processes to do the bulk of the expressive work? So I am not sanguine about semantic modesty as an intermediate position for the semanticist.

So much the worse for the semanticist, one might think. However, if our intuitions about truth and falsity are responsive to processes that are not linguistically controlled, we need an *explanatory* account of information freely provided by context. And it's not clear that such an account is in the offing. The most serious problem facing the advocate

of free pragmatic enrichment to intuitive truth-conditions is that of *over-generation* (cf. Stanley (2002a)). If our intuitions about the truth-conditions of utterances of quantified sentences are due to a process of free pragmatic enrichment, then it would be a mystery why utterances of certain sentences lack certain readings. For example, why is it the case that an utterance of (11) can express the same proposition as an utterance of (12), but never the same proposition as an utterance of (13)?

- (11) Every Frenchman is seated.
- (12) Every Frenchman in the classroom is seated.
- (13) Every Frenchman or Dutchman is seated.

In short, since these "enrichment" processes are not linguistically constrained, they should be constrained only by general pragmatic reasoning. But why do general pragmatic facts allow (11) to express (12) but not (13)?

In the light of these worries with free pragmatic enrichment accounts of intuitive truth-conditions, it is important to investigate the possibility that the intuitive truth-conditions of utterances of sentences such as those in (1)-(10) are due to linguistically determined content. Linguistically determined content is content that is constrained not just by pragmatic means. Particularized conversational implicatures, for example, are constrained only by pragmatic means, and hence are not part of linguistically determined content. In contrast, the value of a term such as "she", relative to a context, is linguistically determined, because the speaker intentions that determine its value must be referential intentions consistent with the literal meaning of "she". So the question posed by such examples is how to establish that the intuitive readings of the problematic sentences in (1)-(10) are due to linguistically determined content.

Section II. Responding to the Challenge

I adopt the conception of semantics at work in Stanley (2000) and spelled out in detail in King and Stanley (2004). The semantic content of a sentence relative to a context is derived by taking the semantic contents of the parts of that sentence, relative to

that context, and composing them in accord with the composition rules governing the syntactic structure of that sentence. The semantic value of a basic constituent of a sentence is what is determined by speaker intentions together with features of the context, in accord with the standing meaning of that lexical item. Given this conception of semantics, the position of the semanticist is then that the source of our intuitions about the truth and falsity of utterances relative to various possible circumstances is due to semantics.

When faced with the claim that a certain construction C has a reading R that *prima facie* does not seem traceable to the semantics, the semanticist has three options. The first option is to establish that the alleged reading is not part of the intuitive truth-conditions of an utterance of that sentence, but is instead due to the pragmatics (as in the discussion of (7), above). The second option is to argue that the claim that reading R is not due to the semantics is due to an overly simplistic conception of the semantic content of some elements of C. When the correct semantics for the relevant expression is given, the reading does emerge from the semantics (cf. King and Stanley (2004, Section 5), on "implicature intrusion"). The third option is to argue that the claim that reading R is not to the semantics is due to an overly simplistic conception of the syntactic structure of C. In fact, C contains covert structure, and once this is recognized, reading R does emerge from the semantics (cf. Stanley and Szabo (2000) on domain restriction, and Stanley (2000) for discussion of other constructions).

So, when faced with a list such as that given in the previous section, the semanticist has, in each case, three alternatives. The first is to reject the semantic significance of the data, the second is to give an alternative semantic assignment to some overt element, and the third is to argue for covert syntactic structure. As I have indicated, it is a construction specific matter which of these options is preferable. The difficulty facing the semanticist's opponent is that she must establish, for each case, that none of the three very different alternatives is available as an account of the data.¹

Of all the constructions on the list, I think the central worry for the semanticist is (9), the case of deferred reference. Not only is there is a strong intuition that the deferred

meaning is part of the intuitive truth-conditions, but the deferred meaning enters into certain linguistic processes, such as anaphora and ellipsis. For example, the natural reading of (14) is one in which the anaphoric element "his" receives its value not from the "literal" content of "the ham sandwich", but from its deferred meaning:

(14) The ham sandwich wants his bill now.

Similarly, when we consider someone uttering (15) in the kitchen of a restaurant, describing the predicaments of two waiters, it is the deferred meaning of "an annoying ham sandwich" that is carried over to the ellided constituent:

(15) Bill served a ham sandwich, and John did too.

In particular, (15) cannot be interpreted as conveying that Bill served a person who ordered a ham sandwich, whereas John served a ham sandwich. Finally, one could argue that the literal meaning of an expression provides a guide for its deferred meaning, and so the deferred meaning is semantic after all.

I think neither of these points show that deferred reference is semantic. In a nominative metaphor such as (16a), we see the same phenomenon as in (14), and in (16b), we see the same phenomenon as in (15), where the metaphorical reading is what is carried over in ellipsis:

(16) a. The pig in the next room wants his check immediately.
b. John is a pig, and Bill is too.

We cannot interpret (16b) to mean that John is a person who is a sloppy eater, and Bill is (e.g.) a pig, perhaps John's pet. But on a standard view of metaphor, metaphor is not semantic. If the metaphorical meaning of an expression does not affect the semantic content of sentences containing it, relative to a context, then the fact that deferred

¹ The only pragmaticist I know of who seems to recognize the daunting challenge this poses to the opponent of the semanticist is Stephen Levinson (cf. Levinson (2000, p.

reference behaves in a similar manner should not lead us to believe that deferred reference is semantic.

Joseph Stern (2000, pp. 69-70) has recently used facts such as (16b) to argue that metaphor affects semantic content, that (as he would put it), there is such a thing as semantically significant metaphorical *meaning*. But, as Elizabeth Camp (ms.) has pointed out, we also see the same phenomenon with *irony*. Consider:

- (17) John: Bill is a fine friend.
Sally: Sue is too.

If John's utterance is intended ironically, then the ellided constituent "fine friend" in Sally's utterance must be understood ironically as well. But this does not show that irony is semantic, or that there is such a thing as semantically significant ironical *meaning*. As Stern (Ibid., p. 232) writes, "Now, whatever controversy surrounds the status of metaphorical meaning, the ironic 'meaning' of an utterance is surely not a semantic meaning." So, such ellipsis facts do not demonstrate that a phenomenon is semantic.

The second argument that deferred reference is semantic is that the literal meaning of an expression in context provides a guide to its deferred meaning. For example, the literal meaning of "the ham sandwich" provides a guide to the deferred meaning of "the ham sandwich" in (14), which is *the person who ordered the ham sandwich*. So if the mark of the semantic is guidance (in some sense) by literal meaning, then there is evidence that deferred reference is semantic.

But it is also the case that the literal meaning of "has nice handwriting", in the context of an utterance of "John has nice handwriting" in a reference letter, provides a guide for the implicated property, *is a bad philosopher*. So the fact that the literal meaning is used in deriving the deferred meaning does not show that the deferred meaning is linguistically controlled in the relevant sense (i.e. semantic).²

I think a more general argument can be given that deferred reference should not be treated as semantic. The mark of the semantic is that semantic content is *constrained*

214)).

² Thanks to Hanna Kim for discussion here.

by linguistic meaning. At the very least, the semantic content of an expression, relative to a context, must be something of which that expression is true. If it is not, it is hard to see how the semantic content of that expression has been constrained by the conventional meaning of that expression. But in the case of deferred reference, that is not true. If deferred reference were semantic, the denotation of "the ham sandwich" would be something of which the predicate "ham sandwich" were not true. So it is hard to see how the deferred reference of "the ham sandwich", in (9), is semantic, since it is not constrained by the conventional meaning of the words used. So, one theoretical consideration that should lead us to deny that deferred reference (as in example (9)) is not semantic is that the deferred reference of an expression is not something of which the conventional meaning need be true.³

A second consideration involves the *scope* of the phenomenon. One reason against taking metaphor to be semantic is that virtually any term can be used metaphorically. This suggests that metaphor has to do with the *use* of a term, rather than the semantics of a particular expression. Similarly, virtually any term can be used with a deferred reference.⁴ This suggests that the phenomenon of deferred reference does not have to do with the semantics of any particular construction. Rather, it involves how we can *use* constructions that have a certain semantics to communicate something different than such constructions semantically express.

A final theoretical consideration that can be brought to bear in arguing that deferred reference is not semantic has to do with the unconstrained nature of any semantic theory adequate to the task. This emerges when one considers the details of the semantic resources one would need to adopt in order to incorporate deferred reference into the semantics. Sag (1981) gives a semantic theory appropriate to the task of

³ As Jeff King has pointed out to me, this distinguishes deferred reference from deferred *ostension*. Suppose, pointing at a parked car festooned with tickets, I utter "That driver is going to be upset." The reference of "that driver" is the driver of the indicated car, even though what I demonstrated is the car. But the driver is still who is denoted by my use of "that driver", because he satisfies the predicative material "driver". This distinguishes deferred reference from deferred ostension; the former is not semantic, whereas the latter is.

⁴ For example, we can have "Two ham sandwiches are getting irritated", "Every ham sandwich is clamoring for her check", "John ham-sandwiched again" (where this latter may mean the same as "John ordered a ham sandwich again").

incorporating deferred reference into semantic content. Sag introduces "sense-transfer functions" into contexts, and then uses them to interpret expressions in a sentence interpreted relative to a context. On his account, an expression is interpreted relative to a sense-transfer function, which can map the meaning of that expression onto any other meaning. The class of sense-transfer functions is restricted only by pragmatics.

Something like Sag's semantic proposal is required to account for deferred reference. But notice what the resulting "semantic" theory has the power to do. In no sense can it be said that semantic content is "constrained" by conventional meaning. Since, as we just discussed, virtually any word can have a deferred meaning, it follows that any word could in principle acquire any meaning, via a sense-transfer function. The available sense-transfer functions are constrained only by pragmatics. So, the resulting semantic theory is one according to which semantic content is unconstrained by conventional meaning. The semantic content of the word "house" could be the property of being a dog -- the only thing that would prevent it from acquiring this semantic content is pragmatic facts about a context.

The moral of this final consideration is that, to capture deferred reference semantically, one would need to adopt a semantic theory where semantic content is not constrained by conventional meaning; in short, an unconstrained semantic theory (that is, constrained only by pragmatics). When capturing a phenomenon within the semantics would result in an unconstrained semantic theory, that suggests that the phenomenon is not semantic. For example, if in order to capture a phenomenon within the semantics, one needs to exploit resources that could allow the semantic content of "Grass is green", relative to a context, be the proposition that snow is white, then the phenomenon is not semantic. This is the principal theoretical reason for denying that deferred reference is semantic.

So, I have given two theoretical reasons for denying that deferred reference is semantic. These considerations are not arguments based on intuitions. As I have already indicated, there is a sense of "intuitive truth-conditions" in which deferred reference enters into intuitive truth-conditions. So one might think that to draw the distinction between semantic content and what is only pragmatic in such a way that the deferred reference of a use of an expression is not part of the semantic content is to abandon the

semanticist's view that semantics is the source of our intuitions about the truth-conditions of an utterance.

I don't think that any reasonable way of delineating the border between the semantic and the non-semantic will deliver results that will satisfy all. The responsibility of the semanticist is rather to provide some way of drawing the distinction that preserves the core semanticist claim that the source of our intuitions about truth-conditions is the semantics. Cases like deferred reference are cases in which tutored intuitions *diverge*. It is certainly the case that the non-deferred meaning of (9) is available to all competent users of the language, as in the discourse:

(18) A: The ham sandwich is getting annoyed.

B: That's absurd; sandwiches do not get annoyed.

In such a case, where the putatively literal semantic content is clearly available to all competent users of the language, it is perfectly permissible to let theoretical considerations decide between the putatively literal semantic content and the enriched content (that is, the content enriched with the deferred meaning). This is consistent with the semanticist's position, since this is a case in which speakers have several intuitions easily available to them.

The case of deferred reference contrasts, then, with the case of comparative adjectives. Suppose that a theorist maintained (cf. Cappelen and Lepore (this volume)) that the semantic contribution of "tall" was something like the semantic content "tall for some comparison class", so that everything in the universe except the smallest thing is tall. Suppose then that I showed a speaker a picture of a tiny dwarfish man, surrounded by normal sized men. Pointing at the dwarfish man, I uttered "That man is tall". On the envisaged theory, the semantic content of my sentence, relative to this context, is a true proposition. The person in question is tall, relative to some comparison class (e.g. the class of mice). But this semantic content is utterly inaccessible to the speaker. Unlike the case of deferred reference, there is no possibility of a sensible discourse along the following lines:

(19) A (pointing at the dwarfish man): That person is not tall.

B: That's absurd; everyone and everything is tall, except for the smallest thing.

What this indicates is that the putative semantic content -- that the indicated person is tall for some comparison class -- is not available to the competent user of the language.

Therefore, it is not consistent with the view I am suggesting to take it as the actual semantic content of the sentence, in context.⁵

I have said that, when the putative semantic content is clearly accessible and tutored intuitions about semantic content diverge, theoretical considerations may enter in to decide where to draw the line between semantic content proper and the rest of what is conveyed in a speech act. As we have seen, deferred reference is one such case. In this case, I gave two theoretical reasons to take the semantic content of a sentence not to be sensitive to deferred meanings. It is instructive to look at another such case in which tutored intuitions may diverge, but theoretical considerations impel us to draw a different sort of line between the semantic and the non-semantic; the case of domain restriction.

When a sentence such as (3) or (20) is uttered, we naturally interpret it with respect to a salient domain of quantification:

(20) Every bottle is in fridge.

For example, (20) could be used to communicate the proposition that every bottle in the house is in the fridge. However, like the case of deferred reference (and unlike the case of comparative adjectives), the unrestricted interpretation is also available to competent language users, as the coherence of the following sort of discourse illustrates:

(21) A: Every bottle is in the fridge.

⁵ There are other powerful objections against the view in question. For example, "tall", like other comparative adjectives, is *gradable*. On a degree theoretic view, the function of an intensifier such as "very" is to raise the contextually salient degree of height that something must meet in order to be tall. But, on the Cappelen and Lepore view, it is mysterious what the semantic function of "very" would be in a sentence such as "Bill is tall, but John is very tall."

B: Well, your fridge couldn't possibly be that large! There are bottles somewhere in the world that aren't in your fridge.

So, like the case of deferred reference, though intuitions are sensitive to the domain of quantification, it is nevertheless possible for competent speakers to detect the unrestricted reading of quantified sentences. If, as in the case of deferred reference, there were overwhelming theoretical considerations that mitigated against building the restricted reading of quantified sentences into the semantics, then it would then be acceptable to do so, consistently with the thesis that semantic content delivers intuitive truth-conditions.

However, there are no good theoretical reasons against incorporating domain restriction in the semantics. As we saw, incorporating deferred meaning into semantic content has two disturbing results. First, the semantic content of an expression may be something that does not satisfy the conventional meaning of that expression. Secondly, in order to treat the phenomenon, one needs to employ resources that trivialize the semantics. In contrast, incorporating domain restriction into the semantics brings no such costs.

On the theory of domain restriction in Stanley and Szabo (2000) and Stanley (2002b), the effect of domain restriction is to restrict the extension of the head noun in a quantified noun phrase. That is, in a sentence such as (19), the effect of domain restriction is just to restrict the interpretation of the property expressed by "bottle", by intersecting its extension with the extension of the property that is the domain restriction. The semantic content of the result will be a subset of the set of bottles. So, the semantic content of the restriction of "bottle" will be something that satisfies the conventional meaning of "bottle". Secondly, in incorporating domain restriction into the semantics, there is no risk of giving the semantics the resources to make "Grass is green" express the proposition that snow is white. The only effect context can have is to restrict the interpretation delivered by the conventional meaning of the head noun in a quantified noun phrase. So, incorporating domain restriction into the semantics is perfectly consistent with the nature of semantic content as intrinsically constrained by conventional meaning.

Since incorporating domain restriction into the semantics does not have theoretical costs, given that domain restriction does affect some level of intuitive truth-conditions, it ought to be incorporated into the semantics. Of course, it is only possible to incorporate domain restriction into the semantics if it is due to semantics, that is, due to the compositional assignment of content to a sentence in context. In previous work, I have argued that there is covert structure in quantified noun phrases to which provision of a domain to the semantic content of the sentence containing that noun phrase is due. One argument I have used for this conclusion (in the case of domain restriction as well as other constructions) is what has since been called *the binding argument*. Note that the sentences in (22) are most naturally interpreted as in (23):

- (22) a. In every room, every bottle is in the corner.
b. Every student answered every question.

- (23) a. In every room r , every bottle in r is in the corner.
b. Every student x answered every question y on x 's exam.

One way to generate the readings in (23) is to suppose that there are bound variables in the structure of quantified noun phrases, whose values, relative to a context, generate a domain of quantification.

More specifically, the theory of domain restriction I favor (Stanley and Szabo (2000), Stanley (2002b)) captures these readings in the following way. Syntactically associated with each nominal are domain restriction indices, of the form ' $f(i)$ '. Relative to a context, ' f ' is assigned a function from objects to properties, and ' i ' is assigned an object.⁶ So, the syntactic structure of the sentences in (23) is similar to the sentences in (24):

- (24) a. Every fireman is tired.

⁶ To my knowledge, the need for such a function variable in an account of domain restriction was first pointed out in Von Stechow (1994, p. 31). Von Stechow's theory differs from Stanley and Szabo (2000) in that his representations associate the domain indices with determiners, rather than nominals.

- b. Every student answered every question.
- (25) a. Every <fireman, f(i)> is tired.
 b. [Every <student, f(j)>]-i answered every <question, f(i)>.

Given an utterance of, e.g., (23b), the speaker intends the value of 'f' to be a function from students to their exams, and 'i' is bound by the higher quantifier 'every student', yielding the desired reading (23b).⁷

If these are the right representations, then domain restriction is due to the semantics, since it is due to the assignment of values to constituents of a sentence, relative to a context. Evidence that these are the right representations comes from the fact that one detects operator-variable interactions involving quantifier domains, and operator-variable interactions are syntactic in nature.

Of course, binding considerations are certainly not the only way to argue that an allegedly non-semantic phenomenon is due to semantics. For example, the view that the phenomenon in question is non-semantic could be due to an overly simplistic conception of the semantics of some overt expression, and so one way of establishing that the phenomenon is semantic is by giving a more complex semantic clause for some overt expression (cf. again King and Stanley (2004) on implicature intrusion). Furthermore, binding considerations are not the only way to establish covert structure, since nothing bars the language system from employing syntactic structures containing covert non-bindable indexicals, akin to the overt non-bindable indexicals "I" and "here" of English. But binding considerations are still one way to argue for covert structure, and one that generalizes to a wide range of constructions (cf. Stanley (2000)). Because such considerations do provide an argument for the semantic treatment of a wide range of data that pragmaticists have long claimed to be non-semantic in nature, they have recently been widely criticized. In the rest of this paper, I will look at some of the criticisms of the argument from binding, to see whether they undermine the status of these considerations as arguments for syntactic structure.

Section III. The Binding Argument

According to the binding argument, if there is a genuine bound reading of a certain construction, that supports the hypothesis that the quantifier in question binds a variable in the syntactic structure of the sentence. For the binding argument to have force, the bound reading must be generated by an expression that is an uncontroversial example of a quantificational expression. The binding argument, *considered as an argument for syntactic structure*, has been interpreted in several different ways. In this section, I discuss the three different ways it has been interpreted.⁸

On the first interpretation of the binding argument, which occurs in unpublished work by Michael Nelson, and in Cappelen and Lepore (2002), the binding argument establishes the existence of covert structure, on pain of ungrammaticality due to vacuous binding. For instance, in the case of (18a), the quantifier "every room" must bind a variable in the syntactic structure of the sentence "every bottle is in the corner", on pain of ungrammaticality. On the second interpretation of the binding argument, a bound reading of a sentence is evidence for syntactic structure, since bound readings are the semantic effect of a syntactic process (see Stanley (2000, pp. 412ff.) for details). On this version of the binding argument, it is not potential ungrammaticality that is at issue. Rather, certain kinds of semantic phenomena (e.g. bound readings, scope ambiguities) have ultimately a syntactic explanation. On the third (and weakest) interpretation of the binding argument, it is an inference to the best explanation. By postulating a covert variable, one can account for the bound reading, and there is no other satisfactory way to

⁷ The values for the domain indices for the first nominal "student" could be, e.g., the classroom (for 'j') and a function from the classroom to its inhabitants (for 'f'). For more discussion of the values of unbound domain indices, cf. Stanley (2002b, p. 371).

⁸ Many authors have used bound readings of various constructions to draw disparate morals. Partee (1989) uses bound readings of relational expressions such as "local" and "enemy" to argue that binding is *not* always represented linguistically (thereby drawing the *opposite* conclusion from such data than Stanley (2000)). Cooper (1993) provides bound readings to argue for the semantic reality of situation variables. Von Stechow (1994) uses bound readings of quantifier phrases to argue that resource domain variables are indexical in nature, but stops just short of arguing that they are syntactically present (Ibid., p. 33). Nevertheless, it's natural to read Von Stechow as endorsing that thesis.

account for it. In Stanley (2002a, p. 152), for argument's sake, I employed this third interpretation in arguing against 'free enrichment' accounts of binding.

There are two basic kinds of challenges to the binding argument. First, there are attempts to argue that, whatever the right account of the data, the methodology behind the binding argument is unsound. However, the point of the third version of the binding argument is that merely objecting to the postulation of variables without providing an alternative account is insufficient. It is one thing to raise faults with the methodology, but quite another to provide an account that is equally adequate to the explanatory task. The second kind of response to the binding argument is to attempt to fulfill this obligation, by explaining bound readings without postulating covert structure.

Since the most important task for the person who objects to the binding argument is to explain bound readings without postulating covert structure, I will focus first on accounts that attempt to accomplish this. But before I begin my discussion of such accounts, I want to discuss briefly two approaches to the data that I will **not** discuss at length: variable free semantics, and free pragmatic enrichment.

A variable free semantic framework can provide an account of bound readings of sentences without postulating covert structure. There are many different versions of variable-free semantics, but I'll briefly focus on the elegant version given in Jacobson (1999). On Jacobson's account, work that might ordinarily be done by postulating syntactic movement or covert structure is done instead by type-shifting in the semantics; a pronoun in the complement of a verb induces a *type-shift* in that verb. A transitive verb has potentially different semantic types, depending upon the number of pronouns that occur within its complement. Complexity in the syntax, on a variable-free account, is replaced by complexity in semantic type assignments to lexical items.

I will not discuss variable-free semantics, because I think the question of whether to implement binding syntactically or semantically is orthogonal to the question at hand, which is whether certain examples demonstrate that intuitive truth-conditions are not generated within the semantics. Both the proponent of variable-free semantics and the more traditional syntactician and semanticist should agree that bound readings of a sentence are of semantic significance. The more traditional syntactician and semanticist should think they are of semantic significance because they indicate hidden syntactic

structure, whereas the variable-free semanticist should think they are of semantic significance because they demonstrate that a lexical item is associated with potentially distinct semantic types. Variable-free semantics does not make it easier to argue that certain readings cannot be generated in the semantics; it is irrelevant to this issue.⁹

The second topic I will not discuss is free enrichment accounts of the data. There are two styles of such accounts. According to the first, when one utters a sentence, via a pragmatic process, the sentence itself is 'enriched' into a longer sentence with the addition of lexical material. We may call this *free syntactic enrichment*. According to the second, when one utters a sentence, the semantic content of that sentence (a proposition or a property or propositional function) is enriched by the addition of additional semantic constituents. We may call this *free semantic enrichment*.¹⁰

I shall not discuss either enrichment account of the intuitive data, not because it is not topical (it clearly is), but because I have discussed such accounts in detail already (Stanley (2002a)). As I have previously indicated, my objection to such accounts is that they over-generate. If free pragmatic enrichment of either kind were a regular mechanism we could appeal to in communication, it would be a mystery why many sentences cannot serve greater communicative functions than they do.

In the next sections, I shall rather discuss two challenges to the binding argument that seek to account for the intuitive data without free pragmatic enrichment, either

⁹ An interesting issue arises with what would be captured as free readings of variables in a more traditional framework (Jacobson, 1999, pp. 134-5). On a variable-free framework, there really are no free variables. Explicit pronouns are semantically empty (express the identity function). The effect of a free variable (or a free reading of a relational expression such as "enemy") is to induce type-shifts so that the resulting sentence expresses a propositional function (e.g., in the case of "enemy", a function from persons to singular propositions). On this view, a sentence containing a free variable does not express a proposition, but rather a function from a certain kind of entity (determined by the type of the free-variable) to a proposition (or truth-value, depending on one's framework). Satisfying this function is not a matter of free enrichment, but rather closer to what Kent Bach (1994) calls "completion".

¹⁰ Some philosophers of language hold that a sentence expresses a structured semantic content, with specific holes that are saturated by context. I do not consider this to be free semantic enrichment. In the envisaged process, the role of context is constrained to supply elements of a particular semantic type. Thus, it is conventionally constrained. In contrast, free semantic enrichment is, by its nature, not so constrained. Elements of any semantic type, consistently with the conversational context, could be added.

syntactic or semantic. If there were viable alternative accounts of some of the binding data, then that would raise worries about the soundness of the underlying methodology, and thereby threaten to rob the defender of the binding argument of a useful tool by which to establish covert structure in a wide variety of problem cases.

Section IV. Binding and Comparative Adjectives

In Stanley (2000), the binding argument was used to argue for the syntactic representation of comparison classes for comparative adjectives, such as "tall" and "old". The target of the arguments there was the following kind of unarticulated constituent clause:

(R) Relative to any context *c*, "old" expresses the property of being old for a thing of the kind that is salient in context *c*.

In Stanley (Ibid., p. 418), I pointed out that (R) cannot capture the most natural reading of a sentence such as:

(26) Most species have members that are old.

The problem with (R) is that it predicts that the occurrence of "old" in (26) must be fixed to a particular species salient in the context of use of (26). But, in the natural reading of (26), the values introduced by the initial quantifier "most species" vary the comparison class to which "old" is applied. (R) cannot account for this reading.

In conversation, I have encountered philosophers challenging this line of argument, by contending that "old" in (26) simply means *old for a thing of its kind*. If so, then (R) produces the desired reading, because the variation is part of the lexical meaning of the adjective "old". However, this suggestion does not rescue (R), as a similar example shows:

(27) Every sports team has a member who is old.

Intuitively, (27) may express the proposition that every sports team has a member that is old for that sport. But on the view we are considering, "old" expresses the property of being old for x's kind; that is $_x(\text{old for } x\text{'s kind})$. But each member of a sports team belongs to many different kinds. So it is unclear how to use this suggestion to obtain this reading of (27).

Perhaps we can use this suggestion to emend (R):

(R*) Relative to any context c , "old" expresses the property $_x(x \text{ is old for } x\text{'s } N)$, where N is the contextually salient property.

Unlike (R), (R*) has no trouble with (26). For relative to a context of utterance for (26), the salient property is *species*. So, relative to a context of utterance of (26), "old" expresses $_x(\text{old for } x\text{'s species})$, which delivers the correct reading. But (R*) also promises to help with (27). Relative to a context of utterance of (27), the contextually salient property is *sport*. So, relative to a context of utterance of (27), "old" expresses the property $_x(\text{old for } x\text{'s sport})$. According to this clause, then, in a context of utterance of the appropriate sort, (27) expresses the proposition that every sports team has a member who is old for his sport. And this seems to be the desired reading.

However, (R*) does not work. It faces what we may call *the Bo Jackson problem*. Someone may play more than one sport. In such a case, (R*) will not deliver a result, since (R*) requires that there is one unique sport played by each person. Nevertheless, relative to such a situation, (27) may still express a coherent and indeed true proposition, namely the proposition that every sports team S has a member who is old for the sport played by S .

Here is a possible repair to (R*) in light of the Bo Jackson problem:

(R**) Relative to any context c , "old" expresses the property $_x(x \text{ is old for some } N \text{ in which } x \text{ participates})$, where N is the contextually salient property.

(R**) evades the Bo Jackson problem, since it does not require, of each thing, that it participates in only one kind of the contextually salient property (in the case of (26), only one kind of sport). However, (R**) also fails.

According to (R**), in a context of the appropriate sort, (26) expresses the proposition that every sports team has a member that is old for some sport he plays. Suppose that there are three sports teams, a gymnastics team, a chess club, and a baseball team. One person, Bob, plays for all three teams. Bob is old for a gymnast, but not old for a chess player or a baseball player. No other members of the teams are old for their sports. Intuitively, what an utterance of (27) expresses, relative to this situation, is false. However, according to (R**), the proposition expressed by (27) should be true in this situation. For each sports team does have a member who is old for some sport he plays. Each sports team contains Bob, who is old for a gymnast.

So there does not seem to be any easy repair of a rule such as (R). If one wishes to capture semantically all of the intuitive judgements about truth and falsity we have discussed, examples such as (26) and (27) seem to require a syntactically represented comparison class (or some other mechanism that imitates binding).

Section V. Quantifying over Contexts

One classic example of an unarticulated constituent analysis of a construction involves example (28), from Perry (1986):

(28) It's raining.

(U) "It is raining(t)" is true in a context c if and only if the denotation of "rain" takes $\langle t, l \rangle$ to the true, where l is the contextually salient location in c .

In Stanley (2000, pp. 415ff.), I used binding considerations against an unarticulated constituent analysis like this. In particular, I used examples such as:

(29) Every time John lights a cigarette, it rains.

The unarticulated constituent analysis suggested in (U) cannot derive the natural reading of (29), where the location of the raining varies with the values of the initial quantifier, "every time John lights a cigarette".

My purpose in giving this argument was not so much to advance my own account of such examples as to reject an unarticulated constituent analysis. But I did propose two positive "articulated" accounts of the data (Ibid., pp. 416-17). According to the first, "rain" is associated with an event or situation variable, which is bound by the initial quantifier "every time". According to the second, "rain" is associated with a pair of variable positions, one of which determines a time, and the other a location, both bound into by the initial quantifier.

Peter Pagin (this volume) seeks to evade the need for either kind of analysis, by treating quantifications such as "every time John lights a cigarette" as quantifiers over contexts. On Pagin's analysis, (29) ends up having the truth-conditions in (30):

(30) For every context c' differing from c at most in its time and location indices, 'if John lights a cigarette, then it rains' is true in c' .

On Pagin's treatment, there is no need for a variable for events or locations, because the initial quantifier is over contexts.

My concern with Pagin's analysis is a familiar one with operators that shift contextual features, noted originally by Lewis (1981, p. 86):

...we need to know what happens to the truth values of constituent sentences when one feature of context is shifted and the rest are held fixed. But features of context do not vary independently. No two contexts differ by only one feature. Shift one feature only, and the result of the shift is not a context at all.

Suppose I utter (29) in a context c . I am the speaker of c . But to obtain the right truth-conditions, we need to quantify over all n -tuples that differ from c only in their location and time indices. But some n -tuples of indices will not be possible contexts of use.

To make Lewis's point vivid, consider the sentences in (31)

- (31) a. Whenever I'm politely listening to someone speaking, it starts to rain.
b. Whenever wind blows through a mountain pass, it starts to rain.

Pagin's truth-conditions for the sentences in (31), considered as uttered in a context *c*, are:

- (32) a. For every context *c'* differing from *c* in at most its location and time indices, "If I'm politely listening to someone speaking, it rains" is true in *c'*.
b. For every context *c'* differing from *c* in at most its location and time indices, "If the wind blows through a mountain pass, it starts to rain" is true in *c'*.

Now consider the contexts *c'* involved in obtaining the correct truth-conditions. To obtain the right truth-conditions, many of these must be packages of indices that are not possible contexts of use. For example, to obtain the right truth-conditions for (31a) via (32a), we need contexts in which the speaker in *c* is the addressee in *c'*, rather than the speaker. But these will not be contexts that differ from *c* only in their location and time indices. A different problem surfaces for (31b). A context is one in which the agent of the context is at the time and location of the context. So Pagin predicts that (31b) is true just in case whenever the speaker in *c* is at the time and location of the relevant mountain pass, it starts to rain when the wind blows through. But clearly, these truth-conditions are too weak. (31b) would be falsified if there are situations with no one around (and hence no agents) in which the wind blows through a mountain pass and it doesn't start to rain at that location.¹¹

So, I'm skeptical that appealing to quantifiers over contexts will help in accounting for bound readings of alleged unarticulated constituents. The problem is that quantifying over contexts results in truth-conditions that are too weak, given the paucity of contexts of use.

¹¹ Pagin discusses a similar problem (see the discussion surrounding principle (I)). But his discussion cannot accommodate (31b), since his approach involves quantifying only over contexts, and to obtain the right truth-conditions for (31b), one needs to quantify over indices that are not possible contexts of use.

Section VI. The challenge from Over-Generation

In the past two sections, I have discussed attempts to capture the binding data without postulating unpronounced structure. I now turn to challenges to the binding argument. The most common sort of objection involves *over-generation*.¹² According to this kind of challenge, if one postulates variables when bound readings are available, what results is an over-generation of variables in the syntax.

In what follows, I will respond to the over-generation concern for the binding argument. But first, I want to note an oddity about the strategy of advancing over-generation objection against the binding arguments. Those who advance such objections

¹² One different type of objection to the binding argument, discussed in Cappelen and Lepore (2002, pp. 276-7) is that the variables the binding argument would have us postulate behave differently than overt pronouns, in particular in anaphora. For example, as they point out, it is odd to follow up an utterance of "Tigers are mammals", by "and it is a big domain" (with the "it" referring to the domain associated with "tigers"). This objection deserves more attention than I can give it here (unpublished work by Adam Sennett and Brett Sherman is important in this regard). But one reason to be suspicious of the argument is that it would apply to an alarmingly large range of constructions. Of course Cappelen and Lepore revel in this fact; as they note (Ibid., p. 279), these considerations would also tell against postulating variables for comparison classes for comparative adjectives (cf. also Cappelen and Lepore, this volume). Their argument would also tell against postulating variables for degrees for adjectival constructions, as witnessed by the oddity of "John is tall, and it is a high degree." More problematically, the considerations also entail that the implicit anaphoric elements associated with relational expressions such as "local" and "enemy" are not syntactically realized. For example, suppose Bill utters "John talked to an enemy in 2004", thereby expressing the proposition that John talked to an enemy of Bill in 2004. It is not possible for someone to follow this utterance up by saying "He has many enemies", where the "he" is genuinely *discourse anaphoric* on the covert variable that refers to Bill. Similarly, suppose Bill utters "John talked to an enemy", meaning an enemy of the United States. It is not possible to follow this up with, "And it is a big country", where "it" is *discourse anaphoric* on the covert variable. So, if this argument were correct, *implicit anaphora* would not be syntactically realized. Similarly, it is plausible to take epistemic modals to involve implicit anaphora; the occurrence of "might" in a token of "It might be raining in Paris on July 19, 2004" is to be taken relative to the knowledge state of the person making the utterance. But one cannot follow up someone's utterance of "It might be raining in Paris" with "He is strange", where "he" is an anaphoric pronoun (contrast this with the acceptability of following "According to John, it might be raining in Paris" with "He is strange", where "he" is uttered with anaphoric intent). The argument therefore proves too much, unless Cappelen and Lepore are also willing to use it to reject the syntactic representation of implicit anaphora.

typically do so in support of *pragmatic* accounts of the bound readings. But pragmatic accounts of the data, as I have emphasized, over-generate more than any other account possibly could. For pragmatic accounts are, by their nature, *unconstrained* by linguistic meaning. Were some pragmatic account to be correct, there would be numerous sentences that would allow readings that they actually do not allow (Stanley (2002a)).¹³

The first over-generation objection I will discuss is due to Cappelen and Lepore (2002, p. 273):

A confused mathematical anthropologist (call her 'Sally') trying to find out if mathematical truths are universal utters (5) as a summary of her findings:

(5) Everywhere I go, $2 + 2 = 4$

Here's the binding argument applied to (5):

Intuitively, (5) says that for every place Sally goes, $2 + 2 = 4$ at that place. So we should present the logical form of (5) along the following lines:

(5*) For all places x , if Sally goes to x , then $2 + 2 = 4$ at x .

The quantifier phrase 'Everywhere Sally goes' is binding a place variable in the logical form of ' $2 + 2 = 4$ ' -- otherwise, there would be nothing for the quantifier phrase to bind. This establishes that the logical form of the sentence ' $2 + 2 = 4$ ' has a freely occurring place variable.

¹³ Indeed, in the thousands of pages that have been written over the last decade arguing for pragmatic (non-semantic) accounts of a wide range of apparently semantic phenomena, I am not aware of a single attempt to provide a response to the threat of over-generation to pragmatic theories. Indeed, I am not even aware, aside from passing footnote references, to a discussion of the over-generation threat facing such theories. Given this silence, there is some irony involved in such theorists' extreme sensitivity to over-generation worries with alternative positions. I hope that the sensitivity such

Since there is obviously *no* variable ranging over locations in ' $2 + 2 = 4$ ', this is a *reductio* of the Binding Argument.

Before responding to this argument, we should be clear about what version of the Binding Argument Cappelen and Lepore have in mind. As I discussed above, there are three versions of the Binding Argument. According to the first, it has to do with grammaticality; one must postulate a place variable in the logical form of a sentence, or else what cannot explain the grammaticality of the larger construction. According to the second, it does not have to do with grammaticality; bound readings are taken to be a reflection of syntactic binding. According to the third reading, it is an inference to the best explanation of the bound reading.

Cappelen and Lepore, in their paper, address the first of these versions of the Binding Argument. Since I am not aware of that version being promoted in published work, I'm not sure it should be the focus of critical attention.¹⁴ That is, the obvious *grammaticality* of example (5) poses no worries for the advocate of the Binding Argument. So I take it that the feature of the example that is supposed to concern the advocate of the binding argument is the claim that the intuitive reading of the example is (5*). If the intuitive reading of (5) is (5*), then it would seem that the advocate of the binding argument is committed to postulating a place variable in the logical form of " $2 + 2 = 4$ ".

However, I do not see that (5*) is the intuitive reading of (5), and I do not see that Cappelen and Lepore even believe that (5*) is a reading of (5). As Cappelen and Lepore (Ibid.) point out, "...it is close to indisputable that arithmetical statements lack hidden indexicals referring to places." Presumably, the reason they are so convinced of this is that it is unclear what it even means to speak of an arithmetical statement being true *at a place*. If that is the intuition, then it is equally hard to see how (5*) is a legitimate reading of (5). Sally may intend (5*) as a reading of (5), because, as Cappelen and Lepore put it,

theorists evince to over-generation objections will soon be reflected in greater attention to these worries with their favored accounts.

¹⁴ Cappelen and Lepore cite Nelson (2001), who seems to have the first interpretation of the Binding Argument in mind. But the relevant passages in Nelson (2001) involve a

she is a "confused mathematical anthropologist." But the fact that someone confusedly believes a sentence has a certain reading does not give that sentence that reading.

Of course, it is uncontroversial that (5) is grammatical. But nothing follows from this, other than a rejection of the first version of the Binding Argument. But since the first version of binding argument is not one that has ever been advanced in print, it is not germane to the issue.

Brehehy (2004) has leveled another sort of over-generation objection against the Binding Argument, this involving what he calls the problem of multiple dependencies. Since Brehehy's arguments are interesting and illustrate important points, it is worth going over them in detail.

Recall that on my favored account of domain restriction, motivated by binding considerations, each noun is syntactically associated with two indices, a function index and an argument index. Relative to a context, the function index is assigned a function from objects to properties, and the argument index an object. So, if I have New Jersey in mind when I say "Every politician is saintly", then New Jersey is the value of the object index, and perhaps the function index is assigned a function from states to the property of being an inhabitant of that State. It is this view that is Brehehy's target.¹⁵

The first sort of example Brehehy discusses is:

(33) Every student was feeling particularly lucky and thought no examiner would notice every mistake.

summary and subsequent critique of Stanley (2000), and in that paper, I certainly did not have the first interpretation of the Binding Argument in mind.

¹⁵ In previous work (Stanley and Szabo (2000), Stanley (2002b)) I took the function index and the object index to 'co-habit' a node with the head noun. I no longer think this is correct. Talk of 'cohabiting a node' with a lexical item suggests that domain indices are part of lexical structure. But the position that domain indices are merely part of lexical structure is not consistent with the general view (advocated in Stanley (2000)) that considerations such as binding, weak-crossover, and ellipsis are evidence for genuine syntactic structure, rather than mere lexical structure. Secondly, both Brehehy (Ibid.) and Timothy Williamson (2004) provide evidence that the domain index must sometimes be outside the scope of adjectives modifying the head noun. Since the domain indices must sometimes be inside the scope of modifying adjectives (cf. Stanley (2002b, pp. 372-3)), this suggests that the domain indices occupy their own terminal nodes that can have different adjunction sites.

Brehehy argues that (33) has the reading in:

(34) [Every student]_x thought [no examiner]_y would notice [every mistake made on a paper x turned in and y examines]_z.

Brehehy argues that this sort of example is problematic, because the approach I advocate cannot generate reading (34) of (33), without postulating more syntactic variables. That is, Brehehy worries that the methodology of the Binding Argument commits its proponent to postulating implausibly many variables, certainly more than just the domain variables discussed above.

But I don't see the worry with this particular example. The noun "examiner" is a relational noun. It is associated (cf. Stanley (2000)) with a syntactically represented index. This fact straightforwardly generates the desired reading. That is, an independently motivated claim about relational expressions, together with an independently motivated theory of domain restriction, straightforwardly predicts that (33) has the reading:

(35) [Every student]_x thought [no examiner, x]_y would notice [every mistake $f(x)$]_z.

We may assume 'f' is assigned a function from students to their exam questions. So, we can straightforwardly predict a reading of (33) according to which every student thought no examiner of that student would notice every mistake on that student's exam. So recognizing the relational nature of 'examiner' enables us to derive the natural reading with only independently motivated resources.

Brehehy also has examples that do not involve relational nouns. However, they also do not raise problems for any view I have defended. Consider Brehehy's "multiple dependence" example (36a), which he claims to have reading (36b):

- (36) a. Every paranoid artist thinks no dealer will stop at selling every forged painting.
b. [Every paranoid artist]_x thinks [no dealer]_y will stop at selling every [[forged [painting by x]] coming into y's possession.

On my account, every noun (or N') in every QNP is associated with a domain index. So, (36a) is predicted to have a reading as in (37):

(37) [Every paranoid artist]_x thinks [no dealer $f(x)$]_y will stop at selling [every forged painting $f'(y)$].

Relative to the envisaged context, 'f' is assigned a function from artists to people in possession of forged paintings of that artist, and 'f'' is assigned a function from dealers to their collections. So, the theory predicts that (36a) has the reading:

(38) Every paranoid artist thinks that no dealer in possession of forged paintings of that artist will stop at selling every forged painting in their collection.

Intuitions are subtle here. But both of my informants think that (36a) clearly has reading (38), and none of my informants think that (36a) has reading (36b). In other words, there is a strong difference among my informants between (36a) and (39):

(36) a. Every paranoid artist thinks no dealer will stop at selling every forged painting.

(39) [Every paranoid artist]_x thinks no dealer will stop at selling every forged painting of his_x.

My informants do not obtain reading (39) of (36a). They only obtain reading (38) of (36a). And that is precisely what an account that only postulates a single pair of domain variables (one argument and one function variable) would suggest.

A final example of Breheny is:

(40) Every company knows that none of the pension fund can be diverted away from any former employee.

Breheny claims this example has the following reading:

(41) [Every company]_x knows that [none of the pension fund]_y can be diverted away from any [[former employee of x] who is due some of y].

But again, no additional variables are needed to capture the appropriate reading semantically. Since "employee" is a relational noun, (40) is predicted to have the reading in (42):

(42) [Every company]_x knows that [none of the pension fund $f(x)$]_y can be diverted away from any former [employee, x].

Relative to the envisaged context, 'f' is assigned a function from (e.g.) companies to their benefit programs. So, (41) expresses the proposition that every company knows that none of the pension funds of that company can be diverted away from any former employee of that company. And this is precisely the desired reading.

Breheeny writes that "if we wanted to pursue a variable-rich approach, given these multiple dependencies, we seem to need to assume that QNP structures contain a plethora of hidden variables at different levels which are vacuously assigned (to what?) when not used." But, as I have argued, this worry is unfounded. The assumption that relational nouns are associated with syntactically realized implicit arguments, together with the assumption that QNPs are each associated with a single domain index, is sufficient to explain all the data. Indeed, the framework even explains why certain sentences *lack* readings, as in reading (36b) of (36a).

There is a final sort of over-generation objection against the Binding Argument that I wish to discuss. Because I think this sort of objection hinges on a confusion about the methodology of semantics, responding to it involves less focus on empirical detail, and more on foundational matters.

According to Francois Recanati, the Binding Argument involves what he calls the "Binding Criterion". The Binding Criterion is that intuitively bound readings must be reflected by bound variables in the syntax. Given his assumption, here is his objection (Recanati (2004), p. 106-7):

We can say:

(15) John is anorexic, but whenever his father cooks mushrooms, he eats.

On a natural interpretation, we understand that John eats the mushrooms his father has cooked. Intuitively, a form of binding is operative here; for the food eaten by John co-varies with the food cooked by his father. Such examples show that intuitive binding, per se, does not entail the existence of a free variable in logical form. The Binding Criterion, on which Stanley's argument rests, must be rejected.

Recanati's argument against the Binding Criterion has two premises. The first premise is that the intuitive reading of his sentence (15) is that whenever John's father cooks mushrooms, John eats the mushrooms his father cooks. Recanati's second premise is that there is no covert variable for *what John eats* in the logical form of "he eats". Assuming these two premises, Recanati concludes that the Binding Criterion, which requires such a covert variable, must be rejected.

Recanati's example is deliberately modeled upon examples discussed in Stanley (2000), such as:

(43) Whenever John lights a cigarette, it rains.

There, I concluded that "it rains" does contain a variable that can be bound by the initial quantifier, "whenever John lights a cigarette." The evidence that it can be is due to the fact that the intuitive reading of (43) is:

(44) Whenever John lights a cigarette, it rains at the location of that lighting.

I also suggested various ways of accounting for the bound reading in question (Ibid., p. 416). Recanati's suggestion is that his example (15) is analogous to (43), so that whatever methodology leads one to postulate an unpronounced variable in "it rains" to account for

the bound reading in (43) should lead one, incorrectly, to postulate an unpronounced variable for what is eaten in "John eats."

However, Recanati's (15) is simply not analogous to (43). In particular, the first premise of Recanati's argument is false (whereas the corresponding premise for (43) is true). As Luisa Marti (ms.) has shown, it is not the case that the intuitive truth-conditions of (15) are what Recanati says they are. The following two discourses, due essentially to Marti, reveal the clear difference between (15) and (43):

- (45) A. Whenever John's father cooks mushrooms, John eats.
B. #No he doesn't -- he eats broccoli when his father cooks mushrooms.
- (46) A. Whenever John lights a cigarette, it rains.
B. No it doesn't -- though it rains somewhere else.

As the oddity of Marti's discourse in (45) clearly demonstrates, it is not permissible to deny the content of A's assertion on the grounds that one thinks that John eats broccoli, rather than mushrooms, when his father cooks mushrooms. This demonstrates that the intuitive truth-conditions of an utterance of Recanati's (15) is not that whenever John's father cooks mushrooms, John eats the mushrooms his father cooks. It is rather that whenever John's father cooks mushrooms, John eats *something*. In contrast, the acceptability of the discourse in (46) demonstrates that the intuitive content of (43) is (44).

According to the binding argument, when a bound reading is part of the intuitive truth-conditions of an utterance, it is the result of a quantifier-variable interaction. The problem with Recanati's example is that (as Marti's argument demonstrates), it is no part of the intuitive truth-conditions of an utterance of (15) that John ate the mushrooms his father cooked. Our intuitions about truth and falsity clearly reveal this to be a reflection of our background assumptions, combined with the semantic content of (15), which is just that in whatever situation John's father cooked mushrooms, John ate something in that situation. In contrast, our intuitions about the truth-conditions of utterances of

sentences such as (43) clearly reveal that relativity to a location parameter is part of their intuitive truth-conditions.

The distinction between a verb like "rains" and a verb like "eats" can be seen even in non-embedded sentences. Suppose Bill has cooked a mushroom dinner. Pointing at a dirty plate on the table, and intending to communicate that John has eaten the mushrooms Bill ate, I utter:

(47) John ate.

Suppose one knew that John had just eaten, but he did not eat the mushrooms Bill cooked. It is still clearly not permissible to follow my assertion with:

(48) No he didn't; he ate broccoli instead.

In contrast, suppose that it is raining in New York City, where I am located. Speaking on the phone with Delia, who is in Ithaca, I utter:

(49) It's raining.

Delia, seeing on television that the sky is clear in New York, utters:

(50) No it isn't. But it's raining here.

Delia's reply is perfectly acceptable, in contrast to the unacceptable (48). This demonstrates that the location is part of the intuitive truth-conditions of an utterance of (49), whereas what is eaten, even when it is salient, is clearly not part of the intuitive truth-conditions of an utterance of (47).

Of course, when we hear an utterance of Recanati's (15), we are liable to assume that John ate the mushrooms his father cooked. The reason we assume this has nothing whatever to do with semantics. Rather, we assume that when your father is cooking a meal, and you're eating together with him, it is expected behavior to eat what your father

has cooked. We also assume, when we hear an utterance of Recanati's (15), that John's father is not a Martian. But it would be absurd to suppose that it is the semantics that "tells" us that John's father is not a Martian. It is similarly absurd to suppose that information we acquire via the background assumption that people generally eat the meals that are cooked for them must be supplied semantically, if the semanticist is right.¹⁶

Similar mistakes to these have been made by other pragmaticists, in their discussions of the view of Stanley (2000) that the source of the intuitive truth-conditions of an utterance are covert structure. For example, Wilson and Sperber (2002) exploit the example:

(51) I must wash my hands: I've eaten [using my hands, rather than, say, being spoon-fed].

According to Wilson and Sperber, the bracketed material is part of the intuitive truth-conditions of an utterance of "I must wash my hands: I've eaten." So, they conclude that anyone who defends the clear and elegant explanation of the source of intuitive truth-conditions is committed to the view that "eat" has an argument place for *the manner of eating* (see also Carston (2002, pp. 203-4)). But it is false that the manner of eating, even in the example they envisage, enters into the intuitive truth-conditions, as the oddity of Bill's utterance in the Marti-style discourse in (52) demonstrates:

(52) John: I must wash my hands: I've eaten.
Bill: No you didn't; you got spoon-fed.

¹⁶ Iliria Frana (ms.) provides an argument against the binding argument that is also undermined by these considerations. Frana considers the sentence "Paolo is a real curious guy; every time he finds something, he opens it." According to Frana, an utterance of this sentence has the intuitive truth-conditions that Paolo always opens the thing he finds in a manner appropriate to that thing. But intuitively, what is said would not be *false* if Paolo always opened what he found in a manner that was not appropriate to that thing. Therefore, this is no part of the intuitive truth-conditions of Frana's sentence. A similar point dispenses with the reply to the binding argument in Stalnaker (2004, pp. 110-111).

Of course, we would naturally assume, given John's discourse, that he ate with his hands. We would also assume that he was not from Mars, and that he was not the product of in-vitro fertilization. It is just as absurd to take it as an objection to the semanticist that such information is not provided by the semantics of the verb "eat" as it is to take it as an objection to the semanticist that the semantics does not provide for a manner of eating. Such information has nothing to do with intuitive truth-conditions of an utterance.

The responsibility of the semanticist is rather to show that speaker intuitions about the truth-conditions of an utterance are due to semantics. Unfortunately, too many objections to the semanticist assume that the responsibility of the semanticist is to generate within the semantics all information that a competent speaker and member of a culture may derive from a communicative act. Such objections seem to presume, absurdly, that the semanticist's position is incompatible with Grice. As Marti's tests show, being more subtle about judgements of truth and falsity can clearly reveal the distinction between what is part of intuitive truth-conditions proper and what is conveyed by the communicative act to a hearer who combines these truth-conditions with her background knowledge about the world.

Conclusion.

My purpose in this paper has been to defend the claim that the intuitive truth-conditions of an utterance are due to semantic interpretation. Many of those who have objected to it have done so by saddling the position with absurd theoretical commitments, such as the position that all information conveyed in any discourse is due to the semantics. Part of my goal in this paper has been to explain what costs the semanticist's position incurs, and what costs it doesn't, by elucidating the target concept of semantics. As I have argued, in certain cases (e.g. that of deferred reference), the semanticist must make decisions about the *defiendum* that are informed by theoretical considerations. But this is the ordinary practice even in the human sciences.

Bibliography

- Bach, Kent (1994): "Conversational Implicature", *Mind and Language* **9**: 124-62.
- Breheny, Richard (2004): "A Lexical Account of Implicit (Bound) Contextual Dependence", in *Semantics and Linguistic Theory* (SALT) **13**.
- Camp, Elizabeth (ms.) "Lodging in the Words"
- Cappelen, Herman and Lepore, Ernie (2002): "Indexicality, Binding, Anaphora, and A Priori Truth", *Analysis* **62.4**: 271-81.
- Cappelen, Herman and Lepore, Ernie (xxxx): "A Tall Tale in Defense of Semantic Minimalism and Speech Act Pluralism". This Volume.
- Carston, Robyn (2002): Thoughts and Utterances: The Pragmatics of Explicit Communication. Oxford: Blackwell.
- Cooper, Robin (1993): "Generalized Quantifiers and Resource Situations", Aczel, Israel, Katagiri, and Peters (ed.), Situation Theory and its Applications. Stanford: Stanford University Press.
- Von Fintel, Kai (1994): Restrictions on Quantifier Domains. University of Massachusetts Dissertation.
- Frana, Ilaria (ms.) "The Constituents of Meaning"
- Jacobson, Pauline (1999): "Towards a Variable-Free Semantics", in *Linguistics and Philosophy* **22**: 117-184.
- Levinson, Stephen (2000): Presumptive Meanings: The Theory of Generalized Conversational Implicature. Cambridge: MIT Press.
- Lewis, David (1981): Index, Context, and Content. In Philosophy and Grammar: Papers on the Occasion of the Quincentennial of Uppsala University. Edited by Stig Kanger and Sven Ohman (D. Reidel, Dordrecht): 79-100.
- Marti, Luisa (ms.): "Unarticulated Constituents Revisited".
- Nelson, Michael (2001): "When is an Expression Context-Sensitive?", unpublished ms.
- Pagin, Peter (xxxx): Compositionality and Context. This volume.
- Partee, Barbara (1973): "Some Structural Analogies Between Pronouns and Tense", *Journal of Philosophy*.
- Partee, Barbara (1989): "Binding Implicit Variables in Quantified Contexts", *Proceedings of the Chicago Linguistics Society* **25**, University of Chicago Press, Chicago: 342-65.
- Perry, John (1986): "Thought without Representation", in Supplementary Proceedings of the Aristotelean Society **60**: 137-52.
- Recanati, Francois (2004): Literal Meaning. Cambridge: Cambridge University Press.
- Sag, Ivan (1981): Formal Semantics and Extra-Linguistic Context. In Peter Cole (ed.), *Syntax and Semantics* **13**: Radical Pragmatics. New York, Academic Press. 273-294.
- Stalnaker, Robert (2004): "Comments on 'From Contextualism to Contrastivism'", *Philosophical Studies* **119 1 & 2**: 105-117.
- Stanley, Jason (2000): "Context and Logical Form", *Linguistics and Philosophy* **23.4**: 391-434.
- Stanley, Jason (2002): "Making it Articulated", *Mind and Language* **17 1&2**: 149-168.
- Stanley, Jason and Szabo, Zoltan (2000): "On Quantifier Domain Restriction", *Mind and Language* **15.2**: 219-261.
- Stern, Josef (2000): Metaphor in Context. Cambridge: MIT Press.

Williamson, Timothy (2004): "Everything", in *Philosophical Perspectives* 26:
Philosophical Linguistics and Philosophy of Language.