

Context, Compositionality and Calamity*
To appear in Mind & Language: do not quote without permission

Abstract.

This paper examines an attempt made in a series of articles (Stanley 2002, a.o.) to create a syntactic placeholder for contextual information. The initial shortcoming of Stanley's proposal is that it does not easily integrate these placeholders with domain-restricting information syntactically encoded elsewhere in the utterance. Thus, Stanley erroneously predicts that a sentence in which quantifier-restricting information encoded in (for example) a prepositional phrase conflicts with quantifier-restriction valued by context is internally incoherent.

I continue by exploring the space of possible solutions to this problem that are available to Stanley, demonstrating how each of these possible solutions results in its own interpretation problem and, ultimately, fails. In doing so, I argue that Stanley's syntactic/semantic approach to context is ultimately untenable.

1. Why Compositionality?

The claim that natural language is compositional is the claim that the meaning of a sentence in natural language is exhaustively determined by the meaning of its parts together with its syntax. Compositionality is foremost a claim about the semantic content of a sentence relative to a context. Many have taken the semantic content of a sentence relative to a context to be the *truth conditions* of that sentence relative to that context (this position is typical among semanticists).

But advocates of *truth-conditional pragmatics* (as the program is called by Recanati) deny that the truth conditions of an utterance of a sentence are due to a compositional process of interpretation. Such

* For help on this paper I am most indebted to Alexander Hughes. I'd additionally like to thank Amy Burke, Sam Cumming, Jerry Fodor, John Hawthorne, Ernie Lepore, Peter Ludlow, Ian Proops, Adam Sennet and Jason Stanley for their comments and suggestions. Thanks also to the Departments of Linguistics and Philosophy at Rutgers University for their support.

theorists disagree with each other with respect to what the semantic content of a sentence relative to a context is when it's not a full proposition; some take it to be a 'propositional radical' as in the work of Kent Bach, or a 'partially articulated conceptual representation' as in the work of relevance theorists such as Sperber and Wilson 1996, Carston 2002. They're unified, however, in rejecting the thesis that the intuitive truth conditions of an utterance are invariably the result of a compositional interpretive process of the sentence uttered.¹

Here is one reason to think that the intuitive truth conditions of an utterance of a sentence are not the result simply of combining the values of the words used in the sentence uttered. Consider an utterance of the sentence *Alex is tall*. This sentence could intuitively mean one thing if Alex is in the fifth grade and another if Alex is a basketball player. But there appears to be nothing in the sentence that explains this potential shift in meaning. Similarly, to take a well-known example, quantified sentences seem to have different intuitive truth conditions relative to different contexts of use. One and the same quantified sentence can intuitively have different truth conditions in different contexts. Unless this difference can be traced to something in the sentence uttered, the claim that intuitive truth conditions are the result of a compositional semantic process is violated.

2. Stanley's Solution.

One way to reconcile the apparent context dependence of truth conditions with compositionality is the following: if relevant features of context were encoded in syntax, then its effects would not violate compositionality. A sentence could still be interpreted via only its constituents and syntax.

¹ The space of positions is of course more complex. Cappelen and Lepore (2005) agree with the advocates of truth conditional pragmatics that the *intuitive* truth conditions of an utterance are not the result of a compositional process, and for many of the same reasons. But using some resources from Davidsonian semantics, they nevertheless contend that semantics always results in truth conditions.

In a series of articles (Stanley & Szabó 2000, Stanley 2000, 2002a, 2002b), Stanley attempts to do just that. His goal is ‘to defend the thesis that all truth-conditional effects of extra-linguistic context can be traced to logical form’ (2000 391). He focuses specifically on nominal restriction as a case study. In examining nominal restriction, Stanley ‘hope[s] to provide convincing evidence of the promise of the project of reducing all apparent effects of context on semantic content to a small number of sources’ (2002a 366).

He does this by postulating a phonetically null syntactic variable, referred to as a *domain variable*, whose value is fixed by context. Given that these variables are a feature of the syntax, the argument goes, compositionality is preserved. The explanation is as follows: every (head) noun is associated with an object variable i and a function variable f . A sentence like (1), then, would have the underlying form of (2):

1. Every man runs.
2. [Every <man, $f(i)$ >] runs

‘The value of i is an object provided by context, and the value of f is a function provided by the context that maps objects onto quantifier domains. The restriction on the quantified expression *every man* in [(2)], relative to context, would then be provided by the result of applying the function that context supplies to f to the object that context supplies to i ’ (Stanley 2002a). Thus, (1), once context supplies values for its variables, can be interpreted as (3) or other such variants depending on context.

3. Every man [participating in today’s marathon from Dexter to Ann Arbor...] runs.

Stanley claims that ‘one might worry that there is no independent evidence for the resources needed to treat this sort of context-dependence. If so, then postulating the mechanisms needed to treat this sort of context-dependence in the semantics may seem *ad hoc*’ (ibid. 381). He relies on what he calls *the argument from binding* to provide syntactic evidence for domain variables.

4. In most of his classes, John fails three students.
5. In most of his classes x , John fails three students in x .

In (4), the quantifier domain of the first QP, *most of his classes*, binds the quantifier domain of the second QP, *three students*. This yields (5), the corresponding intuitive reading. Given the assumption that binding is a syntactic phenomenon, the fact that there is a bound reading above is evidence for a variable. If binding occurs in the syntax, the argument goes, there has to be something in the syntax to bind.

With this proposal, Stanley eliminates the problem of context in a subset of linguistic phenomena. Compositionality is not violated because contextually relevant information is encoded in the syntax at LF. If this analysis is correct, and if it can be extended to other contextually interesting linguistic phenomena, then we can preserve the idea of compositionality with respect to natural language.

3. An Interpretation Problem.

3.1 An Overview.

Unfortunately, Stanley's analysis runs into a few problems. Because the value of these variables is determined by context, Stanley's analysis makes incorrect predictions about the truth conditions of sentences. His analysis needs to be revised so that domain variables can be valued both by context and by domain-restricting information elsewhere in the syntax.

Since Stanley's work is centered on syntactically motivating the *existence* of these domain variables, much of his discussion on the implementation of the variables is inexplicit. In this paper, I spend a fair amount of time exploring and then rejecting possible interpretations of his claims. Since this discussion can get chaotic, I include a diagram of the argument scheme on the last page.

Take the following pair of examples, (2) and (3) in Stanley 2002b.

6. Every student answered every question.
7. [Every student]_i answered every question on his_i exam.

Notice that the pronoun *his* in (7) is anaphoric on *every student*. It should be read as: *For every student X, X answered every question on X's exam*. As explained above, Stanley wants to argue that the quantifier phrase *every question* in (6) contains a domain variable that receives its value from context.² When interpreting (6), then, the hearer uses context to value the domain variable, rendering a reading that parallels the one in (7).

Now let's consider an actual utterance of (7). Presumably, the quantifier phrase *every question* in (7) contains a domain variable just as does its matching quantifier phrase in (6). But notice that it's not necessary for the hearer to appeal to context to resolve the domain of the quantifier *every question* in (7); he *could* receive all the information he needs from the prepositional phrase (the PP) *on his exam*. So it seems that, when there's domain-restricting information in the syntax but outside of the QP, Stanley's variable and explicit material in the sentence are doing overlapping work.

Stanley has two options here (Options I and II on the diagram): I) he can claim that PPs that restrict quantifier domains should be interpreted differently than PPs that don't affect the domain of a quantifier, or II) he can claim that *domain variables* whose QPs have their domain affected by information in the explicit syntax should be interpreted differently than domain variables whose QPs are unaffected by the explicit syntax. I'll discuss each of these options in turn.

3.2 Option I.

With respect to Option I: what does it mean for a PP that encodes domain-restricting information to play a non-standard role in interpretation? Normally, a PP contributes novel information to the interpretation of a sentence. In a non-standard role, then, the PP would either A) not contribute information at all, or B)

² He argues the same for the QP *every student*, of course, but I'll discuss the value of only one of the QPs in an attempt to keep things simple.

contribute information that is not novel. These branches correspond to Options A and B on the diagram.

I'll examine each of these options in turn.

Let's assume Option A: when a PP represents domain-restricting information, this information is not used in interpreting the sentence. Now, let's examine the sentences above, (6) and (7), repeated here, in addition to a new sentence:

6. Every student answered every question.
7. [Every student]_i answered every question on his_i exam.
8. Every student answered every question on John's exam.

Imagine, in discussing these sentences, a world in which there are three students (and only three students):

Adam, Sam and John. Each student's exam has five questions, but no student has a question that is on another student's exam. So Adam has questions 1-5, Sam has questions 6-10, and John has questions 11-15.

And, in this world, it is in fact the case that every student answered every question on his exam. So: Adam answered questions 1-5, Sam answered questions 6-10, and John answered questions 11-15. For the interpretation of (8), hold fixed the context which makes it true that (6) and (7) have the same meaning.

Intuitively, (8) is false in this world and context. It is true if and only if Adam, Sam and John all answered questions 11-15. But in this world, Adam and Sam did *not* answer questions 11-15. Assuming that the PP contributes no information to the interpretation of (8), and assuming this world and context, *Every student answered every question on John's exam* is true. This is because context and context alone values the domain variable in [every <question, f(i)>]. And context will fix it to mean *every question on his exam*.

When we don't allow QP-external syntactic information to restrict quantifier domains, the sentence will be true regardless of the meaning of the PP. Option A clearly makes the wrong predictions about sentences that encode conflicting syntactic and contextual information because (8) is clearly false.³

³ Bach (2000 263-4) makes similar observations about the relevance of contradiction to contextually valued information. He points out that one reason we might think that 'loose talk' is not just an instance in which domain-restricting information is

Let's examine sentences (6)-(8) in the same world and context, choosing instead to go with Option B. Option B, remember, assumes that PPs with domain-restricting information play their normal role in interpretation, but that the information they contribute is just redundant. This is true for (7): the two sources of information are unproblematically redundant. But (8) is a sentence in which the information encoded in the syntax and the information encoded in context are incompatible. Given that, under this option, both sources contribute to the interpretation of the sentence, the sentence itself will be internally incoherent. The domain variable, valued by context, restricts the quantifier phrase *every question* to mean *every question on his exam* (true in this world), while the PP restricts the quantifier phrase *every question* to mean *every question on John's exam* (false in this world).⁴

In demonstrating that Options A and B lead to false predictions, I've demonstrated that Option I leads to false predictions. Stanley therefore cannot claim that he can solve the problem of information encoded in the syntax by altering the contribution of domain-restricting PPs. This forces him to claim that, in his analysis, domain variables behave differently when they occur in sentences with (QP-external) syntactically-encoded QR information (Option II on the diagram).

3.3 Option II.

In Stanley's standard analysis, domain variables are valued, and they're valued by context. However, Stanley must account for domain-restricting information encoded in the syntax by allowing this


encoded covertly rather than overtly is that sentences such as *Jack and Jill went up the hill* pass Grice's test of cancellability. If we were to instead say *Jack and Jill went up the hill, but not together*, we would not be uttering a contradiction.

⁴ Stanley doesn't discuss whether an erroneously-valued domain variable renders a sentence ungrammatical or uninterpretable. If it leads to an ungrammatical sentence, I would guess that the intrasentential conflict above would be analogous to the ungrammatical underlying form **I are sleepy*. This sentence contains incompatible syntactically encoded information: the agreement features on the subject are first-person, while the agreement features on the verb are second-person. This incompatibility yields an ungrammatical sentence. Another analogy, relevant if domain variables effect interpretation instead of grammaticality, is the sentence **She is a bachelor*. This sentence has two different sources of information regarding the subject; one source, the noun in subject position, has phi-features picking out the referent as female. Another source, the phrase predicating the noun, picks out the referent as male. This is an example of intrasentential conflict that is grammatical but uninterpretable.

information, too, to give value to the variable. One way to do so is this: the hearer first appeals to syntactically encoded information to resolve referents in a sentence. If there is domain-restricting information encoded in a PP, as there is in (8), then this linguistic information will fix the value of the domain variable. If there is no such information encoded in the syntax, the hearer resorts to context to fix the value of the domain variable.⁵

On this view, then, the value of the variable in *every question* in (8) is fixed by the information encoded in the PP *on John's exam*. Since this value is hereby fixed, there is no need to appeal to context to value the variable. Information is consistent intrasententially, and the sentence is false in our prescribed world, which matches our intuition.

The diagram below illustrates this two-step process: in Step One, information in the syntax is evaluated to see if it can value the domain variable. If it can, it yields Value 1. If it cannot, then context is evaluated in Step Two, yielding Value 2.



To sum up the preceding: I've pointed out that, in addition to using context to value domain variables, Stanley's analysis needs to account for sentences in which domain-restricting information is encoded in the QP-external syntax of the sentence. I've shown that his analysis cannot account for this by discounting or altering the role of the syntax itself. He therefore has to do it by altering the way his domain

⁵ There may be other options for Stanley to pursue here; he may choose to have the value of the variable determined *partially* by linguistic context and *partially* by non-linguistic context. An anonymous M&L reviewer points out that such a solution would require a supplemental pragmatic theory that Stanley does not have, but nevertheless seems feasible. I'm of the opinion that it would run into problems similar to those I discuss above: What happens if the information encoded in linguistic context clashes with the information encoded in non-linguistic context? How would we interpret a variable that, for some reason, has only been partially valued?

variables are valued. This gives way to a revised version of Stanley's analysis, one that has a two-step interpretation process.⁶

4. Another Interpretation Problem.

4.1 An Overview.

The best solution to the problems raised in the previous section was a two-step analysis. But this solution brings new problems. Stanley's approach requires that a sentence be interpretable when the domain variable has not been assigned a value. However, a domain variable needs to be assigned a value in order for the sentence to be interpretable. This creates a contradiction within the theory and shows that the system is unimplementable. I'll elaborate on each of the above claims in turn.

Under Stanley's revised analysis, a sentence needs to be interpreted while the variable is still not fixed. (I will refer to a variable that has not been given a value by either syntax or context as a non-fixed variable; this will be ambiguous, as we'll see shortly, between a variable with no value and a variable with a default value.) This is because there are some instances in which the hearer needs to interpret the rest of the sentence to determine how to value the variable.

Imagine that A and B are having a conversation at a nursery.

⁶ The following solution might seem *prima facie* attractive to Stanley: syntactically encoded information makes its contribution to the value of the variable *via* context, getting rid of any need for a two-staged approach. On this view, the speaker's utterance is a *part of* the context (so, as an anonymous M&L reviewer points out, under this view it would actually be impossible to hold fixed a context in our discussion of (7) and (8) above). But this theory has disastrous empirical consequences. Information provided by context has a very different status than information provided by the syntax. The former is cancellable in a Gricean sense (see Bach 2000 263-4): if I were to utter (6), *Every student answered every question*, in a context that clearly restricts the QP to mean *every question on John's exam*, I could still follow up my utterance by adding, "The questions they answered were on their own respective exams" without contradicting myself (although the hearer might be temporarily thrown off). The latter is not similarly cancellable. If I utter the sentence in (8), *Every student answered every question on John's exam* and follow it up by adding, "The questions they answered were on their own respective exams," I would be contradicting myself. If Stanley were to endorse the view that syntactically encoded information *does* value the variable but only indirectly, as a contributor to context, he would lose his ability to account for this difference in cancellability. The solution is also unattractive theoretically. If syntactically encoded information can be a part of context then it seems like anything can, leaving the variables actually unconstrained. This surely goes against the strict notion of compositionality that Stanley is trying to preserve by proposing the variables. Thanks to Adam Sennet and Peter Ludlow for helpful discussion here.

9. A: Every baby at this nursery walks around as if it's drunk.
 B: Every baby is born without the ability to control its own muscles.

Speaker A first makes a comment that restricts the domain of relevant babies to those at the nursery. That is, once A makes his statement, a contextual restriction of the quantifier domain is available to B in the salient discourse. But note B's reply. If A, the hearer, were able to fix the value of the variable associated with *baby* without interpreting the entire sentence, he would use the most salient context to do so. Let's assume that A's previous statement is the most salient context relevant to the restriction of the domain of *baby*. So A will fix $\langle \text{baby}, f(i) \rangle$ so that it's restricted to babies at this nursery. He will then go on to interpret the sentence to mean *Every baby at this nursery is born without...* However, that's not how B intended his statement; he intended for the domain of babies in his utterance to be unrestricted. Speaker A has misinterpreted the sentence, and now has an erroneous set of truth conditions for it.

If B had been able to interpret the entire sentence, he would have been able to learn what was being predicated of the QP. He could have used this information to better determine (indeed, correctly determine) what aspect of the context he should utilize to restrict the domain of quantification.

So it seems that, before valuing a domain variable, the hearer interprets the rest of the sentence (i.e. the VP). In light of the interpretation of the VP, candidate domain restrictions can be rejected; the hearer can't determine what syntax or context is relevant for restricting the quantifier unless he knows what is being predicated of the quantifier phrase. We can thereby conclude that, under Stanley's analysis, a sentence must be interpretable while it contains a non-fixed variable.

However, we can also argue that a domain variable needs to have a value in order for the sentence to be correctly interpretable. Take the following sentence:⁷

10. Good steaks are rare.

⁷ Thanks to Ian Proops for this example.

Let's say this sentence is uttered in the context of a discussion about the scarcity of good cuisine in Russia.

The domain variable would alter the NP from [good steaks] to [good steaks in Russia].

However, in order to disambiguate the lexical entry *rare*, one needs to know what *rare* is predicated of. If we were to utter this sentence in a different context, a discussion about what sort of steak one should order at my favorite steakhouse, we could interpret *rare* to mean 'under-cooked,' instead of 'scarce.' So (10) is an instance in which the quantifier domain needs to be properly restricted in order for the rest of the sentence to be correctly interpretable.

To sum up: it seems that we can argue both that a sentence needs to be interpreted in order for context to restrict a domain variable *and* that the domain variable needs to have a value in order for the sentence to be interpreted. But there are two ways Stanley can side-step this contradiction. He can claim either that a non-fixed variable doesn't have a value (Option α on the diagram) or that a non-fixed variable has a default value (Option β on the diagram). I'll discuss each possibility in turn.

4.2 Option α .

If a non-fixed variable just didn't have a value, the sentence would fail to have a truth value. Let's look at an underlying form with a non-fixed variable:

2. Every $\langle \text{man}, f(i) \rangle$ runs.

If we interpret a non-fixed variable as a variable without a value, as in Option α , then the QP has an incomplete referent. Stanley doesn't give a derivation of truth conditions for sentences with domain variables. And as I mentioned in footnote 4, he does not make explicit what the consequences for having an erroneously-valued domain variable might be. If domain variables affect the sentence's interpretation, then asking someone to provide a truth value for (2) would presumably be like asking someone to provide a truth

value for the sentence *Santa Claus is fat*. Since *Santa Claus* lacks a referent, the sentence *Santa Claus is fat* arguably lacks a truth-value. If domain variables affect the sentence's grammaticality, then we can compare (2) to any other type of sentence whose underlying form has an empty node, a NP missing a gender inflection in French, or an object missing the *-o* object marker in Japanese.

We can thus conclude that pursuing Option α , then, leads either to reference failure (a sentence without a truth value) or to a sentence with an incomplete LF (an ungrammatical sentence). So in order to get around this second interpretation dilemma, Stanley has to claim that variables have a default value (Option β). In what follows, I will first discuss how interpretation involving a domain variable with a default value would work under Stanley's original analysis. Then I'll derive a sentence in Stanley's revised analysis. This discussion will lead us to conclude that Stanley's revised system requires that a variable cannot be revalued after Step One. In light of this new requirement, then, I will demonstrate how Option β , and therefore how Stanley's revised approach, fails.

4.3 Option β .


To start, let's assume, arbitrarily, that the default value for a non-fixed variable is the least restrictive value possible (the value for f would be the self-identity function, as in Stanley 2002a fn. 4, and the value for i would be the entire domain). In Stanley's standard analysis, a derivation for (3) above in which a non-fixed variable has a default value would go like this:

3. Every man [participating in today's marathon from Dexter to Ann Arbor...] runs.

Every $\langle \text{man}, f(i) \rangle$ runs $f = \text{self-ID}, i = \text{all possible things}$ Underlying Form	\longrightarrow	Every $\langle \text{man}, f(i) \rangle$ runs $f = \text{participating in}, i = \text{today's marathon...}$ Surface Form
--	-------------------	--


The sentence starts out as meaning, roughly, *Every man in the world runs*, and the hearer then uses context to interpret the sentence as *Every man participating in today's marathon...* However, we've already established that Stanley needs an additional step. Here's a derivation of sentence (8) in the two-step analysis:

8. Every student answered every question on John's exam.



Notice that, in Stanley's revised analysis, it must be the case that if the variable is valued at Step One, it cannot be revalued at Step Two. If this were to have happened in the derivation above, given the Adam/Sam/John world and context discussed above, context would have revalued the variable to read *Every student answered every question on his exam*, rendering (8) true. This is precisely the prediction we designed the two-step program to avoid. We can conclude that, with the adoption of the two-step analysis, Stanley needs to stipulate that a variable cannot be revalued if it has a value at Step One.

Which brings us back to our discussion of Option β . Any sentence whose quantifier needs to be restricted by context will be false because any sentence that is not valued at Step One by syntax will have an unrestricted quantifier (or whatever other default value is given). Suppose I say to you *Every man runs*, but discourse restricts the quantifier domain to refer to every man in the race. The derivation, under Option β , would go like this:



The value of the variable is not fixed by syntax in Step One because there is no quantifier-restricting information encoded in the syntax. Furthermore, the variable is not fixed by context in Step Two because Stanley has to say that context cannot fix the value of a variable if it has a value after Step One.

The sentence, which was intended to mean *Every man in the race runs*, now is only true if every man in the world runs.

5. Conclusion.

I have argued that Stanley is caught in a double-bind: it is both the case that a sentence needs to be interpreted in order for context to restrict a domain variable *and* that the domain variable needs to have a value in order for the sentence to be interpreted. I have also argued that there is no way to characterize a non-fixed variable such that it allows for a sentence to be interpretable before the variable is valued. The variable can neither start out as non-valued or with a default variable. Stanley's theory is therefore trapped.

Of course, I have not suggested a *positive* solution to the problems discussed by Stanley. Perhaps some other semantic theory will account for the data, or perhaps a pragmatic solution employing 'expansion' (Bach 2000) or 'free enrichment' (Recanati 2004, Carston 2002) is viable (though these approaches are criticized in Stanley 2002b). Here is not the place to speculate about the directions such a solution would take.

References.

- Bach, K. 2000: Quantification, qualification, and context: A reply to Stanley and Szabó. *Mind & Language*, 15, 262-83.
- Cappelen, H. and Lepore, E. 2005: Radical and moderate pragmatics: Does meaning determine truth conditions? In Z. Szabo (ed.), *Semantics vs. Pragmatics*. Oxford: Oxford University Press.
- Carston, R. 2002: *Thoughts and utterances: the pragmatics of explicit communication*. Oxford: Blackwell.
- Recanati, F. 2004: *Literal Meaning*. Cambridge: Cambridge University Press.
- Sperber, D. and Wilson, D. 1996: *Relevance: communication and cognition*. Oxford: Blackwell.
- Stanley, J. 2000: Context and logical form. *Linguistics and Philosophy* 23: 391-434.
- Stanley, J. 2002a: Nominal restriction. In G. Preyer and G. Peter (eds.) *Logical Form and Language*. Oxford: Clarendon.
- Stanley, J. 2002b: Making it articulated. *Mind & Language*, 17: 1&2, 149-168.
- Stanley, J. and Szabó, Z. 2000: On quantifier domain restriction. *Mind & Language*, 15 2&3, 219-261.

STANLEY'S STANDARD ANALYSIS:
(Stanley 2000, 2002a, 2002b)

Problem:

How is the variable valued when there is domain-restricting information encoded in the syntax external to the Quantifier Phrase (i.e. in a PP)?

OPTION I:
Exceptional Information:
When there is QP-external domain-restricting information, that information plays a non-standard role

OPTION II:
Exceptional Variables:
When there is QP-external domain-restricting information, the domain variable plays a non-standard role

Option A:
PP does not contribute info to interpretation; is ignored

Option B:
PP contributes info, but its contribution is redundant

STANLEY'S REVISED ANALYSIS:
Two-Step Interpretation

☠ Falsely predicts that sentences with erroneous domain-restricting info in the PP are true

☠ Leads to sentences that are internally incoherent and, plausibly, ungrammatical

Problem:

How are sentences with non-fixed variables nevertheless interpretable?

Option α :
Non-fixed variables have no value

Option β :
Non-fixed variables have a default value

☠ Leads to sentences without truth values (reference failure) or ungrammatical sentences (UFs with empty nodes)

☠ Leads to sentences with incorrect truth conditions