SPi | MILA | mila_12151 | B | Dispatch: June 30, 2017 | Journal: MILA
Journal Name | Manuscript No. | Author Received: | No of pages: 28 | TS: karthik

Logical Form and the Vernacular Revisited

AQ1 ANDREW BOTTERELL AND ROBERT J. STAINTON

Abstract: We revisit a debate initiated some 15 years ago by Ray Elugardo and Robert Stainton about the domain of arguments. Our main result is that arguments are not exclusively sets of linguistic expressions. Instead, as we put it, some non-linguistic items have 'logical form'. The crucial examples are arguments, both deductive and inductive, made with unembedded words and phrases.

1. Introduction

> ... *sub*sentential expressions such as singular terms and predicates... cannot serve as premises or conclusions in inferences (R. Brandom, 2000, p. 40).

Our goal in this article is to revisit a debate initiated some 15 years ago by Ray Elugardo and Robert Stainton (2001) in their paper 'Logical Form and the Vernacular'. Very broadly, the conclusion of that paper was that the domain of arguments cannot be identified with the linguistic, and hence that Vernacularism— the view that only linguistic expressions have logical forms— is false. Here, we revisit that conclusion, together with the argument for it. The result of our revisiting, we think, is an argument that, while resting on less controversial grounds, results in a conclusion of deeper and more obvious philosophical interest.

In what follows, we will try to present a stand-alone argument that can be assessed on its own merits. To that end, we will not presuppose detailed knowledge of the original Elugardo and Stainton paper, although we will at times make reference to it. Here is the game plan. First, we will explain our topic, and indicate what it is—and, importantly, what it is not—about. Second, we will present our new and improved argument and compare and contrast it with Elugardo and Stainton's original. And third, we will consider an array of novel, interesting, and initially plausible objections to the new argument, together with our replies.

Acknowledgements: A first draft of this material was presented at the University of Waterloo on 27 September 2013. It was subsequently delivered at PhiLogica III in Bogotá, Colombia, on 12 February 2014 and at the Universidad del Caribe in Cancún on 26 January 2016. We are grateful to all those in attendance, but must especially single out Santiago Amaya, Tim Kenyon, Kepa Korta and Arthur Sullivan. Written comments from Axel Barceló, Eros Corazza, Ray Elugardo, Peter Pagin, Stefano Predelli, Robert Thompson and two anonymous referees were very helpful, as were conversations with Western Ontario colleagues Jon Life, Nic McGinnis and Ileana Paul. Financial support for this project was provided by the Social Sciences and Humanities Research Council of Canada.

Address for correspondence: 1151 Richmond St., London, ON, Canada, N6A5B6.
Email: abottere@uwo.ca; rstainto@uwo.ca

2. Part One: The Revised Argument

2.1 Our Twin Questions

Let us begin by defining our topic. Our inquiry is guided by two questions, one general, and one specific. The general question is: are there arguments whose premises and/or conclusions are not linguistic expressions? The specific question is: if there are arguments whose premises and/or conclusions are not linguistic expressions, are there any such arguments in which logical structure plays a weight-bearing role?

Both questions stand in need of clarification. First, when we use the term 'argument' we are not asking whether, to use Leo Groarke's (2005) notion, there are 'modes of arguing' that are non-linguistic, or that involve non-linguistic components. We think it fairly clear that there are non-linguistic arguing-behaviours: one can wave one's arms, or jump up and down, or point to a location on a map in order to make one's argumentative point. Rather, we are asking whether the elements that make up the 'argument-thereby-made', its premises and conclusions, can themselves be non-linguistic. In this sense of 'argument', we further clarify, something will count as an argument if full-blooded entailment relations hold among its elements, although something could count as an argument in the very same sense even if it lacked an 'alethic guarantee'.

Second, when we use the phrase 'linguistic expression', we mean to include expressions of both natural and artificial languages. There are, of course, arguments that are made up of purely logical symbols, and so in that sense there are arguments that are not composed entirely of 'linguistic expressions'. But our question is: are there arguments some of whose elements are not symbolic expressions at all?

Third, what do we mean when we ask whether there are arguments where structure plays a weight-bearing role? That is, to coin a phrase that will recur below, what do we have in mind by the expression 'structure-involving'? There are two ways in which an argument could display structure. First, the structure could enter into the picture at the level of propositional content, in the following sense: the same proposition might occur in several different places in the argument, and this might be important for the justificatory relations that hold among the various elements of the argument. In this sense, structure is to be found across the argument as a whole, at the 'macro level'. But there is a second way in which structure could enter into the picture, and that is at the level of sub-propositional constituents: the same object or property could reoccur. Here the structure is to be found *inside* the propositional elements of the argument, i.e. *within* the premises and conclusions of the argument. To see what we have in mind, think of the difference in terms of whether the structure could be captured with the propositional calculus alone, or whether the predicate calculus with quantification would be required in order to reveal the 'micro level' presence of structure. The former kind of structure could be spotted, e.g. in *modus ponens*; the latter in existential generalization.

One final point of clarification. We will not be concerned with the following question: could there be structure-involving arguments whose elements are not linguistic *in the absence of linguistic things entirely*? For all we know, and as far as our

considerations show, the kind of alinguistic arguments that we identify in what follows might be possible only against a backdrop of genuinely linguistic arguments. And if that is the case, then people like Davidson, Gauker (1994, 2013), Sellars and Wittgenstein might still be on the right track (see Kenyon, 1999, for discussion.)

2.2 A Quick Argument for a 'Yes' to Both Questions

So much for preliminaries. Let us now consider a quick argument for the conclusion that there exist arguments whose premises and/or conclusions are not linguistic symbols.

What would be required for an affirmative answer to this question? At a minimum, we would need to identify an argument that has elements that are truth evaluable; that stand in supporting relations to other elements of the argument; and where at least one element in the argument—i.e. at least one member of the set of premises/conclusion—is not a linguistic expression, whether natural or artificial.

With these desiderata in mind, consider a much-debated example from the original Elugardo and Stainton paper. It appears to satisfy all three conditions:

> A discussion has taken place between Alice and Bruce in which Bruce takes the position that there are not really any colored objects. A day or so later, Alice meets Bruce. Having just read G.E. Moore, she offers the following counter-argument. She picks up a red pen, and says 'Red. Right?' Bruce, guileless fellow that he is, happily agrees. Alice continues, 'Red things are colored things. Right?' Bruce nods. At which point, Alice springs her trap: 'So Bruce, there is at least one colored thing. This thing.' (Elugardo and Stainton, 2001, pp. 402–3)

The sub-sentential premise in this argument really is truth-evaluable and really does support the conclusion that there is at least one coloured thing: if the former premise is true, the latter conclusion is more likely to be true (indeed, we think that there is a full-on entailment relation here, although that is not essential). Crucially, however, the first premise in the argument is not a linguistic expression of any kind. Because what is actually tokened in the Alice/Bruce example, namely the word 'red', does not *itself* express a proposition, not even at a context, *it* can't be the thing that is truth-evaluable and stands in a justificatory relationship to other elements of the argument (more on this below.) Perhaps the thing that is truth-evaluable and that stands in justificatory relations is a mental thing, such as a belief; or perhaps it's an abstract thing, such as a proposition. Either way, the point is that although Alice utters the word 'red' in the course of making her argument, the premise she puts forward isn't a linguistic expression. In short, the Alice-Bruce example seems to support the claim that there are arguments whose premises and/or conclusions are not linguistic. (The fact that linguistic expressions are deployed in putting forward the argument is not, of course, an objection to our claim that some elements of the argument thereby put forward are not themselves linguistic in nature.)

But, recalling our second question, more can be said about the foregoing argument. It would appear that the argument contains as an element a singular proposition about an object *a*, to the effect that it has a property F. It would appear that there is a bridge premise, which says that anything that has F has G. It would appear that there is an intermediate conclusion that *a* is G. And finally, it would appear that the inference rule of existential generalization is applied to yield the conclusion that something is G. We thus have both kinds of structure highlighted above, that is, macro-level structure across the premises/conclusion as well as micro-level structure within them. Hence, the phenomenon of sub-sentential assertion yields a positive answer to both of our two guiding questions: first, there exist arguments whose premises and/or conclusions are not linguistic expressions; second, in some of those arguments structure plays a weight-bearing role.

It is worth noting that we can describe additional examples involving arguments the premises of which do not involve an alethic guarantee that their conclusions will be true. Consider the following. Cara and Dan are trying to determine who took the last beer from the fridge. They reason as follows:

Cara: On the door handle. Fresh red paint. Same colour as the porch.
Dan: So, whoever was painting the porch must have taken the beer?
Cara: Exactly! Eloise!

Simplifying somewhat, in this argument 'On the door handle' is a prepositional phrase, 'Fresh red paint' is an adjective phrase, 'Same colour as the porch' is some kind of nominal, and 'Eloise' is a name. None is a sentence. (Or so it appears. The issue will be addressed in more detail at the end of the paper.) There is clearly an argument here, although it does not exemplify an alethic guarantee: it is merely inductive. This reinforces the result that there exist arguments whose premises and/or conclusions are not linguistic expressions, and in which structure plays a weight-bearing role.

Finally, and as an aside, perhaps even informal fallacies would support our conclusion, provided that they are counted as arguments. For even informal fallacies involve structure, though the structure involved does not guarantee any sort of justificatory relationship between premises and conclusion. (Thanks to Santiago Amaya and Arthur Sullivan for raising this point.)

2.3 Compare and Contrast: The Original Argument and the Revisited Version

As we would now use the term, a positive answer to our two questions amounts to the following: that things other than natural language expressions have *logical form*. What we have is an argument for the conclusion that logical forms can exist outside 'the vernacular' of ordinary language. (Where, obviously, the term 'logical form' does not refer to the level of Logical Form (LF) posited generative grammar. Arguably LFs, by definition, are only had by items of natural language.) Consider

now how the present argument against Vernacularism relates to the earlier one in Elugardo and Stainton's paper on the same theme.

First, that previous argument focused on natural language expressions only, whereas we want to extend the conclusion in the following way: there are, we would now maintain, things that aren't linguistico-logical symbols of any kind, yet have logical form.

Second, that previous argument addressed not just whether things other than natural language expressions have logical form, but also, as we would put it now, whether such things have logical forms non-derivatively, i.e. originally. (This is related to the point above, about whether there could be arguments with non-linguistic premises/conclusions with structure even if there were no languages at all.) While this is an interesting issue, it is a red herring given the present topic, which is framed by our broad and narrow questions.

Third, the previous argument was thoroughly, and problematically, psychologized. Elugardo and Stainton proposed the following: if one can recognize an argument and its structure, then the argument has structure. More precisely, Elugardo and Stainton said that if an agent can recognize an argument and its structure, and if she need not recognize that structure by means of natural language, then the argument has its structure non-derivatively. In the present version, we entirely omit reference to agents' recognitional capacities and simply observe that there is an argument containing a truth-evaluable premise that provides support for its conclusion, and where structure plays a weight-bearing role. (The previous argument had the related disadvantage of misleading readers, whose rebuttals equally overemphasized the psychological processing at work. See especially Corazza, 2011.)

Fourth, the previous argument showed that there are things which have both structure and justificatory relations holding between premises and conclusion, but not that the justificatory relations which hold between premises and conclusion do so *because of* the structure. In other words, the previous argument did not establish that the structure of the argument or its constituents was properly weight-bearing.

Fifth, and finally, the previous argument could easily be read as being addressed to a purely terminological issue: what gets to be called 'bearer of logical form'. But the revised argument makes the import of the issue obvious. It is manifestly relevant to argumentation theory, since it has implications for what kinds of things arguments are, and what kinds of things they can be composed of. It has bearing on issues in epistemology and philosophical logic, since it says something about the bearers of truth and justificatory power. It has metaphysical implications, since it commits us to things that Quine, for example, would have disdained: propositions, structured meanings (which are anathema to physicalists generally), and articulated mental states (which are anathema to traditional behaviourists, such as Bloomfield and Watson, to those—such as Stalnaker—who want only holistic contents as properties of agents, and to eliminativists such as the Churchlands). In contrast, proponents of neo-Russellian language-independent structured propositions such as Braun (1998) and Salmon (1991), and the many proponents of a language of thought, such as

Fodor (1975) and Sperber and Wilson (1986), will welcome the ontological import of our 'revisiting'.

3. Part Two: Objections and Replies

3.1 Summary of Part One

Recall our two guiding questions from Part One. We have been arguing that both ought to be answered in the affirmative—with logical structure being found both within the premises/conclusions themselves and across them. Our overall strategy has been to distinguish between the thing conveyed, on the one hand, and the content, at the context, of the symbolic expression used, on the other. We have applied this very familiar distinction in the philosophy of language to arguments, in order to uncover cases where the premise/conclusion conveyed outstrips the content of the corresponding expression-at-a-context. The particular cases in point have been sub-sentential arguments, both deductive (in the Alice-Bruce case) and inductive (in the Cara-Dan case). For instance, we have urged that the content at the context of the bare word 'red' is just the property RED. We maintain this because, or so it seems, this is what the word means wherever it appears, inside a sentence or out. But the premise in question was about the pen, to the effect that *it* was red. So, something other than the expression used (here, 'red') must be the premise—something which is not linguistic.

Sub-sentences aren't the only potential case in point, of course. There are other plausible ways to pursue the overall strategy. Likely, e.g., there are cases where a premise/conclusion may be conveyed with a gesture, a dance or a painting. That said, there is a singular advantage to appealing to sub-sentences in particular when it comes to the issue of the role played by logical structure. This is because sub-sentential examples are 'linguistic enough' that we are assured that there is a logico-linguistic structure (contrast a painting or a dance); at the same time, there is too little present for the premise/conclusion to simply be the linguistic item (even as interpreted relative to contextual parameters).

3.2 A Worry about Logical Structure

We can distill Part One down into a two-premise argument:

> **P1:** There are sub-sentential arguments in which at least one premise or conclusion is not itself a linguistic expression.
> **P2:** In at least some of these arguments, logical structure plays a weight-bearing role.
> **C:** There are structure-involving arguments whose elements are not expressions (whether of artificial or natural language).
> Thus, Vernacularism is false.

We will defend its two premises in reverse order. The reason is expository. Our main focus in this paper is on the nature of arguments, and P2 concerns that very directly.

The worry we'll address is how one can be sure that logical structure really is present. We will provide two tests—thereby defending P2, and at the same time clarifying further what we intend by the claim that 'logical structure plays a weight-bearing role'.

A first way to respond to the worry about P2 is to translate the argument that contains a sub-sentential premise or conclusion into a symbol system in which logical structure is obviously 'doing work'. This does not mean: translate each of the linguistic items. That wouldn't get us the right result, because (if we are right) some of those items are not fully propositional. Rather, the response involves transforming the *content* of each premise and conclusion, i.e. what is conveyed by the speaker, into symbols.[1] If we do this in the Alice-Bruce case, what we get as the translation is something along the lines of:

P1: Red(a)
P2: (x)Red(x) → Coloured(x)
C: ($\exists x$)Coloured(x)

We find, in the translation, logical structure both across the premises (P1 serving as a constant-involving instance of the antecedent of P2) and within the premises (e.g. existential generalization in C of the property appearing in the consequent of P2). Insofar as there is both 'macro level' and 'micro level' logical structure in the translation, there is logical structure in the original. QED.

We ourselves find this first test very compelling. We think it adequately addresses the concern about whether there is weight-bearing structure at play in sub-sentential arguments. But others might worry: how do we know that we haven't illicitly 'read in' logical structure by proposing a particular kind of translation? Compare 'translating' a dance or a diagram into the predicate calculus: does the possibility of such a translation genuinely demonstrate that there was structure implicit in the original?

We have two responses to this concern. Our first reply is that it applies too broadly, since it would call into question whether ordinary language arguments in general, including fully sentential ones, have weight-bearing logical structure. For even there, we frequently need to translate into an artificial language to 'capture the logical structure', and there is nothing obviously problematic about doing so in sentential cases. But a second reply is also available. It relies on the fact that this very specific worry—i.e. about how structure might get smuggled in via translating—cannot arise if translation isn't involved. So, very briefly, let us consider a second test by which we may identify logical structure.

[1] Looking back at the quotation from Brandom (2000) with which the paper began, we can now see that it has both a correct and an incorrect reading. If 'cannot serve as premises or conclusions' is taken to mean that *what sub-sentential words/phrases express relative to a context* is inapt, this is true; if it is taken to mean that *people cannot use plain-old sub-sentential words/phrases to put forward premises/conclusions*, this is false.

The second test involves showing that the argument in question is an instance of the right kind of template. That is, it involves showing that the argument exemplifies a schema with articulated structure, whose place-holders can be filled in by things of the right logico-semantic kind (e.g. Montagovian types). To explain what we have in mind, consider an example of such a template:

P1: α(*a*)
P2: EVERY-α(β)
C: SOMETHING(β)

The first premise-schema involves some object, of logical type <e>, going together with some function of type <e,t>. The resulting whole in P1 is thus of type <t>. The second premise-schema involves a second-order function of type ≪e,t>,t > going together with some function of type <e,t>. Specifically, this second-order function outputs TRUE just when, to put it crudely, β contains α. Again, the resulting whole is of type <t>. Finally, the conclusion-schema has this same quantificational form, but such that the second-order function outputs TRUE given as input some first-order function which itself sometimes has TRUE as output.

It is no accident that we have provided this template as an example, since the Alice-Bruce example is precisely an instance of it. That specific argument can be got by inserting RED for the schematic place-holder α, the pen for the schematic place-holder *a*, and COLOURED for the schematic place-holder β. We must stress that we are not proposing that the predicate 'red' substitutes for a schematic letter α, with a constant name of some sort substituting for a schematic letter '*a*'. Instead, what has the logical structure in P1 is the combination of RED, the *function* of type <e,t>, with the pen itself. The latter *object* is of type <e>. Our thought is that the functions RED and COLOURED are supplied in the familiar way—they are the meanings of linguistic symbols—but that the argument to RED works differently. It, the pen, is supplied directly by the utterance context, not by any symbolic expression.

To drive the point home, consider another instantiation of this template. Suppose Alice wants to argue that some external objects have tastes. She shows Bruce a slice of lemon, and has him place it in his mouth. She says, 'Bitter. Right? Bitter things have a taste. Right? So, there is at least one thing which has a taste. This thing'. This has the same logical shape as the original Alice-Bruce argument. What's more, the parts are contributed in the same sort of way. The instance of α(*a*) comes from putting together the meaning of 'bitter' with the object displayed, the slice of lemon. The instance of EVERY-α(β) is encoded as a matter of standing semantics in the English sentence 'Bitter things have a taste'. And the instance of SOMETHING(β) comes from filling in the function HAS-A-TASTE for the place-holder β. (By the way, this template in fact provides an alethic guarantee, in the sense that whatever objects and functions one puts in for *a*, α and β, the result will be such that if the premises are true, the conclusion must be true. But, we repeat, our conclusion is established even by non-linguistic structure-involving arguments that are merely inductive.)

We have simultaneously expanded upon P2 and defended its plausibility by providing two tests for structure-involving arguments. The first involves translating

them into a format in which said structure is evident on the surface: if the structure is there in the translation, it is there in the original. The second test involves establishing that the argument at issue fits a template in which logical structure is at work. It's worth underscoring its most controversial aspect. The notion of 'instance of a template' that we are invoking here is not the one favoured by Quineans, who would have particular *expressions* substituting for schematic letters. As we are conceiving of this test for logical form, what gets 'substituted' in the logical schema are objects, functions, etc. We have urged that the template-test yields the same result: there isn't merely an argument at work in the Alice-Bruce case (and many others), but one where structure plays a weight-bearing role. (See Stainton, 2000, for related discussion. Patently, this way of thinking conflicts very directly with Vernacularism. That is why, if this is a correct way of conceiving of things like the Alice-Bruce argument, Vernacularism is false.)

3.3 A Worry about Sub-Sentential Assertion: Thompson

Having addressed P2, it remains, in the final part of this paper, to consider objections to P1, according to which there are sub-sentential arguments in which at least one premise or conclusion is not itself a linguistic expression. There are many doubts that we will not address. Some have been dealt with at length before (see in particular Stainton, 2006); and some would take us too far from our main topic of Vernacularism.[2] We'll focus, instead, on four recent papers by Eros Corazza (2011), Jason Merchant (2010), Stefano Predelli (2011) and Robert Thompson (2011).

As an important exegetical caveat, we note that the 'objections to it' we discuss were not directed at the revised argument against Vernacularism: they predate the 'target argument' by several years. More than that, some of these authors seem quite open to articulated mental representations and structured propositions of precisely the sort that stand in conflict with Vernacularism. All the same, one can extract from their insightful papers some novel worries.

We begin with Robert Thompson. He denies that there are genuine sub-sentential assertions. Sub-sentential speech should not, he rightly insists, be assimilated to nudges and kicks under the table (*pace* Stanley, 2000). Nonetheless, insists Thompson, such speech doesn't rise to the level of full-blown assertoric content:

> I do not deny that sub-sentences can be used to communicate in a way that nudges, winks, and kicks under the table cannot; but this does not entail that they can, for example, be used to *assert* a proposition—they are too incomplete for

[2] There have been numerous discussions of Elugardo and Stainton, 2001, specifically, and of Elugardo and Stainton's views on sub-sentential speech in general. In particular, readers are encouraged to consult: Bach, 2008; Barceló, 2011, 2012; Davis, 2005; Harnish, 2009; Kölbel, 2010; Macherie, 2007; Purver, Cann and Kempson, 2006; Valmala-Elguea, 2007; Vicente and Groefsema, 2013; and Zanic, 2008.

> that. But, they can be used to suggest or convey a range of propositions. This sort of speech occurs all the time, but it is not full-fledged […] i.e., sub-sentential utterances are *degenerate* speech acts. […] Asserting a determinate proposition is a special sort of thing, and only an utterance with a certain type of structure can accomplish it. […] What I want to stress is that the sub-sentential utterances that are truly *sub*-sentential are the *wrong sort of item* to assert, ask, and order. They are inhibited by their incomplete syntactic structure, which limits the materials that are available for the speaker to *expect* the hearer to be able to develop into a determinate proposition (Thompson, 2011, pp. 72–3).

We can extract from these passages two senses in which Thompson finds sub-sentential speech 'degenerate' or 'not fully fledged': it is inexplicit (elsewhere, Thompson calls it 'fragmentary' and 'insufficiently guided by language') and indeterminate (elsewhere he calls it 'imprecise', 'coarse-grained', and 'sketchy'). And the worry for P1 is that, as a result, there isn't any genuine premise or conclusion in so-called 'sub-sentential arguments', since genuine premises or conclusions must be the sorts of things that can be asserted.

Our replies revolve around two points: first, Thompson's twin necessary conditions (of explicitness and determinacy) are too stringent; and second, sub-sentential examples (pretty much) meet those requirements anyway.

To begin, we of course grant that there is a sense in which sub-sentential speech is inexplicit. That is an essential part of our plaint. We equally grant that a requirement of full explicitness, as Thompson conceives of the latter, would scotch all of our examples. However, the requirement that a full-blown assertion of *p* requires that the content *p* be fixed by language-at-the-context is too demanding. It manifestly begs the question. In effect, the reason why Thompson denies assertionhood in these cases is because the thing used is sub-sentential (more generally, because the content isn't that of the linguistic expression itself, given the contextual parameters of speaker, addressee, time, place, etc.)

Thompson also requires that in order for something to count as an assertion it must be determinate. But this too is doubtful. On the one hand, it is again too demanding. Many things that strike ordinary language users as assertoric would be ruled out: 'She will soon' and 'They made enough' will often enough result in contents that are imprecise, coarse-grained, and sketchy—but seemingly asserted all the same. On the other hand, even if we grant Thompson this unduly high standard for genuine assertion, sometimes the demand will be met. Indeed, it is met in our recurring 'red' case. The premise therein is a singular proposition, about the pen, to the effect that it is red. And what could be more fine-grained and determinate than that? (Thompson expresses the hope that all those cases which even he must grant have fully determinate contents—the use of 'red' about the pen presumably included—will turn out to involve syntactic ellipsis. Stainton has argued at *great* length elsewhere that this is a vain hope, so we set it aside here.)

If determinacy and explicitness are not the right necessary conditions for genuine assertion, what are? In our view, the seminal issue is whether sub-sentential

assertions rise to the same standards as prototypical sentential ones. If they do, then sub-sentences can indeed be used to assert premises and conclusions. And what are those standards? Here are three, which we will call *reportative*, *phenomenological*, and *normative*. The reportative standard requires that assertions be naturally and correctly reported as 'S said that … ' not just 'S hinted that … ' or 'S led me to understand that … ' The phenomenological standard requires that speakers and hearers not experience full-on assertions as figurative, or hedged, or indirect. And the normative standard requires that assertions be lie-prone, and thus give rise to special moral and epistemic commitments, often including legal obligations. (Lie-prone here is to be contrasted with acts that are merely apt to mislead. See Saul, 2013, and Stainton, 2006, forthcoming, for more on that contrast.) Applying these three standards as the necessary and jointly sufficient conditions, there certainly are sub-sentential arguments.

So as not to put too much weight on the Alice-Bruce example, consider as our case in point the Edward-Fran argument. Fran owns a restaurant, and Edward works there. They say:

> *Fran*: I wonder whether we'll break even tonight.
> *Edward*: [Pointing at a different table each time he speaks] Reserved. Reserved. Reserved. Three tables reserved. Maybe four to six guests at each?
> *Fran*: Looking likely.

This little dialogue meets the reportative condition. One would naturally report it this way: 'Edward told Fran that three tables were reserved. In light of that, Fran said that it was likely that the restaurant would break even that night'. It would be flat out misleading to report their interchange as follows: 'Edward vaguely hinted something about tables being reserved' and 'Fran intimated something or other about likelihood'. The Edward-Fran argument also meets the phenomenological condition, since neither participant need experience the exchange as figurative or indirect. Paying attention to actual talk, one quickly realizes that there is nothing exceptional or noteworthy in such a dialogue: a massive proportion of our speech is sub-sentential in this way. (See Wilson, 2000, for a wealth of empirically attested cases.) Finally, the argument meets the normative condition. For, suppose that Edward knows full well that none of the tables is reserved, but indicates otherwise because he doesn't want to be sent home early. Given this, he lies to Fran each time he says 'Reserved'; he also lies when he concludes 'Three tables reserved'.

To summarize the discussion so far, Thompson's argument might lead us to reject P1 on the grounds that there is no genuine assertion in alleged sub-sentential arguments. On his view, the content of such alleged assertions is both too indeterminate and too inexplicit. Our reply had two parts. First, we argued that Thompson's standards are inappropriate: the requirement that content 'must be closely guided by the language' begs the question against the very possibility of sub-sentential assertion, while the requirement that content be determinate would rule out large swaths of perfectly ordinary and acceptable sentential speech. But second, we have argued that

even if Thompson's second standard is accepted, the requirement of precision can in fact be met by many examples of sub-sentential speech. Finally, when we apply more reasonable theory-independent criteria, the worry about whether there's a genuine assertion evaporates.[3]

Let us conclude with what is perhaps the most important rebuttal. Given what is at issue in the present paper, it is a mistake to think that a genuine premise/conclusion requires 'assertion' *in the sense that Thompson is contesting* (i.e. in the sense in which an assertion is to be contrasted with a truth-evaluable item that is merely presupposed, hinted or implicated). For, even if Alice fails to assert a premise by her use of the bare word 'red', she may all the same convey or implicate one. That, as noted previously, is enough to establish the truth of P1. As Elugardo and Stainton put it back in 2001: 'Whether [Alice] asserts the proposition, or merely implicates it, is not important for present purposes. What is crucial is that she communicates something that is correctly recognized to have structural entailments' (p. 403).

3.4 Another Worry about Sub-Sentential Assertion: Corazza, Merchant and Predelli

Consider now a series of objections to the effect that the premise/conclusion in (alleged) sub-sentential arguments is just a linguistic symbol of some sort.

3.4.1 Corazza. A recent paper by Eros Corazza (2011) affords a novel objection to P1. That paper has a number of conceptual ingredients, many drawn from John Perry, but the most important element for our purposes relies on a contrast between 'reflexive' and 'what is said' content (the latter being sometimes called the 'official' or the 'incremental' truth conditions).

According to Corazza, '[reflexive truth conditions] are inspired by Reichenbach's token-reflexive treatment of indexical expressions. They are associated to utterances in virtue of the meaning of the sentence (the type) uttered' (Corazza, 2011, p. 563). To explain with his example from that same page, the 'what is said' content of 'Jane smokes cigars' is something like $(\exists x)[cigar(x)\ \&\ smoke\ (j)(x)]$. But the reflexive content is different. The reflexive content of 'Jane smokes cigars' is something along the lines of:

There is an individual *x* and a convention *C* such that:

(i) *C* is exploited by this utterance of 'Jane smokes cigars',

[3] As an aside, note that being precise/determinate is not sufficient for being asserted either. The one issue is wholly orthogonal to the other. For example, if someone says 'Andy's car is parked at Betty's house', she may merely hint that Andy himself is there. But it isn't *ipso facto* indeterminate which proposition has been hinted. Or again, suppose that a teacher is prohibited from telling his student what letter grade she received. He says, 'I am not allowed to come right out and state what grade you will receive, but did you know that all apples always appear around April?' The proposition that she scored an A is definitely not asserted in this example, but it is as precise as can be. (More recently, Lepore and Stone, 2015, commit a related conflation between the precise/determinate on the one hand and the conventional on the other.)

(ii) *C* permits one to designate *x* with 'Jane',
(iii) *x* smokes cigars.

In (i), the very utterance 'Jane smokes cigars' appears. So, the utterance itself is a constituent of this kind of truth condition. Hence, 'reflexive'.

Why have this second kind of linguistic content? Largely because its cognitive role is markedly different than the singular proposition which is 'said': e.g. reflexive truth conditions can be understood by someone who doesn't know which Jane is being discussed, and grasping the content can yield very different behavioural proclivities. Importantly for Perry's philosophical project, e.g., although the two have the same 'what is said' content, one can learn something from the reflexive truth-conditions of, say, 'I am John Perry' that one cannot learn from those of 'John Perry is John Perry'—the reflexive truth conditions of the former being something along the lines of 'The person who is speaking this utterance is called 'John Perry' (see Perry, 2001 and Kepa and Perry, 2011.)

So far, we have a distinction between 'what is said' truth conditions and 'reflexive' truth conditions. How might this machinery allow one to reject P1? Drawing on Corazza's proposals for sub-sentential speech generally, the idea would be that 'red', as uttered by Alice, also has reflexive truth conditions associated with it, those truth conditions being something like[4]:

There is an individual *x* and a convention *C* such that:

(i) *C* is exploited by this utterance of 'red',
(ii) *C* permits one to designate *x* as the object discussed,
(iii) *x* is red.

Again, since 'red' appears in (i), this utterance of 'red' is itself a constituent of the reflexive truth conditions for 'red'. Moreover, the conjunctive claim (i)-(iii), which is part of the linguistic content, supports the conclusion of the Alice-Bruce argument that something is coloured, given the other premises. Hence, 'red' itself could be the linguistic item that has the relevant truth conditions, provides the justificatory support, and affords the requisite logical structure. But if that's the case, then the Alice-Bruce example does not, in fact, provide support for our claim that (some of) the constituent elements of arguments are not linguistic expressions.

This is an interesting line of thought. All the same, we think that it misses the mark. We will highlight three sorts of problems: an internal tension which Corazza faces; a general problem with the framework; and certain explanatory failings as applied specifically to sub-sentential speech acts.

Corazza wants to say that sub-sentential speech is genuine, and not elliptical. Indeed, he begins his paper by underscoring the reality of sub-sentential speech:

> Perry (1994) and Stainton (2006) convincingly argue that a single word can be used to make a full-fledged speech act. In general, subsentences such as 'Shoot

[4] See Corazza, 2011, p. 574ff. We have modified his formulation to sidestep some of the aforementioned 'conceptual ingredients', which are not apposite to the present discussion.

> when ready', 'On top of the hill', 'Great player', 'Nice dinner', 'John again', etc., can be used to perform speech acts without their being elliptical. I think this view is basically correct (Corazza, 2011, p. 562).

Corazza reinforces the point later: '[i]t is uncontroversial that speakers can utter ordinary words and phrases in isolation and, in so doing, perform full-fledged speech acts. The examples are limitless' (p. 570). However, despite himself, Corazza ends up denying that genuinely sub-sentential assertion occurs. Instead, on his view, some form of ellipsis is going on. It is not syntactic ellipsis, admittedly, but is rather a kind of semantic ellipsis: on Corazza's view, although we have what sounds like a plain-old word, it is in fact synonymous (on one of its readings) with a compete sentence. Recall: the utterance of 'red' has, as a matter of linguistic content, reflexive *truth conditions*. If one digs beneath the surface appearances, in other words, what Corazza is really proposing is that, for the one sound-pattern (presumably an adjective syntactically) there are two logico-semantic types: <t > when the sound-pattern appears unembedded and <e,t > when the sound-pattern appears embedded. (Take 'reserved' as another example. When unembedded, the adjective is alleged to be (roughly) synonymous with the complete sentence 'There is an individual *x* and a convention *C* such that *C* is exploited by this utterance of 'reserved', *C* permits one to designate *x* as the object discussed, and *x* is reserved'. But this clearly is not what 'reserved' contributes when it is embedded in an expression such as: [$_{S}$ That table is [$_{Adj}$ reserved]]. We will revisit the notion of semantic ellipsis when we discuss Merchant.)

The approach also has more general problems. To begin with, this Perry-inspired account is too conceptually demanding to be what humans in general do when they use and understand language. This is because it requires everyone to master concepts like 'speaker', 'utterance', 'linguistic convention', and the like. (Look back at the alleged reflexive truth conditions for 'Jane smokes cigars'.) This might be fine as an instrumentalist 'rational reconstruction', but we are here asking about what really goes on when something like 'red' or 'reserved' is used to make an argument.

Second, and relatedly, according to Corazza it is supposed to be the case that 'one grasps an utterance's reflexive truth conditions by being linguistically competent' (Corazza, 2011, p. 563). That is, it is an essential aspect of his (and Perry's) approach that reflexive contents are linguistic meanings of a sort: *linguistic* because they are part of what a competent speaker knows; *meanings* to account for their rationalizing cognitive role. This contention, however, is poorly evidenced. To illustrate the problem, let's revisit Corazza's central example, and note his description of how the apparatus works. Says Corazza:

> [Some of the propositions associated with an unambiguous sentence] are reflexive or utterance-bound, with the utterance itself as a constituent. These variegated contents expand from the purely reflexive to the official content: they constitute a family of gradually less reflexive and more incremental propositions.
>
> To highlight this, let's suppose that Jane utters (2) while John utters (3):

(2) I love cigars
(3) Jane loves cigars
The proposition expressed, the official content, by (2) and (3) is the following:
(2/3) a. That Jane loves cigars
(2) and (3) express the very same singular proposition with Jane as a constituent. The official content is the same. This, though, doesn't help in explaining why, in normal circumstances, if Jane wanted John to offer her a cigar, she would utter (2) but not (3): John may even be unaware of Jane's name. Even if John knew that the name of the person in front of him is 'Jane' it would be nonetheless bizarre for Jane to utter (3).

The reflexive content of an utterance like (2) corresponds to what a competent hearer would understand, given her knowledge of English and no other contextual information besides the fact that (2) has been produced. This can be rendered by the following proposition:
(2) b. That the speaker of (2) loves cigars
(2b) is a reflexive proposition for it has (2), the utterance itself, as a constituent. It is, therefore, a singular proposition having an utterance as a constituent. This, though, is not the proposition expressed. It isn't what Jane said in uttering (2). Yet the grasping of it plays a crucial role in Jane's plan (Corazza, 2011, pp. 564–5).

Our point is: there is no reason to think that the alleged reflexive truth conditions, if such there be, are in fact part of the literal linguistic meaning of utterances of 'Jane smokes cigars'. To be sure, sophisticated hearers can glean such token-reflexive facts on the basis of some general-purpose inference, if pressed. But that doesn't make such token-reflexive facts part of the linguistic content: hearers can glean lots of meta-linguistic truths about any given utterance that aren't part of their linguistic meaning. (Elugardo stresses just this point in his 2013 paper.) As Grice (1975) taught us, what hearers can inferentially arrive at is a very poor guide to what a given term means in a language. For example, a typical hearer of an utterance of 'You like small dogs and cats' can recognize that a second person pronoun has been uttered; that this English pronoun can be both singular and plural; that one of the words used may be modifying either the plural word 'dogs' or the phrase 'dogs and cats'; that a natural kind term was spoken; that it can be used for both *Felis catus* and for the large family to which this species belongs. All the same, it's absurd to build information of this kind into what an utterance of 'You like small dogs and cats' semantically encodes in English. In short, we don't actually have any good reason for building 'reflexive truth conditions' into the linguistic content of utterances of unembedded words and phrases.

Why does this complaint about the general framework matter for our purposes? It matters because if the 'reflexive truth conditions' are not part of what *the utterances* 'red' and 'reserved' mean, then the relevant premises in our target arguments cannot be the word (or its linguistically encoded meaning). It is worth pausing to emphasize this. The point is that Vernacularism is not rescued if the reflexive truth conditions

are merely conveyed by the speaker and inferentially gleaned by the hearer—given that the thing used, the word 'red', encodes only a property, even relative to the contextual parameters. As noted when discussing Thompson, if the (alleged) reflexive premise is speaker-meant, and is hence not itself a linguistic symbol, then that renders P1 true.

A final objection. Corazza proposes as a great advantage of this stance that the reflexive content on its own suffices for fixing the truth conditions of sub-sentential speech acts. This is supposed to render his Perry-style approach more simple and elegant than that defended by Elugardo and Stainton, with the latter's appeal to free pragmatic enrichment. But, in fact, his reflexive truth conditions would not suffice to account for sub-sentential speech acts. Pragmatics is still required to settle the illocutionary force of a given utterance. For example, 'reserved' can be used in a range of speech acts: to make the table reserved, to command that someone else list the item as reserved, and so on. Similar considerations apply to the word 'red' and to sub-sentence uses generally. The most important upshot is that once we let pragmatics in to fix illocutionary force, those pragmatic resources will also fix what proposition the speaker means, without the need for linguistically-supplied 'reflexive truth conditions'. The more parsimonious approach, then, is to allow that 'red', inside sentences, but also when unembedded, simply means RED (and 'reserved' simply means RESERVED, and 'Eloise' simply means ELOISE, etc.)

In sum, Corazza's suggestion faces three problems. First, because he implicitly treats 'red', when unembedded, as having truth conditions, he must abandon his own claims that sub-sentential speech does not involve any kind of elliptical sentence. Second, there aren't good reasons to think that (alleged) reflexive truth conditions, associated with utterances of 'red', are genuinely part of their linguistic meaning, as opposed to merely being information which hearers can glean. And third, reflexive truth conditions *qua* meanings are explanatorily idle when it comes to sub-sentences, given that pragmatic machinery is required anyway, and can carry the entire explanatory burden.

3.4.2 Merchant. Let us take stock. We distilled down Part One of our paper into a two-premise argument:

> **P1:** There are sub-sentential arguments in which at least one premise or conclusion is not itself a linguistic expression.
> **P2:** In at least some of these arguments, logical structure plays a weight-bearing role.
> **C:** There are structure-involving arguments whose elements are not expressions (whether of artificial or natural language).

In Part Two, we have been considering ways to resist this 'revisited' anti-Vernacularist argument. We defended P2 by offering two tests for structure-involving arguments, and urging that both applied to the Alice-Bruce example. As for P1, it was first resisted on the grounds that twin necessary conditions for assertion, namely precision

and explicitness, were not met. P1 was next resisted on the grounds that the (alleged) sub-sentence actually does have linguistically-supplied truth conditions, specifically reflexive truth conditions. We have argued that neither objection to P1 holds up under scrutiny. In what remains, we discuss two more attempts to discredit P1.

We begin with one which draws on some ideas from Jason Merchant. Merchant (2010), despairing of a syntactic ellipsis strategy that would apply widely enough, presents a variant on semantic ellipsis for most sub-sentential assertions. (See especially Culicover and Jackendoff, 2005, on Bare Argument Ellipsis and Chapters 5 and 6 of Stainton, 2006, for the grounds for despair.) What semantic ellipsis approaches all share is the idea that a seeming word or phrase (e.g. 'red' or 'on the door handle'), though not syntactically a sentence, is nonetheless of type <t > when it occurs unembedded. As in Corazza, this would allow one to rebut P1 because 'red' *et al.* express propositions when they occur in isolation (i.e. not embedded in any sentence); therefore, those very expressions can be the linguistic item which bears logical form. Merchant's suggestion is along these lines:

> There is a way of construing the semantics of [phrases that seemingly pick out individuals, properties and quantifiers] which I believe is fully consonant with Stainton's points about their interpretations in context, but which makes use solely of commonly assumed, independently posited, semantic combinatorics. The basic idea is to let the semantic value of these expressions (what they 'say') include a variable over the relevant kind of object, and to let this variable receive its value in the usual way, namely by an assignment function … entirely determined by context in the Gricean pragmatic way (Merchant, 2010, pp. 172–3).

What is especially clever, as we'll see, is Merchant's wholly novel implementation of this semantic ellipsis strategy. We will introduce his version by means of contrasting it with two more familiar ones. We will then level some objections.

According to the most familiar and simple variant on the semantic ellipsis approach, when we seem to encounter an ordinary word used to make an assertion, there is actually a lexical ambiguity at work. Specifically, the alleged 'ordinary word' has both a predicational and a propositional meaning (in particular, of the kind found in context-sensitive sentences). To explain with an example, consider the relationships among 'The pen is red', 'It is red' and 'red' used to make an assertion. The first declarative sentence is non-controversially of type <t>. So, less obviously, is 'It is red'. It is assigned to this same logico-semantic type because an utterance of it at a context is true/false—depending upon whether the object, assigned to the free variable by the contextually-supplied assignment, is red. That is, using notation from intentional logic, 'The pen is red' translates to $red(p)$ while 'It is red' translates to $red(x)$. Both are sentences, not just in the syntactic but also in the semantic sense: the first closed, the second open. Turn now to 'red'. The key idea is that this expression has two lexically specified meanings. Inside a sentence, it translates to $\lambda x[red(x)]$. However, if it occurs in isolation, it has the same meaning as 'It is red': that is, it translates to $red(x)$.

This simplest version is patently unattractive, because it makes every word lexically ambiguous. For, performance restrictions aside, any word can be used to make an assertion. There is, however, a better variant on semantic ellipsis, one which will take us closer to the Merchant-inspired objection to P1. Goes the idea, predicates are, as a matter of their lexical semantics, univocally of type <e,t>. Crucially, however, they type-shift to <t > when they occur on their own, unembedded. Put otherwise, a context-sensitive slot is introduced by discourse context, thereby turning an ordinary predicate into an open sentence. We can again illustrate with a sub-sentential assertion using 'red'. Without type-shifting, 'red' is taken to simply mean RED. This is its sole entry in the lexicon. But if that word is used outside a sentence, this triggers type-shifting, and 'red' becomes synonymous with the intentional logic sentence *red*(*x*). (One might think that some kind of 'sub-sentential construction' triggers this. See Stainton, 2006, pp 89–92.) The overall result is the same as in the simplest variant: because the expression uttered is fully propositional (at a context), we get a truth-evaluable assertion. One gets this result, however, without positing massive lexical ambiguity.

There is any number of problems with this second variant, discussed in detail in Stainton's book. We will highlight one, about how type-shifting gets triggered. There are two obvious options, and neither is attractive. On the one hand, one can let the alleged 'sub-sentention construction' do it. That introduces a kind of grammatical construction which—unlike familiar ones such as the causative, which can create a di-transitive from an intransitive (e.g. 'John sneezed' becoming 'John sneezed a tissue across the table')—simply cannot embed in any larger syntactic environment. In addition, if the shifting is not triggered by the need to find a proposition meant but rather by an element of grammar, this predicts that all unembedded uses of sub-sentences will be propositional. But, patently, when a bare word appears on a list of ingredients, the word itself doesn't make a statement. (Additional examples of non-propositional uses of words and phrases appear below.) What's more, if the 'construction' builds the illocutionary force into the standing meaning, to explain why a full-blown assertion occurs, it further predicts that all such uses will be assertoric. On the other hand, then, the trigger could be pragmatic. But then the view looks unduly complex and *ad hoc*, especially when compared to Elugardo and Stainton's pragmatics-oriented story, according to which speakers are simply deploying a perfectly ordinary predicate in a creative way. The type-shifting is otiose, being driven by pragmatic mechanisms which are themselves sufficient to explain the observed usage.

This takes us, at last, to Merchant's variant on applying semantic ellipsis to (what appear to be) sub-sentential assertions. He too claims that predicates are of type <t > when they occur unembedded[5], and that type-shifting is required to explain

[5] It is worth reiterating the import of all of this for our paper. On a standard view, the lexical entry for the word 'red' pairs it with λx[red(x)]. This is more or less what Elugardo and Stainton (2001) take for granted. Given this, when 'red' occurs unembedded in the Alice-Bruce exchange, Alice

the contrasting behaviour of predicates within sentences versus unembedded. His alternative differs in one clever and very key way. Where everyone else assumes that predicates are of type <e,t > as a matter of their lexical semantics, Merchant proposes that they are of type <t>. In other words, the shifting goes in the opposite direction: from <t > to <e,t>. What's more, says Merchant, this shifting only happens when the predicate embeds in sentences, as required by compositional semantics.

Let us apply his idea to our recurring example. The idea would be that 'red' means the open-sentence *red(x)* in the lexicon—the same meaning that 'It is red' has. That word itself is true/false (though only at a context), rather than being true/false *of* things. This, therefore, is what 'red' means when it occurs unembedded. It is when 'red' embeds that it undergoes type-shifting, to $\lambda x[red(x)]$. This has clear advantages over more familiar variants of semantic ellipsis: there is no bizarre lexical ambiguity; at the same time, *ad hoc* type-shifting is avoided because the shift is forced in a frequently-attested way. There are plenty of routine cases where there is a type mismatch between sisters on a tree, such that combining them demands that one change its logico-semantic type. For instance, given an <e > -type meaning for the name 'John', and a general requirement that intersection can only apply to words of the same logical type, to get the meaning of 'John and several girls' the name must shift up to ≪e,t>,t>, only then forming a compound generalized quantifier. Embedding a predicate lexically marked as <t > would merely be a further case in point.

Another way to understand how Merchant's proposal intersects with our larger questions is to recall David Kaplan's (1989) approach to logical relations for context-sensitive sentences. For Kaplan, it is sentences at a context (in the 'narrow' sense of a short list of contextual parameters, such as speaker, addressee, time, place, etc.) that are true or false. And an argument is valid when, if the premises are true at such a context, the conclusion must also be true at the context. Given this, one can understand Merchant as providing a way to extend this familiar machinery to (alleged) sub-sentential arguments. Granted, the word-types 'red' and 'reserved' are not true or false, so those lexical items *per se* cannot be premises; but, the idea would go, 'red'/'reserved' at a Kaplanian context are truth-evaluable. Thus, the examples we have introduced provide no reason for modifying the received understanding of the domain of the logical, beyond what sentences like 'It is red'

produced something that is not itself truth-evaluable, even given a pragmatically-determined variable assignment. Hence, the thing that is true must be something else. This supports P1. On any of the foregoing semantic ellipsis approaches, in contrast, the premise just is a linguistic expression, namely 'red'. In particular, as Merchant writes: 'On this approach, then, there really are more "slots" to be filled … Here the variables are already there, as parts of the meaning of the items used. What their *values* are is determined by context … So this has precisely the same effect as Stainton's account in this way, since it is the pragmatics that does this. But it "semanticizes" the variables in a familiar way. The difference between this account and Stainton's is pretty tiny indeed: the only real difference is that by having the semantic "slots" in the meaning (semantic value) of the phrase uttered, they can all be type <t>, propositional' (Merchant, 2010, p. 177).

already demanded. What is important about Merchant's paper is that it provides a way of resuscitating the otherwise unpromising idea that 'red', 'reserved', *et al.* can be true/false. Unsurprisingly, we are ultimately sceptical. We will now put forward four objections.

As we understand the idea, Merchant's view entails that any predicate 'F' used in isolation will mean the open sentence F(x). This isn't correct, on two fronts. Take 'red'. On the one hand, there are uses of it on which the utterance does not mean a proposition *of the right F(x) kind*. A speaker could use 'red' on its own to mean that his favourite colour is red. This would be of the form $\lambda x[red(x)] = \lambda y[favourite\text{-}colour(y)]$ rather than $red(x)$. A speaker could equally use 'red' to mean that every pen is red. The form there would be EVERY-PEN($\lambda x[red(x)]$), not at all equivalent in meaning to 'It is red'. (These will not arise from type-shifting, by the way, because 'red' is not embedded, and hence it has no sister to coerce its logico-semantic type.) Why does it matter that not all propositional cases are of this form? Because a main advantage of Merchant's sophisticated variant of the semantic ellipsis hypothesis was supposed to be that it avoided positing unmotivated ambiguities. These cases threaten, however, to make a word like 'red' lexically ambiguous after all: it will have several <t > -type meanings, being synonymous not just with 'It is red', but also with 'Red is identical to that property', 'Red exists in that quantity' and (as other examples would show) 'Red has that property'. In contrast, if 'red' just means RED (as we maintain), then all propositional uses of it are explained pragmatically in the same sort of way, and lexical ambiguity is avoided. That's on the one hand. On the other hand, to echo a point made above, there are uses of 'red' on which the utterance does not mean a proposition at all. *A fortiori*, it isn't a proposition of the form F(x). Here is an example. A speaker says, 'Name something you hate'. The interlocutor replies: 'Red'. This referential use of 'red' exhibits a meaning that is not consistent with Merchant's semantics. To clarify, the point is not that the interlocutor fails to make any statement at all. The point, rather, is that he uses 'red' in a way comparable to: 'A: Name someone you hate. B: Steve'. Here it is very implausible to suppose that the proper name 'Steve' is itself of type <t>. A second related example. Georgina lays a bet that no linguists will attend a certain dinner party. Henri points at a place-setting and says 'Barbara Partee'. He then points at a chair and utters 'The editor of *Linguistic Inquiry*'. Even Merchant agrees that his 'limited ellipsis strategy' does not apply to these, which he dubs 'labelling cases'. He is happy to grant this because he is interested in what may be strictly and literally asserted with sub-sentences, and he takes such 'labelling' to involve less-than-fully-assertoric communication of a proposition by the speaker (Merchant, 2010, pp. 165–6). In the context of the present paper, however, the merely-convey/assert contrast is beside the point: as we noted when discussing both Thompson and Corazza, if a premise/conclusion is speaker-meant in a sub-sentential argument, then (assuming, as seems to us obvious, that what is speaker-meant isn't itself a specific linguistic expression), P1 is vindicated. Recognizing this, notice that whether or not he states them, Henri definitely conveys two premises: that Barbara Partee will sit at this place, and that the editor of *Linguistic Inquiry* will sit at that one. Combined with some additional implicit ones, these

premises support the conclusion that at least two linguists will attend the party. And that entails, of course, that it is not the case that no linguists will do so.

Merchant's proposal is ingenious. We have already seen, however, that there are cases which it leaves untouched. So far, these have involved bare predicates. Another class of lacunae involves bare lexical items of other categories. Recall, for instance, that both 'Eloise' and 'Exactly' were used to state a conclusion in the Cara-Dan argument. In a related vein, Edward could have used 'Three' to conclude that three tables were reserved. To successfully refute P1, given such examples, Merchant's apparatus needs to be extended beyond predicates to lexical items generally: all must be of type <t>. He is aware of the problem, and he thus generalizes the notion of λ-abstraction so that it applies, e.g., to proper names. Couching things in our terms, 'Eloise' would be of type <t>, but the name would type-raise to λP.P(eloise) to allow the proper combinatorics (see his p. 176). The trouble with such a generalization of Merchant's idea is that if all words are, as a matter of their lexical semantics, of the same type <t>, the mechanism for coercing type-shifting is lost. Consider, for instance, the sentence 'Every very red pen is expensive'. Lexical assertion, on this generalized version of Merchant's view, would yield something like:

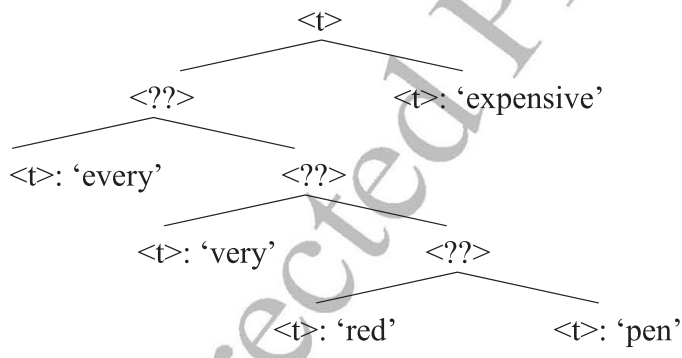

We are at a loss to see what independently motivated combinatorial rules would take these semantic parts and induce the requisite type-shifting.

Two objections issued, two remaining. Another set of examples which Merchant's proposal leaves untouched are phrasal ones. Again, Merchant recognizes that his account must extend beyond bare lexical items. We maintain, however, that once we think through the details, it becomes clear that it cannot reasonably do so. To understand how the objection arises, we need to fill in more particulars about how Merchant-style type-shifting would work. Patently, inside a sentence a predicate cannot (always) remain as type <t>. For instance, 'red' must often mean $\lambda x[red(x)]$ rather than $red(x)$, otherwise 'It is red' would mean something bizarre like [$_{<t>}$ It is [$_{<t>}$ it is red]]. Merchant therefore proposes that λ-abstraction occurs. To spell this out with some examples, consider how 'It is red' would work. Taking 'it' to be of type <e>, and setting the copula aside, the contextually supplied object <e > combines with the λ-abstracted $\lambda x[red(x)]$ of type <e,t>, created by

type-shifting, to yield the complete sentence's semantic value, which is of type <t>. Voilà. To take another example, in 'Every red pen is expensive', the bare quantifier 'every', being of type ≪e,t>,≪e,t>,t≫>, demands an input of type <e,t>. Both 'red' and 'pen' are assigned type <t > in the lexicon on Merchant's view, so this mismatch will trigger type-shifting of some sort. It might go like this: both bare predicates shift to <e,t>, and then a rule for intersective adjectives applies to yield an <e,t > -type for the compound 'red pen'. This, in its turn, serves as input to the meaning of 'every'. Subsequently, 'every red pen' takes the <e,t > meaning of 'expensive' as input. And the whole sentence comes out, correctly, as type <t >.

Now that the mechanism has been spelled out, consider how all of this will apply to bare phrases such as 'On the door handle' and 'Fresh red paint'. Such type-shifting will need to happen within such unembedded phrases, too. For instance, suppose Alice says 'A very red pen' in her argument with Bruce. As above, 'red' and 'pen' both need to shift to type <e,t>, to be combined together, and to thereafter compose with 'very', and ultimately with the bare quantifier 'a'. So, on numerous fronts, type-shifting away from <t > would be triggered within unembedded phrases, just as it is triggered within sentences. But this leads to a host of examples which escape Merchant's net: for instance, given how it can combine with 'is expensive', we know that 'A very red pen' must end up as type ≪e,t>,t>, rather than of type <t>. Hence, as in the original Alice-Bruce dialogue, the phrase used, even at a context, does not match the premise which Alice puts forward. To drive the point home, consider another example. Suppose that Alice had uttered 'Very red'. As we saw, 'red' cannot remain of type <t > if it is going to combine with this intensifying adverb, which is of type ≪e,t>,<e,t≫. But the result of combining the type-shifted 'red' with 'very' is not itself of type <t>, though what Alice means by it is. In short, what is needed to defeat P1 is that every complex, whether sentential or phrasal, end up as type <t>. Merchant's proposal, as it stands, cannot turn this trick. Rather, to achieve the requisite general result, including in phrasal cases where composition has required shifting up from <t>, one would need to always shift the complex back down to <t>. But then we have reintroduced the unhappy *ad hoc* and otiose step his proposal seemed to evade.

A brief recap is in order. The idea that we are drawing on, as providing a potential objection to P1, is nicely encapsulated in a lengthy quotation from Merchant himself:

> A common practice in natural language semantics is to assign λ-terms as the translation of lexical items … But this use of the λ-operator is not a necessary one. Imagine instead that λ-abstraction occurs in the course of, or as part of, the semantic composition, not as stipulated in lexical entries … On this view, then, λ-abstraction occurs as necessary to enable semantic composition, but not otherwise. It is a possible precursor to function application … The result of this view of the semantics is that predicates have a variable in them, but no λ-binder. When used in isolation, they will therefore have a free variable … What the values of these variables will be is determined by the assignment function. So Stainton is right that the pragmatics is crucial, and that our intuitions require

> that it be the context that determines what individual or property is used, but once we admit that the assignment function is responsible for 'slot-filling' of unbound variables, we already have in place the semantic mechanism needed (Merchant, 2010, pp. 174–5).

Now, this novel semantic ellipsis gambit initially looks superior to earlier variants, which were subject to twin worries about unmotivated lexical ambiguities and unmotivated type-shifting. Yet ultimately, we have urged, if Merchant's approach is to cover all the cases he would need to—e.g. where lexical predicates do not mean an open sentence of the form F(x), where non-predicates are employed, and in phrasal cases—the twin worries raise their heads again, albeit in modified form. So, his variant doesn't afford a means of rejecting P1 after all.

We end with a methodological challenge. It is not the default, the null hypothesis, that a word like 'red' is synonymous with 'It is red'. Some reason must be given for embracing such an odd proposal. But, we think, none has been. From our point of view, the proposal's motivation is merely to save the phenomenon. (One might be misled by notation here. If one takes the lexical entry for 'red' to be $\lambda x[red(x)]$, this makes the word's meaning look derived—from something more basic and sentential. Being philosophically scrupulous, however, one ought to take the intentional logic translation of 'red' to be *red*. Surely *that* is the null hypothesis.) What's more, there are good Fregean reasons *why* a < t > -type content should not be the default. We agree with Frege that a natural home of a lexical item is embedded in a sentence, and that when looking for the meaning of a word, it is essential to consider first-and-foremost how it behaves therein. We ourselves combine this with the unorthodox view that lexical items occur outside sentences, and that their meanings must also be consistent with this fact, but we grant to Frege that such truth-apt usage must be derivative in *some* sense: e.g. a heavy dose of pragmatics must be invoked. In contrast, Merchant seemingly takes the unembedded assertoric use of a lexical item to correspond to its fundamental meaning, with some ingenious hocus pocus yielding the non-<t > within-a-sentence meaning. We take his anti-Fregean point of departure to be a methodological misstep.

3.4.3 Predelli. A familiar rebuttal to alleged examples of sub-sentential assertions is that what the speaker really meant, in her so-called use of a 'bare word or phrase', was sentence such-and-such. For example, maybe what Alice intended was the sentence 'This is red'. As applied to our little two premise argument, P1 could then be resisted by holding that the sentence-meant could be the linguistic thing that has the truth conditions, stands in justificatory relations, has structure, and so on.

Of course, there is an equally well-known problem with appealing to a sentence-meant, which we have hinted at above: why pick that specific sentence as opposed to some closely related alternative? Why not 'That is red' or 'It is red' or 'This thing is red'? If we attempt to identify one particular sentence as the linguistic item that expresses the premise/conclusion, we get indeterminacy about which is 'the right one'. Frequently, that is, any particular sentence is too

fine-grained to capture the course-grained content actually intended. This had led to the widespread consensus that, in cases of speaker meaning, a speaker intends a content rather than a form.

Turning now to the final objection to our P1, Stefano Predelli makes a proposal designed to overcome this worry, while retaining the idea that 'logical properties are properties of *expressions*' (Predelli, 2011, p. 576). We'll explain Predelli's basic idea; show how it might avoid P1, if correct; and argue that it isn't a plausible account of how sub-sentences are used and understood.

A crucial notion here is regimentation (p. 573*ff*). Predelli portrays interpretation as moving from sound pattern to surface form to semantic input. Regimentation comes in at the second step because, in general, surface form underdetermines a particular sentence. Happily, he insists, there are regularities of speech behaviour, and features of 'wide' context (as opposed to the 'narrow' list of objective parameters such as speaker, time, place of utterance and addressee) which help fill the gap. Turning to sub-sentences in particular, these regularities yield a class of sentences that are truth-conditionally equivalent relative to the narrow context. Predelli's signal idea is that one can take their shared core to capture what is meant:

> ...the semantic input appropriate [in sub-sentential cases] is not a particular logical form, but, to coin a term, a *Constraint on Logical Forms* (*CLF*)... CLFs may be interpreted as classes of logical forms for any given [set of contextual parameters] c – in the simple case of a CLF mapping an object o to a node n, the class of logical forms mapping a singular term t to n iff t refers to o with respect to c (Predelli, 2011, p. 581).

How does this provide grounds for rejecting P1? According to Predelli's idea, the thing that has the logico-linguistic structure—the premise/conclusion—would not be the surface form (which, he grants, isn't fully propositional). Rather, the thing that has logico-linguistic structure would be the shared core, the CLF, associated with the surface form—itself determined by the set of sentences which would capture the truth conditions of what the speaker stated. And that shared core itself, that linguistic item, is fully propositional, given an appropriate context of utterance.[6]

Although there is much more to be said about Predelli's proposal, we have not spelled out all of its details here. The reason is that, despite its sophistication, Predelli's

[6] Here is another way of characterizing the approach, although it isn't a formulation Predelli himself offers. Think of superficial form as corresponding to something like a linear string of lexical items, not yet assigned syntactic categories or a hierarchical structure. For instance, take flying-planes-can-be-dangerous. This string needs to be correlated with a full-blown tree that will be input into compositional semantics, e.g. $[[_{NP}$ Flying planes$][_{VP}$ can be dangerous$]]$. That string-to-tree transition is regimentation. Predelli's key innovations are, first, his recognition that the larger discourse patterns can play a role in establishing the right correlation; and, second, his insistence that the input to semantics doesn't need to be a unique particular tree, but can instead be something like a class of trees, or more exactly a 'constraint' on trees met by all members of the class.

proposal falls prey to objections which plague less sophisticated set-of-sentences approaches. The first problem is that Predelli's view gets the order of explanation wrong. What makes it the case that such-and-such sentence is a close enough paraphrase of a speaker's use of a bare word is the proposition the speaker stated, not vice versa. As Elugardo and Stainton put the point in 2001:

> To give an analogy, there are two ways of understanding the location of the bull's-eye vis-à-vis an arrow. One might say, 'The bull's-eye is located just here because this arrow came very close, that arrow missed by a few inches, and that arrow was nowhere near'. One might also say: 'This arrow came very close because the bull's-eye is just here, and the arrow struck just there'. In the case of the bull's-eye, it's clear that the former gets the order of explanation wrong: the bull's-eye does not come to have its location because such-and-such an arrow missed. Now consider: in virtue of what is 'This thing is red' a quite good paraphrase of what Alice meant, whereas 'This doohickey exhibits red' isn't very good, and 'My plane is late' is nowhere near? Is 'the target' determined by which things come close and which are far away? Or is the closeness of the paraphrase determined by where the target is? We take the latter line. (And, we believe, only someone antecedently in the grip of Vernacularism would insist otherwise.) In which case, the paraphrases are not prior to—do not determine—the thing-meant (Elugardo and Stainton, 2001, p. 409).

That is, as applied to Alice's use of 'red', the CLF turns out to be fixed by the actual premise meant, rather than vice versa.

A related objection is that Predelli's proposal involves an unnecessary, and psychologically dubious, additional step. As Alison Hall (2008, 2009) shows, all such set-based proposals require the hearer to find the content meant by the speaker, and *then* to find a series of sentences that express that meant-content. Returning to our example, the hearer must first find the singular proposition, about the pen, that it is red, and must then find all of 'This is red', 'That is red', 'It is red', and so on. She subsequently extracts the 'shared core', which in turn serves as the semantic input. But this introduces an absurd epicycle: having once grasped what is meant, including even the relevant structure of the proposition meant, why look for sentences that express that proposition? The only motivation seems to be to save Vernacularism.

A third and final objection. Like its simpler kin, Predelli's sophisticated variant on the set-of-sentences approach is psychologically implausible in another sense. In general, human hearers simply do not find such sets, especially not ones including candidates for the relevant equivalence class such as 'This is red and $2 + 2 = 4$' (which is included by Predelli in the CLF because it does indeed get the truth conditions right). In particular, it does not fit with evidence from speakers who cannot process any sentences, but can process sub-sentential statements, such as adults who have cognitive deficits. To be fair, this may not move Predelli, who writes that 'more than a few topics of contention in the current debate, for instance having to do with communication, language understanding, and speech-act theory, are tangential to

my topic' (Predelli, 2011, p. 571). As he further highlights regarding his limited aims, 'the theory passes muster as long as it yields the intuitively correct truth-values for the content in question across points of evaluation' (Predelli, 2011, p. 581). But for those of us who are interested in an empirical account of the logico-linguistic facts, as opposed to instrumentalist modelling under conditions of restricted evidence, the facts about psychological processing carry serious weight (cf. Elugardo, 2007).

4. Conclusion

This has proven a very long and winding road. We end, therefore, with an overview of what we have and haven't done. The main aim was to revisit a debate about the domain of arguments—as initiated by Elugardo and Stainton (2001)—drawing on examples of sub-sentential speech. We have focused on two questions: Do premises and conclusions need to be linguistic expressions (natural or artificial)? More specifically, are there non-linguistic premises/conclusions in which logical structure plays a weight-bearing role?

The 'revisited argument' for a positive answer on both fronts can be condensed into two brief premises:

P1: There are sub-sentential arguments in which at least one premise or conclusion is not itself a linguistic expression.

P2: In at least some of these arguments, logical structure plays a weight-bearing role.

These were supported by a series of examples, whether exhibiting an alethic guarantee or lacking it. On that basis, we reiterated Elugardo and Stainton's original anti-Vernacularist conclusion—i.e. it isn't only symbolic expressions which have logical form—while overcoming a series of weaknesses in their original paper.

Part Two addressed objections of two sorts: doubts about whether the example arguments really are structure-involving, and doubts about whether there really is a non-linguistic premise/conclusion at work. We responded on behalf of P2 with two tests for structure which sub-sentential arguments meet: they translate into artificial logical languages in a way which highlights structure, and they instantiate structure-involving templates. The defense of P1 was longer and more various, but the seminal points pertained to: the appropriate criteria for being an asserted premise/conclusion; the fact that, given the debate at issue in the present paper, a premise/conclusion needn't be full-blown asserted in any case; the descriptive and methodological inadequacies of semantic ellipsis (whether it treats sub-sentences as having reflexive truth conditions, or as lexically of type <t > prior to type-shifting); and the descriptive and methodological inadequacies of set-of-sentence strategies, even when implemented by means of 'constraints on logical form'.

As noted, there is much that we haven't done. We did not offer a complete defence of the genuineness of sub-sentential speech acts. We left open the ontological status of the 'non-linguistic elements' in arguments—maybe they are structured propositions, maybe articulated mental representations, maybe something else again. And we did not attempt to show that there can be things with logical form in the complete absence of linguistic expressions. Even so, we take our 'revisiting' to have buttressed some important consequences of sub-sentential arguments for argumentation theory, philosophical logic, epistemology, and ontology.

Department of Philosophy
The University of Western Ontario

References

Bach, K. 2008: Review of Words and Thoughts. *Mind*, 117, 739–42.

Barceló, A. 2012: Word and images in argumentation. *Argumentation*, 26(3), 355–68.

Barceló, A. 2011: Sub-sentential logical form. *Crítica*, 43 (129), 53–63.

Brandom, R. 2000: *Articulating Reasons*. Cambridge: Harvard University Press.

Braun, D.M. 1998: Understanding belief reports. *Philosophical Review*, 107(4), 555–95.

Corazza, E. 2011: Unenriched subsentential illocutions. *Philosophy and Phenomenological Research*, 83(3), 560–82.

Culicover, P. and Jackendoff, R. 2005: *Simpler Syntax*. Oxford: Oxford University Press.

Davis, S. 2005: Quinean interpretation and the anti-Vernacularism view. In R. Elugardo and R. Stainton (eds.), *Ellipsis and Nonsentential Speech*. Dordrecht: Springer.

Elugardo, R. 2013: Sub-sentential speech acts, reflexive content, and pragmatic enrichment. In L. Goldstein (ed), *Brevity*. Oxford: Oxford University Press.

Elugardo, R. 2007: Minimal propositions, cognitive safety mechanisms, and psychological reality. In G. Preyer and G. Peter (eds.), *Context-Sensitivity and Semantic Minimalism: Essays on Semantics and Pragmatics*. Oxford: Oxford University Press.

Elugardo, R. and Stainton, R.J. 2001: Logical form and the vernacular. *Mind and Language*, 164, 393–424.

Fodor, J. 1975: *The Language of Thought*. Cambridge, MA: Harvard University Press.

Gauker, C. 2013: Inexplicit thoughts. In L. Goldstein (ed), *Brevity*. Oxford: Oxford University Press.

Gauker, C. 1994: *Thinking Out Loud*. Princeton: Princeton University Press.

Grice, H.P. 1975: Logic and conversation. In P. Cole and J.L. Morgan (eds.), *Speech Acts*. New York: Academic Press.

Groarke, L. 2005: Going multimodal: what is a mode of arguing and why does it matter? *Argumentation*, 29(2), 133–55.

Hall, A. 2009: Free enrichment or hidden indexicals? *Mind and Language*, 23(4), 426–56.

Hall, A. 2008: Subsentential utterances, ellipsis, and pragmatic enrichment. *Pragmatics and Cognition*, 17, 222–50.

Harnish, R.M. 2009: The problem of fragments: two interpretive strategies. *Pragmatics and Cognition*, 17, 251–82.

Kaplan, D. 1989: Demonstratives. In J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press.

Kenyon, T. 1999: Non-sentential assertions and the dependence thesis of word meaning. *Mind and Language*, 14(4), 424–40.

Kölbel, M. 2010: Literal force: in defence of conventional assertion. In S. Sawyer (ed), *New Waves in Philosophy of Language*. London: Palgrave Macmillan.

Korta, K. and Perry, J. 2011: *Critical Pragmatics*. Cambridge: Cambridge University Press.

Lepore, E. and Stone, M. 2015: *Imagination and Convention*. Oxford: Oxford University Press.

Macherie, E. 2007: Review of Words and Thoughts. *Notre Dame Philosophical Reviews.*

Merchant, J. 2010: Three kinds of ellipsis. In F. Recanati, I. Stojanovic and N. Villanueva (eds.), *Context-Dependence, Perspective, and Relativity*. Mouton: De Gruyter.

Perry, J. 2001: *Reference and Reflexivity*. Palo Alto, CA: CSLI Publications.

Perry, J. 1994: Davidson's sentences and Wittgenstein's builders. *Proceedings and Addresses of the American Philosophical Association*, 68(2), 23–37.

Predelli, S. 2011: Sub-sentential speech and the traditional view. *Linguistics and Philosophy*, 34(6), 571–88.

Purver, M., Cann, R. and Kempson, R. 2006: Grammars as parsers: meeting the dialogue challenge. *Research on Language and Computation*, 4, 289–326.

Salmon, N. 1991: *Frege's Puzzle*. Cambridge: MIT Press.

Saul, J.M. 2013: *Lying, Misleading and What Is Said*. Oxford: Oxford University Press.

Sperber, D. and Wilson, D. 1986: *Relevance*. Cambridge, MA: Harvard University Press.

Stainton, R.J. forthcoming: Full-on stating. *Mind and Language,* 31(4), 395–413.

Stainton, R.J. 2006: *Words and Thoughts*. Oxford: Oxford University Press.

Stainton, R.J. 2000: Objects and senses and substitutions. *Dialogue*, 39(3), 593–99.

Stanley, J. 2000: Context and logical form. *Linguistics and Philosophy*, 23, 391–434.

Thompson, J.R. 2011: Speech acts and sub-sentential speech. *Crítica*, 43(129), 65–91.

Valmala-Elguea, V. 2007: The syntax of little things. XVII Colloquium on Generative Grammar. Girona, Italy. June 13-15, 2007.

Vicente, B. and Groefsema, M. 2013: Something out of nothing? Rethinking unarticulated constitutents. *Journal of Pragmatics*, 47, 108–27.

Wilson, P. 2000: *Mind the Gap: Ellipsis and Stylistic Variation in Spoken and Written English*. Harlow, Essex: Pearson Education.

Zanic, J. 2008: Review of Words and Thoughts. *Prolegomena*, 7(1), 113–17.

AUTHOR QUERY FORM

Dear Author,
During the preparation of your manuscript for publication, the questions listed below have arisen. Please attend to these matters and return this form with your proof.
Many thanks for your assistance.

Query References	Query	Remarks
AQ1	Please confirm that given names (red) and surnames/family names (green) have been identified correctly	
AQ2	We would very much like to retain the full first names of authors if this is at all possible. I can explain at length why, if you are interested.	