7

Prudent Semantics Meets Wanton Speech Act Pluralism

Elisabeth Camp

1. SEMANTIC METHODOLOGY

A cautious semanticist will be as minimalist as she can. A minimal theory is lean, clean, and easy to comprehend and implement. Each new parameter we introduce into the theory diminishes these virtues, complicating the lexicon and increasing the theory's computational burden. By these measures of virtue, Cappelen and Lepore's semantic theory verges on the saintly. They maintain that only a highly restricted class of expressions, like "I", "today", "this", "actual", and "local", are context-sensitive; they claim that all other expressions, including "every", "tall", and "good", which are widely believed to be context-sensitive, make the same contribution to every sentence in which they occur.

However, theoretical simplicity comes at a cost. As a matter of empirical fact, distinct utterances of what appear to be the same sentence can communicate dramatically different contents in different contexts. We have to explain somehow how speakers manage to converge on common interpretations of these utterances; and the less of this explanatory work is done by the semantics, the more must be shouldered by the pragmatics. To the extent that one shifts the explanatory burden onto the pragmatics in this way, pragmatics becomes the locus of all the interesting action.

Further, the more one does this, the more indirect the empirical support for one's semantic theory becomes. Our evidence for semantic theorizing must ultimately come from the utterances that ordinary speakers make and the interpretations that ordinary hearers assign to them. However, we never find utterances with meanings sitting around on their own; they are always embedded

Thanks to John Hawthorne and François Recanati for discussion, and to an anonymous reviewer for helpful comments.

within a more complex structure of other speech acts. As Cappelen and Lepore say, "We don't know what it is to have intuitions about truth-values of utterances as such. If we are asked to have intuitions not about what an utterance says, asserts, claims, etc., but just about its truth-value, we are at a loss" (2005: 98).

This fact isn't a problem for a "maximalist", because he maintains that, with few exceptions, the intuitive content of a speaker's primary speech act—what he says, asserts, etc.—just is the semantic content of the uttered sentence. The primary challenge for him is to construct a principled, systematic, and plausibly realizable account of how semantically encoded meaning can produce the dramatically shifting contents we seem to find in different contexts. The minimalist doesn't face this latter challenge, because she rejects the "mistaken assumption" (2005: 53) that the semantics itself should explain ordinary intuitions about the contents of speech acts. Instead, the minimalist faces a different challenge: she must find some other, more indirect way to bring intuitions about speech act content to bear on her semantic theory, on pain of cutting it off from all substantive empirical constraints.

Thus, just like everyone else, Cappelen and Lepore must base their theory on intuitions about speech acts. Specifically, their proposal for isolating semantic content consists of three tests, all of which they admit rely upon intuitions about speech act content:

> These three tests all have the following form: *An expression* e *is context sensitive only if competent speakers have certain intuitions about uses of certain sorts of sentences containing* e. These tests appeal to fundamental features of linguistic communication. (2005: 87–8)

> We use communicated content . . . to 'get at' semantic content. . . . There is of course no other way to proceed. The purpose of the tests is to generate contexts in which semantic content is salient. . . . They are ways to get the audience to notice semantic features of sentences uttered. They create contexts in which our attention is drawn to features of the semantic content expressed by the utterances in question. (2005: 113; cf. also 122, 207)

Given that they employ this familiar methodology, Cappelen and Lepore are not entitled to the blanket claim that

> [I]ntuitions about, and other evidence for, speech act content are . . . not even prima facie evidence that p is the proposition semantically expressed by u. This is so no matter how refined, reflected, or 'equilibriumized' the intuition in question might be. (2005: 145)

As they say explicitly in the passages cited above, their tests are precisely means for focusing on or refining certain intuitions about speech act content in order to establish something about semantic content.

Because Cappelen and Lepore take the connection between speech act and semantic content to be so tenuous, it's especially important that the intuitions about speech act content that they focus on do accurately reveal semantic content. In §2, I argue that, as they stand, their tests fail to do this, because they suggest that metaphor is semantically context-sensitive. Although some

theorists do believe that metaphor is semantic, it is hardly a minimalist view. In §§3 and 4, I'll argue that the culprit is their Massively Permissive version of Speech Act Pluralism (MPSPAP). By ignoring important differences among various indirect reports and speech acts, and in particular differences between *saying* and *claiming*, MPSPAP makes our actual communicative practices seem mysterious and unmotivated. If we draw some salient distinctions among speech acts and reports, however, and restrict the application of Cappelen and Lepore's tests accordingly, then those tests become considerably more reliable.

2. THE TESTS

In this section, I argue that the way Cappelen and Lepore apply their proposed tests makes metaphor appear to be semantically context-sensitive. Many theorists would respond to this by dismissing metaphor, and non-literal speech generally, as an irrelevant distraction. I suspect that Cappelen and Lepore are no exception: they say, for instance, that it's crucial that their arguments against Moderate Contextualism not rely on "irrelevant" factors like non-literality, ambiguity, or vagueness (2005: 42).

However, such a dismissive response would be inappropriate. First, many of the people Cappelen and Lepore would call Radical Contextualists (e.g. Bezuidenhout 2001; Carston 2002; Hills 1997; Recanati 2001) believe that metaphor should at least be included within "what is said", if not within semantic content *per se*. And Josef Stern (2000), who counts as a moderate contextualist if not a minimalist by Cappelen and Lepore's standards, maintains that metaphor is semantically context-sensitive. Given that metaphor is a pervasive linguistic phenomenon whose theoretical status is contested ground, it's incumbent upon the minimalist to have some argument against those views. Further, Cappelen and Lepore themselves appeal to non-literal speech in support of SPAP, as in their example of "the moronic clown" (2005: 196). If non-literal speech can constitute evidence in favor of SPAP, then it should be relevant to discussions of semantic content as well.[1]

My general strategy in this section is as follows. Cappelen and Lepore invariably run their tests on intuitions about the truth of utterances *per se*, or else about the truth of indirect reports of what a speaker *said*. However, they should be equally happy to run their tests on what a speaker *claimed*, since they are committed to the view that 'say', 'claim', 'assert', etc. are all equivalent. And when we apply the tests to metaphorical utterances using 'claim' reports, the tests suggest that

[1] As officially formulated, Cappelen and Lepore's tests merely present necessary conditions on context-sensitivity. Thus, granting that metaphor passes the tests wouldn't force them to conclude that metaphor really is semantically context-sensitive. However, they regularly take the fact that an expression passes the tests to constitute positive evidence for its actually being semantically context-sensitive (cf. e.g. ibid. 88). The same standard should therefore apply to metaphor.

metaphor is semantically context-sensitive. By contrast, the evidence delivered by 'say' reports is considerably weaker. Cappelen and Lepore thus face a choice: either grant that we have evidence for metaphor's being semantically context-sensitive, or acknowledge a significant distinction between saying and claiming. At a minimum, they need to provide some explanation for the shift in intuitions produced by the different illocutionary verbs.

Test 1

> "An Expression is Context Sensitive Only if it Typically Blocks Inter-Contextual Disquotational Indirect Reports": "If the occurrence of an expression *e* in a sentence tends to block disquotational indirect reports (i.e., render such reports false), then you have evidence that *e* is context sensitive" (ibid. 88).

I claim that disquotational 'claim' reports of metaphorical utterances are blocked whenever the utterance and reporting contexts are relevantly different. Here are three examples, designed to mimic Cappelen and Lepore's discussion as closely as possible:

Example 1.1

> C: Alex says, to Bill: "You were right to dump Jane. She's one of those long-stemmed roses. She's overwrought and showy, and her thorns poke you if you try to get too close."
>
> C′: Charlie says, to Bill: "I think I'm in love with Jane. She's soft, fragrant, and always classy. She's a real long-stemmed rose." Bill says: "Alex claims that Jane is a long-stemmed rose, too." Alex rightfully objects that he didn't claim anything like what Charlie did; Bill has misreported him.

The previous evaluation of Bill's indirect report is situated within a real context, C′. Further, reflecting on both contexts from my own current (rather boring) context C″, in which no special assumptions about roses or romance are operative, my intuition is that all three of the following are false:

($1.1.1_C$) Alex claimed that Jane is a long-stemmed rose.

($1.1.2_C$) Charlie claimed that Jane is a long-stemmed rose.

($1.1.3_C$) Alex and Charlie both claimed the same thing: that Jane is a long-stemmed rose.

By Test 1, this should be evidence that metaphor is semantically context-sensitive. By contrast, however, utterances of the following all seem true to me:

($1.1.1_s$) Alex said that Jane is a long-stemmed rose.
($1.1.2_s$) Charlie said that Jane is a long-stemmed rose.
($1.1.3_s$) Alex and Charlie both said that Jane is a long-stemmed rose.

For instance, ($1.1.3_s$) is clearly true if followed by a clause like "but they meant totally different things by it". Admittedly, there are contexts in which ($1.1.3_s$) seems false: for instance, if it's followed with something like "so they agree about Jane". In such contexts, 'say' is being used in a way equivalent to 'claim'. The important point is that there is *a* use of 'say' on which ($1.1.1_s$)–($1.1.3_s$) are true, and that intuitions about 'say' reports are at least considerably more equivocal than intuitions 'claim' reports.

Example 1.2

C: Jane tells Charlene, "I dumped Bill because he's such a gorilla. I'm looking for someone who can hear people disagree with him without yelling and getting all hot under the collar."

C′: Alice and Charlene are watching *Gorillas in the Mist*, viewing footage of mountain gorillas lazing about and eating leaves. Alice says dreamily, "You know, Bill is such a noble, gentle force of nature, sometimes I think he's a gorilla." Charlene says: "Jane told me that Bill is a gorilla, too. Funny that you ladies agree on so much when you say you hate each other." Charlene has misreported Jane's utterance.

Again, here now in C″, it seems to me that all of the following are false, while the analogous reports involving 'say' are true or equivocal, as above:

($1.2.1_C$) Jane claims that Bill is a gorilla.
($1.2.2_C$) Alice claims that Bill is a gorilla.
($1.2.3_C$) Jane and Alice both claim that Bill is a gorilla.

Example 1.3

C: Cappelen and Lepore have a habit of calling anyone who accepts a semantic referential/attributive ambiguity a moronic clown, in the nicest possible way. When Jerome's new paper comes out, Cappelen tells Lepore, "It turns out that Jerome's a moronic clown too. I'm surprised, given his work on complex demonstratives. Still, his paper has the best arguments I've seen for the view—quite ingenious."

C′: Meanwhile, James is trying to write a letter of recommendation for Jerome, but he doesn't think Jerome's a good philosopher. Finally, he decides to just be honest. He writes, "You should hire Jerome if you're looking for someone to demonstrate to students how not to do philosophy. He can't construct a valid argument to save his life. Jerome is a moronic clown." Later, Bill is giving Alex a rundown on the latest gossip. Bill says, "Well, I heard that Cappelen claimed that Jerome is a moronic clown. And James claims the same thing, too." Bill is wrong: his reports are inaccurate.

Again, from C″, all three of the following seem clearly false, while the analogous reports involving 'say' are true or equivocal:

($1.3.1_C$) Cappelen claimed that Jerome is a moronic clown.
($1.3.2_C$) James claimed that Jerome is a moronic clown.
($1.3.3_C$) Cappelen and James both claimed that Jerome is a moronic clown.

How might Cappelen and Lepore deal with these cases while holding on to their test? The most obvious option would be to deny that metaphor counts as semantically context-sensitive according to Test 1, by insisting that cross-contextual disquotational 'claim' reports of metaphorical utterances, like ($1.1.1_c$) and ($1.1.2_c$), are ***true***. And I think we can find contexts in which such reports do seem true. As Cappelen and Lepore insist, though, it's not enough that there be some such contexts; the reports must be true even in contexts that differ dramatically from the original one. Thus, in particular, both ($1.1.1_c$) and ($1.1.2_c$) should be assessable as true in the same context, and ($1.1.2_c$) should follow automatically from them in that context. This seems highly implausible in its own right.

Further, anyone who does insist that such reports are generally true must decide how to reconcile this with our ordinary assumption that someone who claims *p* thereby commits themselves to obvious consequences of *p*. Can we conclude from the truth of ($1.3.1_C$) that Cappelen is committed to Jerome's being a moron? Can we conclude from the truth of ($1.2.1_C$) that Jane is committed to Bill's being a member of an endangered species? How do such reports function in chains of further reasoning: for instance, how do we assess the truth-value of statements like ($1.1.4_C$)?

($1.1.4_C$) If what Alex claimed is true, then Jane belongs to a species of flower that is cultivated in greenhouses and is more expensive than carnations.

Anyone who grants that ($1.1.4_C$) is true needs to explain how the consequent can be compatible with other claims that Alex presumably accepts, such as that

Jane is a woman and that women don't grow in greenhouses. Finding such an explanation will be challenging, to say the least. But on the other hand, if we deny that sentences like ($1.1.4_C$) are true, while still retaining the assumption that ($1.1.1_C$) and ($1.2.1_C$) are true, then we will be forced abandon the assumption that knowing that sentences like ($1.1.1_C$) and ($1.2.1_C$) are true tells us anything about what *else* the reported speaker might believe or have committed himself to. Given that nearly any word or phrase can be used metaphorically, and given that nothing in sentences like ($1.1.1_C$) and ($1.1.4_C$) indicates that the original utterances are metaphorical, it seems that we should *never* make claims about conditionals like ($1.1.4_C$) unless we are acquainted with the original context of utterance, because only such direct acquaintance would inform us of the inferential import of a speaker's claim. This seems like an excessively high cost to pay. Instead, it seems clear that we should deny that sentences containing 'claim' that disquote metaphorical utterances are true, and abandon the use of 'claim' reports in general as a test for semantic context sensitivity.

Test 2

> "Context Sensitive Expressions Block Collective Descriptions": "If a verb phrase *v* is context sensitive . . . then on the basis of merely knowing that there are two contexts of utterance in which 'A *v*-s' and 'B *v*-s' are true respectively, we *cannot* automatically infer that there is a context in which '*v*' can be used to describe what A and B have both done . . . it doesn't *follow* that there is a true utterance of 'A and B both *v*". (ibid. 99)

As formulated here, Test 2 relies on intuitions about utterances of the sentences 'A *v*-s' and 'B *v*-s' from distinct contexts, with the collective description located either in one of those two contexts or in a third. This certainly seems like the right way to run the test. Thus, consider a case involving a paradigmatically context-sensitive expression.

Example 2.1

C: George says, "Jim bought that car". (Pointing at a red Honda Civic)

C′: Charlie says, "Alex bought that car". (Pointing at a blue Lexus)

Supposing that Jim bought the car that George ostended, and that Alex bought the car that Charlie ostended, it doesn't follow that there is a true utterance, in C, C′, or my current context C″, of

(2.1.1) Jim and Alex both bought that car.

Nor can I legitimately pick up on Charlie's utterance in C′ by saying "And Jim did too."[2]

Stern (2000: 69) uses data of precisely this form to argue that metaphor is semantically context-sensitive. He claims that sentences like

(2.1.2) The largest blob of gases in the solar system is the sun, and Juliet *is, too.*

(2.1.3) Juliet is the sun, and Achilles *is, too.*

are "semantically ill-formed". I've argued (Camp 2005) against Stern's interpretation of the data, and specifically against the claim that we need to explain the badness of utterances like (2.1.2) and (2.1.3) semantically. However, Stern is undeniably correct that there is *something* true about appropriate utterances of "The largest blob of gases in the solar system is the sun", "Juliet is the sun", and "Achilles is the sun", and something very wrong with utterances of sentences like (2.1.2) and (2.1.3). This certainly seems like an application of Test 2. If we want to resist the conclusion that metaphor is semantically context-sensitive, we need to think more carefully about just how to apply Test 2.

The official formulation of Test 2 above appeals to the intuition that certain *utterances* are true.[3] If we use this criterion, then the test will support Stern's analysis. Consider the following example.

Example 2.2

C: Jane is an overwrought and showy woman, who only feels comfortable with people of her own socio-economic class, and who takes umbrage at anyone who disagrees with her. After enumerating these facts, Alex says: "Jane is a long-stemmed rose. I'd rather go out with a daisy, you know?"

C′: Kim is a relaxed but sophisticated lady. She has an effortless, impeccably elegant sense of style, a soft and demure way of interacting with others, and she always wears a lovely perfume. After lauding these virtues of Kim's, Charlie says: "I think I'm in love with Kim. She's a real long-stemmed rose. It's rare to meet such a class act."

[2] Cappelen and Lepore claim that for a sentence containing a paradigmatically context-sensitive expression, such as "Frank bought this, and Martha did too", "there are no available interpretations . . . on which Frank and Martha bought different things" (ibid. 101n. 7). Notice, however, that there can be considerable variability in what counts as *the same thing* for a phrase like "that car": say, the same token vehicle (a specific red Honda Civic), or two cars of the same specific type (red Honda Civics), or two cars of a quite general, contextually salient type (the fully loaded edition of a Japanese company's mid-level model). As a result, Frank and Martha might buy different things on some construals of "thing" but not others.

[3] Because Cappelen and Lepore deny that we have intuitions about truth-values of utterances as such, they aren't entitled to run Test 2 this way in any case; this is an independent reason to adopt the modified version I suggest below.

Like Romeo's utterance, Alex's and Charlie's utterances each seem intuitively true in their respective contexts. Further, it would be quite natural for someone who believed that Jane is overwrought and showy to respond to Alex's utterance by saying something like "You're right" or "That's true" (cf. Hills 1997); and it seems plausible that such comments pick up on utterance content—supposing that there is such a thing. However, when the truths of those respective utterances are used as premises to derive collective descriptions, then the conclusion is clearly false, as in (2.2.1):

(2.2.1) Alex's utterance of "Jane is a long-stemmed rose" is true. And Charlie's utterance of "Kim is a long-stemmed rose" is true. Therefore, Jane and Kim are both long-stemmed roses.[4]

Thus, if we frame Test 2 in terms of *utterances'* truth-values, metaphor appears to be semantically context-sensitive. However, when Cappelen and Lepore actually apply Test 2 themselves (2005: 100), they follow a rather different protocol. There, after presenting a true assertion of "Smith weighs 80 kg" in C and a true assertion of "Jones weighs 80kg" in C′, they shift to their own context C″, note that assertions of those very same sentences seem true in that context as well, and then show that we can combine the descriptions in C″. The analogous application to our metaphorical assertions from Example 2.2 gives us (2.2.2):

(2.2.2) Jane is a long-stemmed rose. And Kim is a long-stemmed rose. So Jane and Kim are both long-stemmed roses.

And indeed, I believe the syllogism in (2.2.2) is valid; it's just not sound, because the premises are both false. If that's right, then (2.2.2) doesn't suggest that metaphor passes Test 2.

However, construed in this way, Test 2 gives us little or no evidence about context-sensitivity beyond that already provided by Test 1. Once we've successfully disquoted "weighs 80 kg" from C and from C′ into the single context C″, the possibility of collection shouldn't be surprising. An expression like "today" can't be disquoted to relevantly different contexts; but for those context-pairs for which it can be disquoted, collection then follows automatically.

[4] It's true that we can say things like "Jane and Kim are both long stemmed-roses, albeit in very different senses"; perhaps the collective description in (2.2.1) and (2.2.2) can be heard as elliptical for this. But such a reading precisely *avoids* claiming that the collective description is true in virtue of any common properties that Jane and Kim possess. Thus, it doesn't show that we can "describe what they all have in common by using '*v*' " (2005: 99), as Cappelen and Lepore rightly require. With a gradable adjective like "tall", it's perhaps possible to insist that the Empire State Building, Osama bin Laden, and Billy the third-grader do all share a very general property: being tall. I'm even less of a metaphysician than Cappelen and Lepore, but I don't know what real property or common way of being Jane and Kim could share that would underwrite a substantive reading of "are both long-stemmed roses".

Therefore, insisting on using a disquotational formulation of Test 2 threatens to beg the question against a semantic theorist of metaphor like Stern: in order to apply the test at all, we must restrict our attention to contexts for which collection will go through.[5]

To capture what's true about the original metaphorical utterances in a way that doesn't suggest that metaphor is semantically context-sensitive, we need to appeal to more specific speech acts than Cappelen and Lepore are willing to do. Contrast the following pair of syllogisms:

(2.2.3) What Alex claimed (by uttering "Jane is a long-stemmed rose") is true. And what Charlie claimed (by uttering "Jane is a long-stemmed rose") is true. Therefore, Jane and Kim are both long-stemmed roses.

(2.2.4) What Alex said (by uttering "Jane is a long-stemmed rose") is true. And what Charlie said (by uttering "Jane is a long-stemmed rose" in context C) is true. Therefore, Jane and Kim are both long-stemmed roses.

In (2.2.3), the premises again seem clearly true, and the conclusion false.[6] In this case, however, unlike with (2.1.3) and (2.2.1), we can explain why the conclusion doesn't follow in a way that doesn't support semantic context-sensitivity. By uttering "Jane is a long-stemmed rose", Alex didn't claim that Jane is a long-stemmed rose, but rather that she possesses some set of properties *P* that happen to be saliently associated with long-stemmed roses in C. And by uttering the same sentence in his context C′, Charlie claimed that Kim possesses some other set of properties *Q* associated with long-stemmed roses in C′. There's no reason to think that the largely disjoint sets *P* and *Q* should be collectible under a common description, in C″ or any other context. I think this explanation is highly plausible. But it entails we shouldn't use 'claim' reports to test for semantic context sensitivity.

By contrast, I hear the premises in (2.2.4) as false. And if that's right, then metaphor doesn't generate a failure to preserve truth through collection, as

[5] Further, it's not obvious that cross-contextual disquotation *can't* go through for the premises in (2.2.2). I think we can, albeit with effort, hear the premises as true by importing the relevant assumptions for each sentence from its original context; this would be analogous to an utterance of "Jim bought *that* car and Alex bought *that* one" where the speaker mentally ostends distinct cars from memory. To the extent that we can indeed hear the premises in (2.2.2) this way, then the syllogism's conclusion is once again blocked, in the way that Test 2 requires for semantic context-sensitivity.

[6] One could deny, as Davidson (1978) does, that the speakers claimed anything true. But this view flies in the face of ordinary intuitions about metaphorical utterances and their 'claim' reports. In conversation, we do take speakers who have spoken metaphorically to have made speech acts with contents that we can hold them responsible for. Further, hearers can easily deny those claims using the same form of words as the speaker. For instance, within its context of utterance, the most natural way to deny that Jane is overwrought and showy is to utter the negation of the sentence Alex uttered (cf. Hills 1997; Bezuidenhout 2001). An adequate theory of speech acts must explain this.

Test 2 requires: we just have another valid but unsound syllogism. Again, there is undeniably an ordinary use of 'say' on which the premises are true; and in that case we again have a blocked inference, *à la* Test 2. But as I'll argue in §4, there is *an* ordinary use of 'say' on which the premises are false. Thus, if Cappelen and Lepore want a version of Test 2 that doesn't make metaphor appear semantically context-sensitive, they need to employ *this* use of 'say'; neither direct appeal to utterance truth, nor disquotation, nor 'claim' reports, will do.

Test 3

> "Context Sensitive Expressions Pass an Inter-Contextual Disquotational Test": "*e* is context sensitive only if there is a true utterance of an instance of the following schema for Inter-Contextual Disquotation:
>
> (ICD) There are (or can be) false utterances of 'S' even though S.
>
> (Alternatively, run the test in reverse). (Cappelen and Lepore 2005: 105)

As formulated, Test 3 again relies upon intuitions about the truth-values of utterances as such. But because this is not a way of talking that Cappelen and Lepore are entitled to, I'll also run Test 3 on what speakers claim by making their utterances. I reuse context C from 2.2.

(3.1.1) There can be true utterances of "Jane is a long-stemmed rose", even though Jane is not a long-stemmed rose.

(3.1.2) What Alex claimed (by uttering "Jane is a long stemmed rose") is true, even though Jane is not a long-stemmed rose.

(3.1.3) What Alex said (by uttering "Jane is a long stemmed rose") is true, even though Jane is not a long-stemmed rose.

My intuition, insofar as I have intuitions about the truth-values of utterances *per se*, is that (3.1.1) is true. (3.1.2) is, I think, obviously true, assuming that Jane really is overwrought and showy. Finally, I think (3.1.3) is false, or at least equivocal and considerably worse than (3.1.2). At least, then, we find here the same pattern of shifting intuitions depending on the particular verb employed; this again suggests, *pace* Cappelen and Lepore, that 'claim', 'assert', and 'say' are not all equivalent.

3. INDIRECT REPORTS AND ORIGINAL UTTERANCE CENTRISM

Insensitive Semantics is a book about semantic methodology, not about the semantics of specific expressions. Given this, Cappelen and Lepore could simply

stick to their guns, and take the discussion of their tests in §2 as *prima-facie* evidence that metaphor is semantically context-sensitive. However, most semantic theorists, and especially minimalists, would be quite surprised if metaphor did turn out to be semantic. I don't believe that it is semantic, and I think that a more general, independently motivated theory of the relation between speech act and semantic contents can explain why it is not. In §4, I'll argue for a distinction between asserting and saying, understood in a specific sense, which supports this view. In order to draw that distinction properly, though, we first need to constrain 'what is asserted' so that it more accurately reflects our broader communicative practices. In this section, then, I ignore the distinction between saying and asserting, and focus on getting a clearer view of saying/asserting/claiming, where "say" is used equivalently to "assert" and "claim".

As Cappelen and Lepore themselves emphasize (2005: 199), indirect reports are just as subject to "speech act pluralism" as any other utterances. Therefore, we must be extremely careful about assuming that an intuition that an indirect report communicates something true is driven by that report's semantic content. Recall, for example, context C of Example 1.2. Suppose that Tom had previously said, "Bill is incapable of delivering criticisms in a respectful and non-hurtful way." Then, in C, someone could felicitously respond to Jane's utterance by saying "Tom says that Bill's a gorilla, too." We wouldn't normally conclude from the acceptability of such a report that Tom himself said or claimed something with the semantic content of "Bill is a gorilla." So at least sometimes, it can be dangerous to follow Cappelen and Lepore's methodology of "tak[ing] our practice of indirect reporting at face value and assum[ing] that the speakers have said [or claimed] what we have the reporters saying that they have said [or claimed]" (ibid. 50).

As Cappelen and Lepore have argued forcefully (1997, 2005), our reporting practices are massively permissive across the board, and not just when it comes to figurative speech. We often accept indirect reports that communicate content which only follows from the uttered sentence's semantic content when that content is combined with quite substantive additional assumptions. Suppose that Cokie, a news reporter, utters (4.1.1):

(4.1.1) Today's nasty weather is likely to depress voter turnout, which will almost certainly favor the incumbent party.

Someone who'd heard (4.1.1) could easily accept the following report if she shared the relevant political assumptions:

(4.1.2) Did you hear that? She's saying/claiming that we're going to let a little rain prevent us from sending those corrupt hooligans back to their suburban McMansions where they belong!

Reports like these undeniably do constitute part of our normal communicative practices. The trouble comes from combining this fact with Cappelen and Lepore's emphatic rejection of Original Utterance Centrism (OUC) (2005: 201). Denying OUC places a report like (4.1.2) on a par with (4.1.1) itself as evidence about what a speaker actually said or claimed. But if Cokie were told about the report in (4.1.2), she would strenuously object to that characterization of what she said; and ordinary hearers who didn't share the specific assumptions operative in the reporting context would think that she was obviously right to do so.

The same thing can happen if we stick to less tendentious reports. Consider the following example:

Example 4.2

George says to Alex, "I've run into Jim coming out of the Bluebird Diner the last three Monday nights." Alex knows that the Bluebird is rented out to the local chapter of the Democratic Party on Monday nights, so when he sees Bill, he says "George says Jim's been going to the Democratic party meetings at the Bluebird the last few Mondays." Later, when Jerome mentions that the local Democrats' sole project is agitating to raise the city's minimum wage, Bill responds, "Apparently, George says Jim has been going to the weekly Democrats' meetings. Surprising: I always thought he was one of us." When Jerome sees Jim, he says, "The word from George is that you're a big 'living wage' guy." Jim confronts George, saying "How could you reveal my political leanings like that?" George objects that he didn't do or say any such thing; indeed, he didn't even know that the Democrats meet at the Bluebird. George is right to object: he didn't say anything like what Jim said he did.

Each alteration in the successive reports in 4.2 capitalizes on background assumptions that are common and obvious within the specific reporting context. As a result, each successive report would reasonably be accepted as true in its conversational context by a hearer who shared those assumptions and who had heard and accepted the utterance upon which it is immediately based (for instance, by a hearer of Bill's report who had heard Alex's utterance). If relevant features of the reporting context can be "constitutive of" the content of what was originally said, as Cappelen and Lepore maintain (2005: 206), then there should be no need for the hearer to know about George's original utterance. However, through a series of individually acceptable reports, we end up with a highly distorted report. The more links in the chain of 'Telephone', the less substantive the assumptions underwriting each shift in reported content need to be in order to end up with a wildly inaccurate report.

Although our ordinary reporting practices are often quite loose, the ultimate standard for accuracy is clearly fixed by the original context of utterance. In a legal context, for instance, hearsay about what someone said is not admissible as evidence: the court requires direct evidence of the original utterance. More prosaically, anyone on the receiving end of a chain of gossip who wants to figure out what was really said/asserted will find out as much as they can about the original utterance and its surrounding context. This privileging of the original utterance and context reflects the basic purpose of making utterances, which is to commit oneself to certain contents: paradigmatically, to contents that one believes, and wants others to believe. By making a claim, one puts oneself on record as accepting a certain content, and makes oneself liable for either defending its truth or else retracting one's claim (cf. Lewis 1979; Brandom 1983). The purpose of reporting utterances, in turn, is to transmit information about the contents to which speakers have committed themselves across contexts, so that others know what those speakers can safely be assumed to believe, and what they can legitimately be held responsible for.

If assumptions that are operative in dramatically different contexts could be constitutive of the content that I myself said/asserted, then it would be utterly pointless for me to say anything, because the content of what I said/asserted could always shift in ways that were unrecognizable and repugnant to me. Worse, speaking would be extremely dangerous, because I would be potentially liable for all those contents. Speakers can indeed unintentionally commit themselves to contents they're not aware of. But in such cases, the determinative factor is still what a reasonable hearer, located within the original context of utterance and armed with all the assumptions in play in that context, would take the speaker to have committed herself to.

At the very least, then, we need to restrict 'what is said/asserted' to contents that are recoverable by a reasonable hearer of the original utterance. Reports of what is said/asserted that don't meet this standard can still communicate something true, but they are not themselves true.

In fact, however, even this standard is too loose. Although speakers are usually quite open about their communicative intentions, in principle there can always be a significant gap between what a reasonable, charitable hearer would assume the speaker was saying/asserting and the content to which the speaker has strictly speaking committed herself. Genies and oracles, Jesuits in Elizabethan England, insurance companies, and politicians are all infamous for constructing their utterances in surgically precise ways so that they can legitimately deny the validity of what most hearers would assume to be accurate reports of their speech acts. Such utterances are certainly misleading, but we don't ordinarily count them as *false*. For instance, in a legal setting a speaker who knowingly makes such a statement is not liable for perjury if the content assigned on the broader interpretation turns out not to be true.

Condoleeza Rice's statements about US policy on torture offer a particularly blatant example of this. Many reasonable hearers took her statements that "The United States does not transport, and has not transported, detainees from one

country to another for the purpose of interrogation using torture", and that the US "will not transport anyone to a country when we believe he will be tortured", to be "unequivocal denunciations" of torture, and the Bush Administration clearly intended the statements to be taken this way. But others noted that rendition for the purpose of interrogation which just happened to end up involving torture, or rendition to a country where the US merely suspected that the detainee might be tortured, are not strictly speaking covered by these statements.[7] As a result, as evidence surfaces of people being tortured during interrogation after US rendition, Rice can legitimately respond that these are not counterexamples to her claim, so long as the relevant officials didn't believe the torture would definitely occur. Rice's original utterances were deplorably misleading; but if those stricter conditions are indeed met (something that is itself far from obvious), then it wasn't strictly speaking false, and hence not a lie.

Cappelen and Lepore deny that our ordinary linguistic practices support any such notion of "what was strictly speaking said/asserted". They claim that the only way to interpret questions about what someone "really", "literally", or "strictly speaking" said is as questions about whether the indirect report would be true considered as a *direct* quotation (2005: 51). But this is false. It's an essential, pervasive part of our ordinary linguistic and legal practices that we do have such a restrictive notion of speech act content, which is distinct both from direct quotation and from the more permissive standards that we more typically employ. Even Clinton had to admit that by uttering "I did not have sexual relations with that woman", he said/asserted that *he* did not have sexual relations with *Monica* (cf. Saul 1999). The assignment of reference to indexicals was never in doubt; the debate came over what conditions of satisfaction were required for them to count as "having had sexual relations", strictly speaking.

A strict interpretation of what someone said/asserted still leaves considerable room for speech act pluralism. In nearly any context, there are substantive, context-specific assumptions in play that even the most casuistic speaker cannot legitimately deny. These assumptions combine with the semantic content of the sentence uttered to produce commitments to multiple propositions that are syntactically and semantically unrelated to the uttered sentence. Thus, we still

[7] See e.g. Richard Bernstein, "Skepticism Seems to Erode Europeans' Faith in Rice", *New York Times* (7 Dec. 2005); see also http://lawofnations.blogspot.com/2005/12/rice-offers-legal-defense-of-rendition.html: "In parsing Secretary Rice's words, look closely at the specific intent or express knowledge that would be required for her statements to be false, and you'll see the role played by lawyers in the drafting of her speech. *The United States does not transport, and has not transported, detainees from one country to another for the purpose of interrogation using torture.* This careful statement is thus true so long as the purpose of an extraordinary rendition was just to interrogate a detainee, with the knowledge that he might be tortured, or was even likely to be tortured. So long as the purpose of the rendition wasn't that he be rendered for "interrogation using torture", Rice is being truthful. Deceptive, obviously, but truthful. . . . Again, this isn't the meaning being attributed to her by favorable commentators."

have a version of SPAP for assertion (and analogous speech acts, like promising).[8] But it is a version of SPAP that makes significantly better sense of our actual conversational practices than Cappelen and Lepore's MPSPAP does. It is also a version that is better suited for semantic theorizing. In fact, I suspect it supports a fairly minimal semantics; but like Cappelen and Lepore, my focus here is on methodology rather than particular cases.

4. WHAT IS SAID AND WHAT IS ASSERTED

The discussion in §3 relied neither on non-literal speech nor on the distinction between saying and asserting; the restrictions on SPAP I advocated there were motivated solely by the need to reconcile our analysis of indirect reports with our broader communicative practices. However, if we want to achieve a fully general understanding of those communicative practices, and in particular if we want to uncover a reliable connection between speech act and semantic content, then we also need to distinguish saying from asserting.

Although Cappelen and Lepore argue for a radical separation between speech act and semantic content, even they don't deny that there are *any* "interesting and informative connections between intuitions about" the two types of content (2005: 57); if they did, as we saw at the outset, then they would cut themselves off from all empirical evidence for their theory. Not only do their tests depend on such a connection; they also say that this connection is essential to explaining why communication isn't miraculous (p. 204). So what is the crucial connection between speech act and semantic content? They claim it is this: a speaker always asserts the semantic content of the sentence she utters, even if that content is typically swamped by a swarm of further propositions:

> The semantic content of a sentence S is the content that all utterances of S share. It is the content that all utterances of S express no matter how different their contexts of utterance. (p. 143)
>
> One of the many propositions asserted by an utterance is the semantic content of that utterance (the proposition semantically expressed). We argued for this in chapters 7–10. (p. 200)

The problem with this claim, however, is that the postulated connection between semantic and asserted content does not always obtain. Intuitively, we want to say, the speaker of a metaphorical utterance doesn't assert or claim her uttered sentence's semantic content. Romeo, for instance, doesn't claim that Juliet is *the sun*, but rather that she is a certain way *S*, which can be described metaphorically as being the sun. If Romeo were really claiming that Juliet is the sun, then he would also be committed to Juliet's being an enormous hot gaseous body. But

[8] For further discussion see Camp 2006*a*.

this is absurd; and Romeo would strenuously deny this commitment, without feeling any pressure to retract his original claim.[9]

What Romeo does do is *say* that Juliet is the sun; and he cannot deny a report that says he said this. What he can and should respond, if he wants to avoid going on record as committed to Juliet's being a hot gaseous body, is that he didn't *mean* (or intend to be taken to mean) what he said—he meant something else instead (Camp 2006*a*). Similarly, two people who utter "That's a fantastic idea" about the same suggestion, where one is sincere and the other sarcastic, do both *say* that the plan is a fantastic idea, but only one means it. The undeniability of such collective 'say' reports, together with the intuitive falsity of the corollary collective 'claim' reports, reflects a theoretical distinction that we need anyway, whatever we call it: the distinction that Austin (1962) described in terms of locutionary and illocutionary speech acts.[10]

Cappelen and Lepore expressly deny the utility of any distinction in this general vicinity. They consistently treat "saying, claiming, asserting etc." as freely interchangeable; and they say repeatedly that the Gricean distinction between what is said and what is implicated "is superficial and of no significance in trying to find semantic content" (2005: 57). The notion of 'what is said' that I'm elucidating is not quite Grice's. On Grice's view, saying is an illocutionary act that requires sincere commitment; otherwise, one merely "makes as if to say". But my locutionary notion, which is close to that employed by others like Bach and Harnish (1979), does much of the work that Grice wanted. In any case, we need some such distinction to explain how non-literal utterances can commit

[9] It's possible that Cappelen and Lepore would simply insist that speakers who speak metaphorically *do* assert the proposition semantically expressed, but that this assertion is charitably ignored in favor of more salient and communicatively relevant propositions. This line of defense is suggested by their claim that "Nothing even prevents an utterance from asserting (saying, claiming, etc.) propositions incompatible with the proposition semantically expressed by that utterance" (2005: 4). By itself, this position is already far from intuitive, especially when it's applied to sarcasm, where the speaker means something like the opposite of what she says. Further, though, Cappelen and Lepore follow this claim with the following: "From this, it further follows that if you want to exploit intuitions about speech act content to fix semantic content, then you have to be extremely careful in so doing. It can be done, and we'll show you how, but it's a subtle and easily corrupted process" (ibid.). What they "show us how" to do is employ the tests discussed in §2. But those tests suggest precisely that we don't ordinarily take speakers of metaphorical utterances to have claimed or asserted the proposition semantically expressed. Further, although non-literal speech makes the case most clearly, the case against treating semantic content as universally asserted content need not rest exclusively on non-literal speech; see e.g. Soames (forthcoming) for evidence from co-referring expressions.

[10] Some theorists deny that two people who utter "That's a fantastic idea" about the same plan, one sincerely and the other sarcastically, or two people who utter "She is the sun" about the same woman, where one means that she is beautiful and the other that she is dangerous, do *say* the same thing. Such theorists then need a further term for what the two people did in common; as I note below, 'utter' will not do, because the sense in which two people 'did the same thing' goes beyond mere quotation to assign values to context-sensitive expressions. I am focusing here on drawing the distinction between locution and illocution, and prescinding from whether and which forms of 'meaning enrichment' should be included within the locutionary act. As I mention below, locutionary 'what is said' is not itself equivalent to semantic content, and so it would not be surprising if it included some enrichment.

their speakers to determinate contents in the manner of assertion without thereby committing them to the semantic content of the uttered sentence.

Given Cappelen and Lepore's skepticism about "what was strictly speaking said", as discussed in §3, we might predict that they will insist here too that the only use of 'say' that's not equivalent to 'assert' is one on which it functions as a device of direct quotation. This claim might seem more plausible here than it did in the context of §3, given that the locutionary notion of 'saying' is so minimal and seems to be largely driven by fairly recondite theoretical concerns. However, it's important to see that 'say' does have an ordinary use denoting a merely locutionary act. We can observe this in the utter naturalness of reports like the following:

(5.1.1) John: "This man here is your department's new quarterback."

(5.1.2) Alex: "When John said that Bill was our new quarterback, I didn't realize he meant it literally: I thought Bill was our new chairman. But now our whole department is practicing the flea-flicker."

(5.2.1) Bill: "She sounds like a real winner. Let's bring her on board right away."

(5.2.2) Alex: "Bill *said* that we should hire Jane immediately. But I think he was being sarcastic—you can never tell with those deadpan Brits. She doesn't seem like his kind of candidate at all."

Because these reports assign values to contextually sensitive terms and permit some substitution of other expressions, they go significantly beyond reporting just the words uttered; they are definitely not direct quotations. But they also explicitly distinguish what was said from what was meant or asserted. Similarly, two speakers A and B, who say "You are my sunshine" to C and D respectively, have uttered, but not said, the same thing: A said that C is A's sunshine, while B said that D is B's sunshine. Neither has claimed that their addressee actually is sunshine; instead, perhaps, A has claimed that C keeps him honest, while B has claimed that D makes him happy when he is sad. To keep track of what is shared and different across utterances like these, then, we need to appeal to all three levels of content: of what is uttered, what is said, and what is asserted.

Even this restricted notion of 'what is said', identified by what can be reported with utterances like (5.1.2) and (5.2.2), still isn't equivalent to semantic content; for instance, it permits substitution of co-extensive expressions with distinct modal profiles. But locutionary 'what is said' is well-equipped to play the role in communication that Cappelen and Lepore assign to semantic content: thus, it is "content the audience can expect the speaker to grasp (and expect the speaker to expect the audience to grasp, etc.) even if she has mistaken or incomplete information" (2005: 184). In particular, it is content that the hearer has access to even if he is mistaken or ignorant about whether the speaker intended to

assert the semantic content of her uttered sentence, and that he can use (in conjunction with further assumptions) to determine what she could plausibly have meant. Locutionary content is thus a definite, reliable "starting point" (185) for determining illocutionary (and perlocutionary) content. It is also content that is psychologically real, not least in the sense that hearers appeal to it when asked to justify why they think a speaker could plausibly mean *this* by saying *that.*[11]

5. CONCLUSION

Once we are armed with the distinction between saying and asserting, we can fairly easily revise Cappelen and Lepore's tests so that they no longer suggest that metaphor, and non-literal speech more generally, is semantically context-sensitive. (I leave open the question of whether they are otherwise reliable.) We should formulate the tests using the locution "what the speaker said by uttering S", and check that the reported content meets the standards I sketched in §§3 and 4: it should be acceptable to both the original speaker and to a reasonable hearer of the original utterance who was unsure or mistaken about whether the speaker meant what she said. This is a fairly nuanced standard, to be sure, but it is, I have argued, a standard to which our ordinary communicative practices are indeed sensitive.

In the end, I doubt that shifting from 'what is said/asserted/claimed etc.' to 'what is said' construed restrictively will offer much comfort to the contextualist. Nor does it force us into Speech Act Monism (SPAM). But it does require abandoning the most wanton form of SPAP. On a more temperate view, speakers enact a series of nested speech acts by making any utterance: uttering begets saying, which begets asserting/ordering/promising, which begets indicating/pressuring/warning, and so on; each of these speech acts may itself contain multiple contents. I suggest that this moderate pluralism is both compatible with our ordinary communicative practices and nuanced enough to underwrite a reasonably simple and systematic semantic theory.

Cappelen and Lepore are correct that unrestricted maximalism ultimately collapses into radical relativism. They are also correct that indexicals like "I" and "today" form a clear class of particularly well-behaved context-sensitive terms. But if the minimalist wants to ground her overall semantic theory empirically, then she must draw some distinctions within a pool of fairly messy data that fall between these two extremes. Our reporting practices are one source of such data, but they are not the only source; we also need to examine how speakers and hearers actually use and respond to utterances in conversation. Cappelen and Lepore may not like my way of divvying up the data. But without some such distinctions, we end up with a theory of speech acts that makes communication—that is the process of making and exchanging commitments about how the world is—appear utterly baffling.

[11] For further discussion of how to identify the relevant content and why it can play this theoretical role, see Camp 2006*a*. For a review of evidence that literal meaning does play a role in the processing of metaphor, see Camp 2006*b*.

REFERENCES

Austin, J. L. (1962), *How to Do Things with Words.* Cambridge, Mass.: Harvard University Press.

Bach, K., and R. M. Harnish (1979), *Linguistic Communication and Speech Acts.* Cambridge, Mass.: MIT Press.

Bezuidenhout, A. (2001), 'Metaphor and What is Said: A Defense of a Direct Expression View of Metaphor', *Midwest Studies in Philosophy*, 25: 156–86.

Brandom, R. (1983), 'Asserting', *Noûs*, 17/4: 637–50.

Camp, E. (2005), 'Josef Stern, Metaphor in Context', *Noûs*, 39/4: 716–32.

____ (2006*a*), 'Metaphor, Contextualism, and What is Said', *Mind and Language*, 21/3: 280–309.

____ (2006*b*), 'Metaphor in the Mind: The Cognition of Metaphor', *Philosophy Compass*, 1/2: 154–70.

Cappelen, H., and E. Lepore (1997), 'On an Alleged Connection between Indirect Speech and the Theory of Meaning', *Mind and Language*, 12/3–4: 278–96.

____ and ____ (2005), *Insensitive Semantics.* Oxford, Blackwell.

Carston, R. (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication.* Oxford: Blackwell.

Davidson, D. (1978), 'What Metaphors Mean', in S. Sachs (ed.), *On Metaphor*, pp. 41–58. Chicago: University of Chicago Press.

Hills, D. (1997), 'Aptness and Truth in Verbal Metaphor', *Philosophical Topics*, 25/1: 117–53.

Lewis, D. (1979), 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8: 339–59.

Recanati, F. (2001), 'Literal/Nonliteral', *Midwest Studies in Philosophy*, 25: 264–74.

Saul, J. (1999), 'Substitution, Simple Sentences, and Sex Scandals', *Analysis*, 59/2: 106–12.

Soames, S. (forthcoming), 'The Gap between Meaning and Assertion: Why What we Literally Say Often Differs from What our Words Literally Mean', in M. Hackl and R. Thornton (eds.), *Asserting, Meaning, and Implying.*

Stern, J. (2000), *Metaphor in Context.* Cambridge, Mass.: MIT Press.