

Context Dependence and Compositionality[□]

Francis Jeffrey Pelletier

Abstract: Some utterances of sentences such as ‘Every student failed the midterm exam’ and ‘There is no beer’ are widely held to be true in a conversation despite the facts that not every student in the world failed the midterm exam and that there is, in fact, some beer somewhere. For instance, the speaker might be talking about some particular course, or about his refrigerator. Stanley and Szabó (in *Mind and Language* v. 15, 2000) consider many different approaches to how contextual information might give meaning to these ‘restricted quantifier domains’, and find all of them but one wanting. The present paper argues that their considerations against one of these other theories, considerations that turn on notions of compositionality, are incorrect.

Word Count (including footnotes): 6162

[□] Thanks to Greg Carlson, Bernard Linsky, Paul Rusnock, and Richard Vallée for discussions about Stanley and Szabo’s article generally and for comments on the topic of this paper in particular. Ray Elugardo, Ernie LePore, and Rob Stainton have also read earlier versions of this paper and given sage advice. Special thanks go to Jason Stanley and Zoltan Szabó for their comments on an earlier version of this response. I’m sure they still don’t believe me—but then, that makes us even. Three anonymous referees have also made many helpful suggestions.

Address for correspondence: Department of Philosophy, University of Alberta, Edmonton, Alberta, Canada T6G 2E5.

Email: jeffp@cs.ualberta.ca

1. Introduction

Stanley & Szabó (2000)—henceforth S&S—consider the problem of quantifier domain restriction, as indicated by sentences like (1), which can be true in a conversation despite the fact that not every student in the world failed the midterm exam.

(1) Every student failed the midterm exam.

For instance, the speaker might simply be talking about every student in some particular course, or, maybe, every one of his/her students as opposed to a colleague's students. There is another sentence that could, perhaps, be an 'expanded' version of (1) and which might capture what the speaker is talking about:

(2) Every student in my elementary logic course this term failed the midterm exam.

S&S view quantifier domain restriction as a special instance of the problem of context dependence. In the present case, the interpretation of the quantified NP 'every student' is somehow restricted to a certain contextually specified group. (It might be also noted that these sentences manifest another instance of the general problem of context dependence, for the interpretation of the definite NP 'the midterm exam' is also in some way specified by context).

S&S exhaustively consider many different approaches to the problem of how to use contextual information to give the meaning of restricted quantifier domains, and find all but one of them wanting. This is sure to stir theorists to resist these arguments against their favourite theories; and indeed, the very issue of *Mind & Language* in which S&S appeared contains commentaries by Stephen Neale and by Kent Bach, each trying to deflect S&S's argumentation away from their favourite theory.

I wish to do the same for my favourite theory of how to characterize the role of context in restricting quantifier domains. I do not aim at giving a thorough account of my theory and defending it from yet further attack. Rather, I wish to show that the one specific criticism leveled against it by S&S is ‘just plain wrong’, and so my theory emerges as a viable contender even after the scorched earth methodology employed in S&S. In addition, I will close by appealing to some pretheoretic intuitions we have about ‘context dependency,’ with an eye to urging my favourite account in preference to that of S&S.

2. Accounts of Quantifier Domain Restriction

S&S favour a ‘semantic account’ for explaining the quantifier domain restriction phenomena. For the exact details of what this means and how it might differ from syntactic accounts and from pragmatic accounts, the reader is directed to S&S. Within the various ways of giving a semantic account, S&S distinguish ‘metalinguistic’ from other treatments, rejecting the various ‘metalinguistic’ ones. They are thus led to posit variables directly in the logical form (this is what they mean by ‘non-metalinguistic semantic accounts’); these variables are used to represent the information that is provided by the context. They consider sentences in which the contextual information relevant to an embedded quantified NP varies according to which value of a wider-scoped embedding quantified NP is chosen, such as

(3a) In every room in John’s house, every bottle is in the corner.

(3b) Whatever John does, most of the class falls asleep.

In such sentences, the quantifier domain relevant to the embedded NP ('every bottle', 'most of the class') varies according to the choice of a wider-scoped NP ('every room in John's house' or 'whatever John does'). They say (p. 251)

What the phenomenon of quantified contexts shows is that, assuming standard treatments of binding, quantifier domain restriction must be treated with the use of variables... We must associate variables with some part of quantifier expressions, such as 'every bottle'. Distinct non-metalinguistic [semantic] treatments of the problem of quantifier domain restriction correspond to different ways of associating variables with quantifier expressions.


S&S present four ways to do this 'association'. They champion one of these ways and present arguments against the others. My purpose here is to argue that their considerations against *my* favourite way—which is one of these other three—are incorrect.

In S&S's favoured theory, a 'contextual variable' is associated with each common noun. Thus, for example, with the common noun 'student' there might be associated $f(i)$, where the value of i is an object provided by the context and the value of f is a function also provided by the context that maps objects onto quantifier domains. So, the restriction on the quantified NP 'every student' in (1), relative to a context, is then provided by the result of applying the f -function (as supplied by the context) to the object i (as supplied by the context).


S&S's idea is that the semantic interpretation of a common noun like 'student', as relativized to a context c , is

$$(4) \quad \llbracket \langle \text{student}, f(i) \rangle \rrbracket_c = \llbracket \text{student} \rrbracket \cap \{x: x \in c(f)(c(i))\}$$

Considered as a lexical item then, ‘student’ has an interpretation—perhaps the set of all students, perhaps a function on possible worlds to the students in that world— but when it is to be used in a sentence, this immediately invokes the context variable(s). S&S in fact make this invocation be the terminal symbol of their parse tree (see their Fig. 2, p. 251):



My favourite analysis of the matter would associate the context material not with a lexical item or a terminal leaf of a syntactic analysis, but rather with the entire quantified NP. Thus, I would recommend a framework like this:



In this framework there is a domain restriction on the quantified NP, but not on the common noun (nor on the determiner, nor is there a separate node for $f(i)$).

I shall not say anything about the specific items S&S use in their representation of the contextual material, viz., f and i , for I take it that these could easily be changed if it were decided that more or different information should be in the context. Instead my discussion concerns *where* and *how* this information—whatever it may be—should come into the analysis of a sentence.

According to S&S, the sorts of theories we have been outlining here—with variables in the syntactic structure—are ‘semantic’ theories because they give voice to the intuition that sentence (1) expresses a proposition and indeed that it can express the *same* proposition as the unabbreviated (2), given the appropriate context. (See their discussion on pp. 229-233). As they recognize, this is not the only notion of ‘semantic’ that is in the literature (‘...there are other uses of “semantic”, with which our use should not be confused’ pp. 229-230). There is much that could be said about their notion of ‘semantics’ and how it fits with other conceptions, and whether it is legitimate to eschew ‘intermediate representation languages’ at the cost of putting semantic markers in the syntactic representation. Some of these topics are taken up by Bach and Neale, in their commentaries on S&S. But this is not my present concern.

3. S&S’s Argument

S&S think that my favourite proposal is non-compositional. They think one cannot generate the appropriate meaning of NP_i in (6) as a function of the meanings of Det and N. Of course, in order really to prove that this cannot be done, one needs to know some meanings for NP_i , Det and N, and then show that there is no function that fulfills the task. But this is not the direction that S&S take in discussing the issue. Instead, in explaining my favourite proposal, they say only this (p. 255, renumbering their example):

Let ‘ $\llbracket \alpha \rrbracket_{M,c}$ ’ denote the denotation of α with respect to the model M and context c. As is standard, verb phrase denotations and noun denotations are functions from objects to truth-values; noun phrase denotations are functions from such functions to truth-values; and determiner denotations

are functions from functions from objects to truth values to noun phrase denotations...

$$(7) \llbracket_{\text{NP}_i} \text{Det N} \rrbracket_{M,c} = \llbracket \text{Det} \rrbracket_{M,c}(\llbracket \text{N} \rrbracket_{M,c} \cap c(i))$$

It is the composition rule (7) that we find objectionable. What it does is apply the function denoted by the determiner to the result of intersecting the denotation of the noun with the value of the contextual variable. Our worry with composition rules of this sort is that they violate *compositionality*.

And just why, according to S&S, does (7) violate compositionality?

According to (7), the denotation of a noun phrase consisting of a determiner and a noun is determined by more than just applying the function expressed by the determiner to the denotation of the noun. Rather, in different contexts, the denotation of the noun phrase is a different function of the determiner and the denotation of the head noun. Therefore, (7) violates compositionality.

In contrast, they say, in their own proposal the contextual information is already part of the noun information, and so their version of (7) is just

$$(8) \llbracket_{\text{NP}_i} \text{Det N} \rrbracket_{M,c} = \llbracket \text{Det} \rrbracket_{M,c}(\llbracket \text{N} \rrbracket_{M,c})$$

One might usefully question why S&S think that a rule like (4) above, repeated here as (9) (see their p. 253), which interprets *nouns* as being modified by context *is* compositional while a rule like (7), which treats only NP's as being modified by context, is *not* compositional.

$$(9) \llbracket \langle \text{student}, f(i) \rangle \rrbracket_c = \llbracket \text{student} \rrbracket \cap \{x: x \in c(f)(c(i))\}$$

After all, isn't the meaning of some item in the syntactic tree being augmented by contextual information that is not present in any syntactic subnode of that node just as much in their proposal as it is in mine? Isn't it the case in their proposal that 'in different contexts the denotation of the noun is different'? Why should S&S think their Noun-interpretation rule to be compositional while the parallel NP-interpretation rule is not? Is it merely because the one case (theirs) involves a lexical item but the other doesn't? Perhaps that is the point S&S are making.

They say (p. 256):

Rules such as [(7)] allow the composition rules associated with a single linguistic structure to vary freely as a function of *extra-linguistic* context. An indefinite number of composition rules that vary not just according to linguistic context, but also according to extra-linguistic context, seems in tension with learnability considerations. If contextual variables were allowed to occur in non-terminal nodes, then the semantic theories required to interpret the resulting structures would violate compositionality in this quite drastic manner. Therefore, the distribution of contextual variables should be restricted to terminal nodes.

I myself don't see the force of this 'argument'. It is just not true that 'an indefinite number of composition rules that vary ... according to extra-linguistic context' are called for in this proposal but are not called for in S&S's proposal. S&S's proposal, as I have explained it, can perhaps be implemented in two different ways. In the one way, individual words of the lexicon are ambiguous amongst all the possible different meanings that could be generated by different domain restrictions. 'Student', for example, would be entered in the lexicon many, many different times: 'student₁', ... 'student₃₀₉₄' and the like. Although this proposal would be quite

clear in its effects, especially for compositionality (it is obviously a compositional account), S&S correctly rule it out of consideration on learnability-of-the-lexicon consequences (p. 232fn16). The other way, the one actually adopted by S&S, has words in the lexicon bear their ‘regular meanings’ but when they are employed in a sentence as a common noun, they are entered with the $f(i)$ variable which will accept values of the context. But if their proposal allows for such a variable to be entered for every common noun, there seems to be no reason to prohibit the very same variable to be entered for noun phrases. Both proposals seem equally committed or uncommitted to compositionality.

But I will not pursue this line of argumentation further, for the fact of the matter is that the proffered NP-interpretation rule (7) *is* compositional.

4. Compositionality

The Principle of Compositionality (PoC) is most commonly stated as:

The meaning of a (syntactically) complex expression is a function of the meanings of its parts and its (syntactic) mode of combination (i.e., the ‘way those parts are put together’)

The point of adding the clause about ‘mode of combination’ is to acknowledge that sentence-pairs like the following have different meanings despite their employing the same ‘parts’.

(10a) Cats eat mice.

(10b) Mice eat cats.

(11a) the average fat American

(11b) the fat average American

(12a) The man in the park seen with binoculars

(12b) The man in the park seen with binoculars¹

As we see, the syntactic manner of combining plays an important role in the construction of the meaning of a syntactic whole. There is no particular ‘proportion of the meaning’ of the whole that is due to the meanings of the parts versus the effect of the rule of combination. The same rule of combination can take the same meaning-of-parts but ‘put them together differently’ to yield different meanings, as in (10). Or the same meanings-of-parts could be ‘put together’ by different rules, as in (11), where *fat* is an attributive adjective in (a) but an appositive adjective in (b). Or the same meanings-of-parts could be ‘put together’ differently, but in a way that is not ‘visible’ in the string itself, as in (12). In all of these cases the difference in meaning between (a) and (b) is attributed to the effect of the syntactic rule of combination. Yet no one would want to claim that these constructions exhibit non-compositional behaviour. And the reason for this is clear: the syntactic rule that describes how the parts are combined can itself contribute a semantic effect to the meaning of the whole. Compositionality allows (or even, requires) that there is more to the meaning of the whole than merely the meanings of the parts.

The semantic effect that is attributable to the ‘manner of composition’ can be to impart material into the meaning of the whole which is not at all ‘in the meanings of the parts’ of this whole. Consider even simple adjective-noun combinations that form common noun phrases:

(13) $[[_{CN}Adj N]](x) = [[Adj]](x) \& [[N]](x)$

Here the meaning of the complex Common Noun Phrase is claimed to be the conjunction of the meanings of the parts. It is clear that there is no violation of compositionality here, despite the fact that the ‘&’ does not ‘occur in the meaning’ of either part. Instead, this conjunction is the

¹ A structural ambiguity employing different attachments of *with binoculars*, giving different meanings from the same parts.

contribution of the syntactic rule that combines adjectives with nouns to form CN-phrases. Consider also mass terms. Many theorists believe that a semantic feature of mass nouns is divisibility...that the mass noun is true of arbitrary parts of what it is true of. Thus *beer*, for instance, is not only true of the whole barrel of stuff before me, but also is true of each half of it, each glassful of it, and indeed each thimbleful of it. Yet certain syntactic constructions that contain *beer* not only lack this feature but in fact have contrary semantic features which are not in any of the parts of these constructions, as for example in

(14a) This is *a tasty beer*

(14b) There are *two beers* on the table

and so forth. There is no divisibility here, and indeed the italicized NPs are true of discrete objects.² And this semantic feature of ‘countability’ is added by the rule that puts together the relevant parts in the two cases.

Throughout the development of compositional theories of semantics, starting from Montague Grammars and continuing in the most recent pages of the various journals of linguistic semantics, it has always been recognized that there are semantic effects of whatever syntactic rule is being employed to form the complex expression. And no one has complained, when looking at the results from these works, that the suggested theories are non-compositional because ‘there are items in the meaning of the whole that are not in the meanings of the parts.’ For, to repeat, the ‘extra added items’ are due to the effect of the mode of combination. It is simply false that such constructions as (10)-(14) are inherently non-compositional. To be compositional is to obey PoC; and in these cases PoC is maintained by according an appropriate

² There might be disagreement as to what the objects are—portions, standard servings, kinds, amounts, ordinary objects, etc. But they are objects on all accounts.

semantic effect to the modes of combination that give rise to (10)-(14), thereby incorporating the ‘and the way the parts are put together’ clause of the PoC.

Objections concerning the compositionality of a node in a ‘logical form tree’ merely on the grounds that it contains information that is not a part of any of the subnodes is therefore misguided, unless it can be shown in addition that this information is not due to the manner in which the parts are combined. Or to put the point somewhat differently: *so long as there is a function* which takes semantic information of the sort possessed by the parts and generates the semantic value of the whole, then this account does not violate compositionality. (It may be an implausible theory for other reasons. But it is nonetheless compositional.)

Let us now return to (7), (8) and a generalization of (9), repeated here

$$(7) \quad \llbracket_{\text{NP}_i} \text{Det N} \rrbracket_{\text{M},c} = \llbracket \text{Det} \rrbracket_{\text{M},c} (\llbracket \text{N} \rrbracket_{\text{M},c} \cap c(i))$$

$$(8) \quad \llbracket_{\text{NP}_i} \text{Det N} \rrbracket_{\text{M},c} = \llbracket \text{Det} \rrbracket_{\text{M},c} (\llbracket \text{N} \rrbracket_{\text{M},c})$$

$$(9^*) \quad \llbracket \langle \text{N}, f(i) \rangle \rrbracket_c = \llbracket \text{N} \rrbracket \cap \{x: x \in c(f)(c(i))\}$$

and to S&S’s objection that (7), as opposed to the combination of (8) and (9*), is not compositional. Recall that what is required by compositionality is merely that the meaning of each syntactically complex item be *some* function of the meanings of its (syntactic) parts. Surely the function called for by (7) that takes a Det-meaning (in a model and in a context) and an N-meaning (in that same model and context) can yield an NP-meaning which is just exactly the same thing that they generate using (8) and (9*). The difference is merely that they inserted ‘ $f(i)$ ’ at the terminal node by employing (9*), and allowed it to ‘percolate up’ to the NP node using (8), whereas I recommend that it not be at the terminal node but instead be added when constructing the NP-meaning. The claim that $f(i)$ is not ‘in the subparts’ is out of place, as I have been at pains to explain. In the case of (7), the relevant $c(i)$ that gets employed when constructing the

NP-meaning would come from the fact that NP_i is being interpreted with respect to context c —we are evaluating $\llbracket_{NP_i} \text{Det N} \rrbracket_{M,c}$ after all, and c is therefore available for use in defining the relevant function. *Of course* different c 's will make the NP_i denote different functions. What else would anyone expect? That's precisely the role of the subscripted c , and is what it does in their (9*), so it could do this in (7) as well. If the point of a *semantic* interpretation of NPs with respect to a context is to give rise to distinct *semantic* values in different contexts, then my favourite theory, as given by (7), precisely and exactly fits the bill—and it does so in a compositional manner. That is, given a context c , the meaning in that context of any syntactically complex item is a function of the meanings in that context of that item's syntactic parts and the syntactic manner in which these parts are combined.

Does my proposal violate compositionality in any way whatsoever? It is true that in constructing the NP-meaning a piece of information was added that was not already present in the component meanings, but was introduced by the 'rule of combination'. As I have said, this is a misguided objection. Could it be that S&S object to my adding $c(i)$ and $c(f)$ in (7) on the grounds brought up in the above-mentioned quote from p.256 – that there would be 'an indefinite number of such rules'? This also is a misplaced objection. Rules like (7) introduce exactly the same context variable information that S&S encode as part of their lexical items. If these are variables in their version ...which will get their specific values when evaluated in the context at hand ...then they can be variables in my version. If instead S&S intend that the specific values for the function and objects are entered as part of the initial interpretation of lexical items, and this information is used to compute the value of the NP, then in my proposal these two separate stages of interpretation would be combined into one stage, in which this interpretation is done while the NP is evaluated. So, if there is one rule in their proposal, then there is one rule in

mine; if there are many rules in mine, then there are exactly as many in theirs. As I said before: either both are compositional or neither are. (Well, of course they both are compositional!)

5. Comparing Favourite Views: Semantic Intuitions and Linguistic Evidence

I don't believe we have very many clear intuitions about *recherché* areas of semantics – intuitions that are so clear that any semantic theory which does not honour them is *eo ipso* to be rejected. And I think the case of the present phenomena is the same: there are no absolutely compelling intuitions. Still, there are *some* intuitions that we would want to embody in our theories, if possible. Here's one intuition. Individual nouns – lexical items – can be semantically ambiguous. E.g., 'student' means 'a person engaged in a formal course of study' and it means 'an observer' and it means 'a beginner'. But another intuition is that individual nouns – lexical items – are simply *not also* contextually ambiguous. If we fix on one of the lexical meanings of 'student', say 'person engaged in a formal course of study', then it seems absolutely clear that the noun 'student' is not *further* ambiguous as to whose students are being discussed on this or that occasion, or which class of mine we are discussing.

The foregoing would mean that the sort of 'contextual ambiguity' that S&S allege to be attributed to nouns cannot be a part of their lexical meaning. Instead S&S must want the contextual variable information to be somehow inserted as a part of the terminal N-node that lexical items occupy in these interpreted logical form descriptions. So they make a distinction between a lexical item (with its syntactic and semantic information) and the 'meaning' of an item in their theoretical descriptions of sentences. Their claim is that there will always be $f(i)$ added to the entry of any noun. Of course, the precise values of f and i will differ from case to case

when the same sentence is used in different circumstances, and precise values of f and i will differ even in the same context as we consider the same noun in different sentences or even different nouns in the same sentence. But it is claimed that all this is compositional, because we are always adding the same *type* of value. (The semantic value for a common noun in a context would always be a set, because $f(i)$ is a function from objects to sets.) But once again this does not automatically make my favourite theory non-compositional. Although it is true that I haven't said what I take the meaning of NP's to be in general, so I have not yet committed myself to specifying what, exactly, f and i would have to be in my favourite theory, I don't think I need to. For it seems obvious that my favourite theory can eventuate in precisely the same NP-meanings as S&S's theory, but without claiming that the individual N-meanings need vary in this manner from sentence to sentence, and from occurrence to occurrence. Do we *really* want to say that it is the *noun* 'student' that changes its 'interpretation' as we turn our attention from considering all of Jason's students to considering all of Zoltan's students? Or as we turn our attention from considering Jeff's best logic students to considering his best Greek philosophy students? I would say not: while the context may dictate that we are looking at different NP-meanings ('noun phrase denotations are functions from functions from objects to truth-values to truth-values' S&S say, p. 255), it seems simply unintuitive to claim that the interpretation of the same *noun* changes from context. 'Student' means the same thing, and is interpreted the same way, in the different contexts.

On the other hand, the intuition that motivates the entire area is that sentences like (1) are true. And it seems part of that intuition to say that the reason (1) is true is due to the fact that the speaker is using 'every student' to refer to some subgroup of all the students. (To see the naturalness of this intuition, look back to the first paragraph to see that this very claim is used in

describing the phenomenon under discussion!) But that is exactly *my* proposal, and is not theirs! It is *my* claim that the NP refers to some contextually determined subgroup of all students when the sentence is uttered in context; *S&S's* proposal is that 'student' is thus used. But once again, go back to the first paragraph to see whether their proposal can be reasonably substituted into the description given there. I think most readers will concur that it cannot be.

So, the upshot of this series of intuitions is that the noun phrase *is* contextually ambiguous but that the contained noun, even considered as being in the sentence and not just as a lexical item, is not. And isn't this precisely what a further intuition says? A simple noun such as 'student', both as a lexical entry and as used in a sentence, is never *ambiguous* between 'person studying in Jeff's logic class' vs. 'person studying in Jeff's Greek philosophy class'; nor is it ambiguous among 'person studying in Jeff's class' vs. 'person studying in Jason's class' vs. 'person studying in Zoltan's class' vs. 'person studying in Bernie's class' ... That's just *not* what (semantic) ambiguity is. The context, rather, affects the interpretation of *the entire quantified noun phrase* – what 'every student' designates changes from context to context. But the context does not affect the interpretation of 'student' – which continues to mean 'person engaged in a formal course of study' no matter whether they are in Jeff's or Jason's or Zoltan's or Bernie's class; and it continues to mean this even considered as used in different sentences.

To put this intuition in a nutshell: in the relevant sense, 'student' means 'person engaged in a formal course of study' no matter whose classes we are talking about, and 'every' means 'all of the...'. Neither of these terms changes meaning from sentence to sentence. What *does* change is the interpretation of 'every student'. But that is just what my favourite theory says and what *S&S's* theory denies!

There even seems to be some linguistic evidence for my favoured view—some phenomena that are at least more naturally explained by the view that the meaning of the head noun does not vary with the context, even though the interpretation of the encompassing NP does vary. Consider sentences like (15)

(15a) Every student of Jason's passed the course, but not every one of mine.

(15b) The best student in Jason's course is a physics major; the best one in mine is a philosophy major.

Note that 'every student' in (15a) is contextually restricted: the students are restricted to those of Jason's students who are taking the course. But it is quite unnatural to say that the noun 'student' is restricted in this manner, because then it becomes awkward to see how 'one' in the second clause could work. Certainly the most natural interpretation of 'one' is that it means the same thing as 'student' in the first clause. But it does not do so under S&S's proposal—as opposed to mine. Note that (15b) is *not* interpreted as saying that the best student in my class, who is also in Jason's class, is a philosophy major. But under S&S's proposal that is what is called for under a normal interpretation of *one*-anaphora.

Of course, there may be some other formal account of *one*-anaphora that could be made to mesh with their proposal. One alternative is merely to say that *one*-anaphora is somewhat more lax in its restrictions than normal pronoun anaphora. In Stanley (2002) the example

(16a) John owns a bicycle. He rides it every day.

(16b) John is a bicycle-owner. *He rides it every day.

is cited to show that pronominal reference is not possible to incorporated nouns, like *bicycle* in (16b). However, *one*-anaphora seems not to be similarly restricted:

(16c) John is a bicycle-owner. Fran owns one too.

easily allows for *one* to refer to *bicycle*.³ In the present case, this suggestion might be fleshed out by saying that *one*-anaphora is ‘really’ to the lexical item—in the case of (15), to *student*—and not to the node containing that lexical item plus its contextual restriction. Then, the second sentence could employ its own, new contextual variable to associate with *one*.⁴ Still another option open to a defender of S&S’s theory would be to say that in the case of (15) I have mis-described the nominal restriction. It is *not* restricted to Jason’s students who are in the class, as I had said and as suggested by the examples in S&S and Stanley (2000, 2002), but rather just to students who are in the class. Indeed, it might be *always* possible to find some suitable antecedent, with a sufficiently *ad hoc* methodology. But presumably S&S wish to have some systematic method. Such a method is not given in the present article nor, it seems to me, is it likely to be as natural as the proposal that fits naturally into mine according to which ‘one’ picks up the semantic value of ‘student’—which is ‘(set of) persons engaged in a formal course of study’. Some further related examples that any proposal should give an account of are:

- (17a) Every student of Jason’s passed the course, because they were well-prepared.
- (17b) Every student of Jason’s passed the course, but not every one of mine, because they had been out partying before the final exam.

In their discussion of another of the semantic proposals for quantifier domain restriction, namely the one where contextual variables are associated with the quantifier, S&S consider

- (18) Most people regularly scream. They are crazy.

Considering a situation in which we are discussing ‘the set of things in a certain village’, they claim (p. 257) that there are two anaphoric readings of ‘they’: ‘On the first reading, it refers to all

³ Thanks to an anonymous referee for pointing this out.

⁴ Thanks to Ray Elugardo for this suggestion. He wishes it to be known that he does not endorse it.

of the people in the village. On the second reading, it refers to those people in the village who regularly scream.’ In their discussion, S&S are concerned to show that the alternative which makes the domain variable be associated with the quantifier cannot account for either of the readings. I do not wish to discuss this claim but rather wish to see how my favoured account deals with (18). It seems clear that my proposal yields the second reading in just the way that S&S wish (p. 257): ‘Ideally, one would wish to say that cross-sentential anaphora of this sort requires antecedents that are constituents (nodes) of a preceding logical form.’ And in my proposal ‘most people’ has the required domain restriction so that ‘they’ can refer to the appropriate constituent, as per their ‘ideal.’ The problem with my favourite approach comes with the alleged first reading, according to which ‘they’ refers to all the people in the village. In my proposal there is indeed no node that has this semantic value, and thus my proposal does not yield the first reading in their ‘ideal way.’ The case would be different were the sentences (still talking about ‘the set of things in a certain village’)

(19) Most of the people regularly scream. They are crazy.

Here there is a node, one immediately dominating ‘the people’, which will do the trick. (And again, the node immediately dominating ‘most of the people’ will serve to get the other reading.)

I think, in fact, that S&S have confused (18) with (19) when it comes to ‘the first reading’. I think that sentence (18) has no such reading at all, but only has the reading where ‘they’ is anaphoric to those people who regularly scream. In support of this assessment, consider a context where we are discussing instructors at institutions of higher education, both colleges and universities.

(20) Most college professors teach five or more courses per year. They are underpaid.

(and other sentences of this form). I find it quite difficult to make ‘they’ refer to all instructors in institutions of higher education, and I think readers will agree. Rather, it refers to those college professors who teach five or more courses per year, or perhaps to all college professors. It simply does not seem to have any reading where it picks out all the college *and* university professors. Despite the fact that sentence-pairs like (18) and (20) do not naturally license the anaphoric connection of ‘they’ with some group that is unrestricted by the assertion in the preceding sentence (i.e., they do not normally license the ‘first reading’ of such sentences), we can nevertheless construct examples where this non-license is overcome by sufficiently strong background knowledge. Consider

(21) Most female bears give birth every third year. They hibernate for up to three months.

The ‘they’ here seems most naturally to be anaphoric to female bears, or perhaps to female bears who have given birth in that year. But in the common background we know that hibernation in bears is not restricted to bears who have just given birth, nor is it sex-specific to females. And so this background information ‘overcomes’ the linguistically available anaphora, and we can understand the second sentence of (21) to be about all bears. Surely, however, this goes against the linguistic form of (21) and even causes one to pause for a moment to decide whether he or she has understood correctly. Sentences like this constitute one of those peculiar cases that seem to provide such-and-so reading upon first hearing, but then this reading seems to recede further and further into the background the more one thinks deeply about it. For this reason I think S&S’s ‘first reading’ of (18) cannot be used to argue against my favourite proposal. And even more strongly, since their proposal allows both readings equally, I would argue that the facts of these cases actually tell *against* their proposal.

6. Conclusion

S&S argue against a huge number of different analyses of the phenomenon of quantifier domain restriction. One of their targets is my own favourite view. But I find their (single) argument against my proposal to be flawed.

Department of Philosophy

University of Alberta

References

- Bach, K. (2000) Quantification, qualification and context: A reply to Stanley and Szabó. *Mind & Language*, 15, 262-83.
- Neale, S. (2000) On being explicit: Comments on Stanley and Szabó, and on Bach. *Mind & Language*, 15, 284-94.
- Stanley, J. (2000) Context and logical form. *Linguistics and Philosophy*, 23, 391-434.
- Stanley, J. (2002) Nominal restriction. In G. Preyer & G. Peter (eds) *Logical Form and Language*, Oxford: Oxford University Press, pp. 365-88.
- Stanley, J. & Z. Szabó (2000) On quantifier domain restriction. *Mind & Language*, 15, 219-61.