

8

An Abuse of Context in Semantics

The Case of Incomplete Definition Descriptions

1. Introduction

According to Russell (1905, 1919), any sentence of the form [The F is G] is true just in case *there is exactly one* F and it is G.¹ And so ‘The president of the United States of America in 2000 is happy’ is true just in case there is a unique US president in 2000, and he is happy. So understood, this sentence expresses a perfectly general proposition and not a singular one about Bill Clinton.

Critics and champions alike have fussed and fretted for well over half a century about whether Russell’s treatment is compatible with acceptable uses of incomplete definite descriptions,² where a description [the F] is incomplete just in case more than one object satisfies its nominal F, as in (1).

- (1) The table is covered with books.

If Russell is right, it follows that unless every table but one is destroyed, (1), and so all of its tokens, is false. Yet utterances of (1) are often taken to be communicating something true, even though everyone knows the world is table abundant. Since using incomplete descriptions need not compromise defective conversational exchange, how could Russell be right?

Some authors conclude he cannot (e.g., Strawson, 1950; Donnellan, 1968; Devitt, 1981; Wettstein, 1981), arguing that an utterance of (1) says something true only

¹ This chapter was presented at various universities and conferences and I’d like to thank all these audiences for putting up with me as I tried to make sense of this debate. Special thanks to my colleague Stephen Neale for provoking my interest in this topic, and also to Mark Baker, Emma Borg, Ray Elugardo, Jerry Fodor, Lou Goble, Bob Hale, Kent Johnson, Brian Loar, Murali Ramachandan, Paul Pietroski, Barry Smith, Agustin Rayo, François Recanati, Nathan Salmon, Stephen Schiffer, Zoltan Szabó, Peter van Inwagen, Alan Weir, Howard Wettstein, Tim Williamson, Crispin Wright, Zsafia Zvolenszky, and most especially, Herman Cappelen, Jeff King, and Jason Stanley for extensive and often exhausting exchanges about these and other topics.

² Quine, 1940; Reichenbach, 1947; Strawson, 1950; Sellars, 1955; Vendler, 1967; Kripke, 1977; Hornsby, 1977; Peacocke, 1975; Lewis, 1979; Sainsbury, 1979; Grice, 1981; McCawley, 1981; Davies, 1981; Devitt, 1981; Wettstein, 1981, 1983; Evans, 1982; Salmon, 1982, 1991; Barwise and Perry, 1983; Recanati, 1986; Soames, 1986; Bach, 1987; Blackburn, 1988; Reimer, 1992; Neale, 1990; Schiffer, 1995; Larson and Segal, 1995; to name just a few of the contributors to this debate.

when ‘the table’ is a singular term that refers to (and does not quantify over) a specific table.³ And, more likely than not, were someone to utter (1) directly in front of the sole table in a room, without any other table having already been rendered salient, his auditors would take him quite naturally to be talking about this table. Of course, unlike a proper name, not every token of ‘the table’ picks out the same table, so definite descriptions are at best context-sensitive singular terms.

The presumption is that if incomplete definite descriptions are singular terms, why then not unite them, and treat complete ones as singular terms as well? But incompleteness by itself surely could not establish singularity. For even if this line is right about some cases, it cannot accommodate every unobtrusive use of an incomplete definite description, since some known incomplete descriptions are used without anything being potentially referenced. For example, finding a gruesomely mangled body, a detective exclaims (2).

(2) The murderer is insane.

The description ‘the murderer’ is incomplete, since murder is rampant. Yet our speaker need not have any particular murderer in mind. His attitude might be that, regardless of whoever committed this crime, what was expressed by his utterance of (2) is true, having based his conclusion entirely on the state of a victim’s body, and not on any particulars about the identity of whoever committed this crime. Explaining how (these so-called attributive) uses of incomplete descriptions (Peacocke, 1975, 117; Davies, 1981, 150; Recanati, 1986, 67; Soames, 1986, 278) can be used to say something true is one of the ultimate aims of this chapter.⁴

In short, unexceptional uses of incomplete definite descriptions that are not singular terms are endemic. But the traditional view that all definite descriptions are

³ Many authors recommend that some uses of incomplete descriptions are complex demonstratives: ‘it seems to me to be likely that [“incomplete”] descriptions such as ‘the table’ present difficulties for a Russellian analysis. It is somewhat tempting to assimilate such descriptions to the corresponding demonstratives (for example, “that table”), replacing the “the”, say, in “the table”, with “that” as in “that table” (Kripke, 1977, 22; see also Peacocke 1975, 209; Hornsby, 1977, 33; Brinton, 1977, 400–2; Devitt, 1981, 517; Wettstein, 1983, 52–8; Lycan, 1984, 190; Millican, 1990, 179–80; Schiffer, 1995, 43). *Pace* these authors, in addition to not being semantically innocent, such assimilation does not advance a case for even some uses of incomplete descriptions being (demonstrative) singular terms, since expressions of the form ‘That F’ are themselves not singular terms; but rather (restricted) quantifier expressions. See Lepore and Ludwig, 2000.

⁴ This sort of example is often invoked to separate referential from attributive uses of incomplete definite descriptions. But, unfortunately, it relies essentially on an unexplicated notion of ‘what a speaker has in mind’. However, it is easy enough to devise examples that do not so rely. Suppose someone utters ‘Every US president is married to the woman to his right’, in a context where every living past and current US president is standing directly to the left of his wife. Knowledgeable people will take the speaker to have said something true. But suppose that everyone on stage is lined up so that the president first in line has several women to his right, only one of whom is his wife. This alignment renders ‘the woman to his right’ incomplete. In this case, does it even make sense to posit a singular reference? A speaker might not know how many presidents are assembled, basing his utterance on what he knows about protocol on such occasions. To claim he referred to, say five women with a single utterance leaves reference thoroughly unintuitive and theoretically useless. (Indeed, had he prefaced his utterance with ‘it is possible that’, assigning it wide scope, he would, on this account, have referred to indefinitely many possibilia!)

quantifiers contributing only general uniqueness conditions to propositions expressed by their use is jeopardized by incomplete ones if they on occasion denote. So, what are we to do in the face of these commonplace linguistic facts?

Giving up Russell's achievement has proven enormously difficult. The definite article 'the' behaves grammatically and in at least some cases uncontroversially semantically like standard quantifier expressions, so much so that it is hard to see how a semantic theory of complex noun-phrases could proceed systematically were we to deny definite descriptions quantificational status. In response, two sorts of strategies have emerged in the literature for protecting standard quantificational treatments of definite descriptions from the phenomena of incompleteness. One strategy is to treat definite descriptions as harbouring hidden indexical expressions, so that whatever descriptive meaning alone leaves unfinished its context of use can complete.⁵ The other strategy concedes that incomplete definite descriptions never denote, and so (1) and (2), for example, are not true. But, just the same, our psychologies compensate for the shortcomings of our language and an overabundant world to determine what true proposition a speaker is trying to get at when he uses an incomplete description.

The former strategy is semantic, insofar as it aims to show how seemingly incomplete descriptions can be both quantifiers and denoting, and so utterances of (1) and (2) can be true. (In this regard, this strategy shares a goal with the strategy that tries to treat incomplete definite descriptions as singular terms. They all aim to render utterances of sentences like (1) and (2) true.) The latter strategy, however, is pragmatic, inasmuch as, though it endorses Russell's account, it concedes that nothing is ever denoted (or referenced) by an incomplete description—and so sentences like (1) and (2) are false. Still something true might get conveyed (or 'speaker meant') by a particular use of one of them.

In what follows I will first present schematically various semantic strategies, charging that each suffers from what I shall call an overgeneration problem, namely, each permits uses of sentences with incomplete descriptions to express propositions that they cannot reasonably be held to express. In addition, I will argue that semantic strategies saddle definite descriptions with more ambiguity than they actually exhibit. The moral of these two criticisms is significant: it is to practise temperance with respect to how much work we allot to context in semantics. Unfortunately, this moral is ignored by a great deal of what goes on in contemporary philosophy, with the study of incomplete definite descriptions being but one instance.⁶ And so I proffer this chapter, in this regard, as a caution against positing indexicals in order to resolve philosophical debates.

⁵ 'Hidden' because they are not pronounced, even though they are real components of uttered sentences.

⁶ An appeal to hidden indexicals can be found, e.g., in the literature on the semantics of propositional attitude verbs like believe and know (with hidden indexicals referring to modes of presentation or standards of evaluation: Richard, 1990; Cohen, 1999; DeRose, 1992, 1995; Lewis, 1996), comparative adjectives (with hidden indexicals referring to comparison classes: Partee, 1973), vague predicates (with hidden indexicals referring to precise properties or classes: Lewis, 1983b; Soames, 1986), in moral theory (with hidden indexicals referring to moral standards or set of preferences: Dreier, 1990; Davies and Humberstone, 1980).

Next I will turn to so-called pragmatic approaches. Pragmatic strategies agree with semantic ones in two chief respects: they all hold that when someone speaks, something unique is (literally) said and further that whatever is said (and even what is conveyed) by an utterance of a sentence is determined exclusively by the context of utterance. Both claims, I will argue, are mistakes and exposing them will render transparent that unobtrusive uses of incomplete definite descriptions pose no threat whatsoever to the traditional quantificational treatment of definite descriptions. So, in effect, my ‘solution’ to the phenomenon of incompleteness is radical. I agree with the pragmatist that no incomplete definite description is denoting, but neither pragmatics nor semantics offers a workable solution to the problem of incompleteness, and practitioners of both share a mistaken picture about what is said. Eliminating that picture, I will argue, also eliminates the alleged problem of incomplete definite descriptions.

2. Semantic Proposals

We begin with several semantic proposals that try to explain how someone *can* denote *à la* Russell with an incomplete definite description. The first group tries to render incomplete definite descriptions denoting by showing how the nominal F in [The F] can be interpreted relationally, in context, as elliptical for a richer (complete) description. The other semantic proposal we will consider subsumes incompleteness under a more general problem of domain selection—e.g., how can an utterance of ‘Everyone ate potatoes yesterday’ be true, since at least one person did not eat potatoes yesterday?—and then tries to show how in context the domain of an apparently underspecified quantifier can be restricted with ‘the’ being a limiting case.⁷ I don’t pretend that my survey of semantic proposals in response to incompleteness is exhaustive. That would be to a priori limit philosophical imagination.⁸ It is, however, intended to be representative of what’s currently on offer in semantics (and pragmatics), and if I’m right, they all suffer from excessively high expectations of both semantics and pragmatics.

2.1 *Completion by appeal to an implicit descriptive qualification*

A natural and old reaction to the phenomenon of incomplete definite descriptions is that a speaker who knowingly uses one is speaking elliptically. She is denoting what she would denote were she to have uttered a different description, whose linguistic meaning together with the way the world is uniquely denotes (see, e.g., Quine, 1940, 146; Vendler, 1967, 45–6, 52, 55ff). So, e.g., it is natural to understand someone who tokens (3) and (4) in sequence,

⁷ Absorbing the problem of incomplete descriptions into the more general problem of underspecified quantification, if legitimate, provides another reason for resisting any inference from incompleteness to singular reference. See Sainsbury, 1979; Davies, 1981; and Westerstahl, 1985; and Neale, 1990.

⁸ However, the only other semantic proposal for accommodating the problem of incompleteness I’m aware of is discussed and criticized in note 23. So, my discussion is as far as I know *de facto* complete.

- (3) John saw exactly one man in Mary's kitchen last week.
 (4) The man wore a hat.

as denoting with 'the man' in (4) what 'the man whom John saw in Mary's kitchen last week' does, a description with a recovered 'restricted adjunct based upon a previous occurrence of the same noun in an identifying context' (Vendler, 1967, 45–6; also, 52). However, since (1) and (2) can initiate a conversation, this *linguistic* rule cannot govern every potential completion.

Alternatively, when linguistic context alone fails to determine unarticulated completing descriptive materials, it's natural to think that non-linguistic contextual cues do (e.g., Quine, 1940, 146; Vendler, 1967, 55ff). If a speaker tokens (2) in a discussion of Nicole Simpson's murder, she might be saying what she would have said had she instead uttered (5).

- (5) The murderer of Nicole Simpson is insane.

Or, suppose a bloody knife monogrammed with the letters 'OJ' was found next to the victim in question. We might then take the speaker to have said what she would have had she instead uttered (6).

- (6) The murderer who used the bloody knife monogrammed with 'OJ' is insane.

If this strategy is intended as a semantic solution to the problem of incomplete definite descriptions, we need to be told how the semantics (or syntax?) can introduce alleged elliptic materials. One suggestion is that we treat the interpretive truth conditions for a sentence like (2) along the lines of (2'),

- (2') [The x: murderer (x) & R(x)](x is insane)

where 'R' is a placeholder for whichever completing material gets contextually invoked by an utterance of (2) to conjoin with the meaning of its nominal to render its incomplete definite description denoting. On this view, then, what a speaker expresses with an utterance of (2) is what she would express had she instead used a sentence in which the relevant unarticulated completing descriptive content is explicit.

The view that some uses of incomplete definite descriptions still denote because these uses have unarticulated completing descriptive content assigned to them in context has met with a torrent of criticism over the years (e.g., Donnellan, 1968, 204; Wettstein, 1981; Larson and Segal, 1995, 331–2). A sampling of these criticisms will help motivate proposals that exploit indexicality in trying to solve the problems created by incomplete definite descriptions.

2.2 *Problems with positing unarticulated descriptive completions: overdetermined, false, or impoverished completions*

One chief worry with this proposal concerns how the completing property is to be determined. For example, since every (attributively) used incomplete description has

indefinitely many non-synonymous (even co-) denoting completions, shall we conclude that a speaker fails to say any truth in particular? Or, should we conclude that she expresses each distinct completion with her single use of an incomplete definite description? Surely, there are limits to what a speaker can express with an utterance of (1) or (2).

Suppose we limit acceptable descriptive completions to those properties determined by what a speaker *has in mind* when she speaks. That will eliminate some, but if her mind is chock full of descriptive completions,⁹ which one (or ones) determines what she has said? Can it be a combination thereof? Does a speaker who utters (2) with both of the completions ‘the murderer of Nicole Simpson’ and ‘the murderer who used the bloody knife monogrammed with “OJ”’ in mind express what she would have had she instead uttered not (5) or (6) but instead (7) in that same context?

- (7) The murderer of Nicole Simpson who used the bloody knife monogrammed with ‘OJ’ is insane.

Suppose there are more than two. Do they all somehow compose into a single determining property—the murderer who is F and G and H and I ...?

Ignoring worries about how psychologically feasible this proposal is the more overdetermined a denotation is by what’s in a speaker’s mind, the suggestion is foiled by false or impoverished beliefs. Suppose what a speaker has in mind when she tokens an incomplete description determines a failed completion (i.e., whatever is determined by what she has in mind applies to nothing). A speaker might err about a victim’s identity or about which weapon was used, mistaking one victim or weapon for another. Do her mistaken beliefs render her utterance of (2) false? Note that ordinarily misidentification does not suffice to render false your utterance of a sentence with a *complete* definite description. An utterance of ‘The positive square root of four is even’ is true, regardless of whether the speaker falsely believes the number in question is four. Why should it be any different for utterances of (1) or (2)?

Even if a speaker were infallible, what he has in mind might still fail to determine a completing property (Donnellan, 1968, 204); for example, your local environment might be more populated with qualitatively indistinguishable tables or copycat killers than either you or your auditors suspect. In such cases, are your utterances of (1) and (2) rendered false?¹⁰

Confronted with worries about overdetermination, underdetermination, false, and conflicting beliefs, some authors have forsaken the strategy of seeking a *descriptive* completion contextually determined by what’s in a speaker’s mind and instead have looked to implicit indexicality to convert an incomplete definite description into a denoting one.

⁹ Schiffer (1995, 376), among others, states that relevant completions must be restricted not just to what a speaker has in mind, but also what *auditors* can retrieve from context. Schiffer is not conflating epistemology with semantics, but he is, at least for his discussion, prepared to swallow this much intention-based semantics. I’m not, but nothing I argue hangs on my refusal.

¹⁰ For further criticisms of descriptive completions, see Larson and Segal, 1995, 331–2.

2.3 Indexicality and completion

That context can contribute to what is expressed by collaborating with and yet *not adding* to descriptive content is old hat. When you and I each utter ‘I am happy’, what we express differs, though the linguistic meaning of our words is identical. What you express may be true, while what I express may be false, because in our respective contexts of use, the indexical expression ‘I’ refers to you and to me respectively. Divergence is achieved precisely because of what ‘I’ means, namely, that every token of it refers to whoever uses it, and this obtains regardless of how confused or deluded either of us is. In effect, indexical expressions like ‘I’ and ‘now’ are over and underdetermination, and false and inconsistent belief resistant. What has this got to do with incomplete definite descriptions?

Russell was aware of the role indexicality can play in determining a denotation for a definite description, as with overtly indexical descriptions like ‘the *present* king of France’, ‘the table *here*’, and ‘*my* father’ (Russell, 1959, 239; see, also, Reichenbach, 1947, 258).¹¹ Tokens of these descriptions can denote regardless of how confused their users are about when, where, or who they are.

The next two semantic proposals we will consider contend that indexicality can render denoting some uses of an incomplete definite description in a way that sidesteps issues about overdetermination, underdetermination, false, and inconsistent belief. On one proposal, a contextually salient item gets indexed; on another, a contextually salient restriction on a domain of quantification gets indexed.¹² Unlike the first proposal, neither of these requires that the contextually determined items or restrictions invoke descriptive (or qualitative) identifying properties. We will examine each proposal in turn.

2.4 An implicit argument place

Suppose, e.g., someone who utters (2) is saying what he would have said had he instead uttered (2’), where the speaker’s token of ‘her’ picks out the victim.

(2’’) Her murderer is insane.

Likewise, in uttering (1), a speaker might have said what he would have said had he instead uttered (1’) or (1’’).

(1’) The table there is covered with books.

(1’’) Her table is covered with books.

The general story from which these instances are drawn is that the nominal in each description can latch onto the most contextually salient object that completes the

¹¹ By an ‘indexical’ here and throughout is meant any expression-type whose semantic value can vary across use, i.e., can vary from context to context.

¹² I do not here want to limit how these completions get developed. Some proponents have developed hypotheses about the relationship between logical form and indexicality; others have not. My criticisms are intended to stick however these details get worked out. However, elsewhere I’ve developed criticism of so-called ‘unarticulated constituent’ approaches, where the semantics is run not on sentences but on utterances (cf., Cappelen and Lepore, 2000).

description (Evans, 1982, 324; Soames, 1986, 279; Neale, 1990, 100–1). In the present cases, it might be a location (as in (1')), or a person (as in (1'') or (2'')). If no such item is available, the utterance is without a truth-value—in much the same way that (1') and (2'') are without truth-value in a context where no female is demonstrated.

The advantage of this suggestion is supposed to be that context sensitivity does its semantic work regardless of how confused, misinformed, over or underinformed a user is. A speaker succeeds in picking out himself when he utters 'I'm happy', even if he may believe himself to be Napoleon (Kaplan, 1989a). So too a speaker can succeed in denoting a table when he says (1) even if he's confused about its exact position or possessor, as long as some contextually salient item anchors its nominal and renders it uniquely applicable.

We cannot evaluate this proposal without being told more about how indexing can be achieved with uses of apparently non-indexical incomplete definite descriptions. In the relatively sizeable literature where this sort of suggestion is recommended, or at least hinted at as an option, no proponent (or critic) develops a procedure for determining how an item in a context gets indexed when someone uses an incomplete definite description. Rather, recommendations are alleged in specific cases. In some writings tokens of incomplete definite descriptions are treated as indexing locations, as with some uses of 'the table' mapping onto the table *there* (Sellars, 1955, 200–1; cf., also, Grice, 1969, 142; Vendler, 1967, 55; Sainsbury, 1979, 122; Evans, 1982, 324). In others, tokens are treated as indexing times, as with some uses of 'the president' mapping onto the *present* president (Recanati, 1986, 61); some as indexing people, as with some uses of 'the car' onto *our* car (Soames, 1986, 301), 'the murderer' onto *his* murderer (Wettstein, 1981, 250–2; Salmon, 1981, 44).

None of these recommendations is contextually invariant. It's easy to imagine circumstances where any incomplete description can be unobtrusively used but its recommended contextualization fails. For example, when a parking attendant tells a customer, 'The car is out front', the latter would not take him to be denoting their jointly owned car. Or, when a detective wonders out loud, 'Did the murderer also murder *him*', no listener would take him to be wondering whether *his* murderer murdered *him*.

So by what means do speakers cancel potential completions and favour others? All of the familiar examples suggest that their proponents may have in mind a general proposal according to which sentence-types (1) or (2) share truth conditions with (1*) or (2*),

- (1*) The table (i) is covered with books.
- (2*) The murderer (i) is insane.

where 'i' is an indexical, the semantic value of which, in a context of use, is whatever most contextually salient item combines with the linguistic meaning of an incomplete description to render it denoting. If no such item exists, that use of the incomplete definite description is non-denoting. In short, what actually gets indexed can vary from context to context—sometimes a speaker, sometimes a time, an addressee, a location, or a perceptually or topically salient object, whether a victim, a town, a country, a

lap, or whatever—contingent on what happens to be relevantly most contextually salient is in a particular context of use.

Tantalizing metaphysical, conceptual, and psychological questions remain about what determines which item in any context is most contextually relevant for rendering a token of an incomplete definite description denoting, but this proposal has kept up its end of the semantic bargain by sketching how an incomplete description might denote on an occasion of use.

Before evaluating this proposal, I want look at the other route some authors have taken in order to exploit context sensitivity, trying to cast the phenomena of incomplete descriptions as subsumable under the more general problem of underspecified quantification.

2.5 Contextually selected domains

Teachers who assert (8) are rarely challenged, and are almost always taken to have said something true.

(8) Every student must hand in homework.

Yet it's unlikely that anyone who utters (8) expects every student in every course in every school in each school district in every country to hand in homework. How can this intuition be squared with the ordinary semantics for quantifiers? On an occasion when a speaker utters (8) can she succeed in asserting (9)?

(9) Every student here must hand in homework.

But how is a restriction to every student *here* to be effected? Standard semantic accounts of quantification assume that a domain for a quantifier expression is determined by its nominal alone, say, students for the noun phrase 'every student'. Relevant uses of (8) are supposed to show that a domain of quantification can be determined by context as well. With an utterance of (8), the quantifier 'every' doesn't quantify over all students, only discourse relevant ones. Of course, a restriction on 'every' needn't be as narrow as in (9). Context is supposed to determine just how broadly or narrowly a domain of quantification is to be restricted. Accordingly, paralleling the last suggestion, a narrower quantifier domain might get selected, say, a contextually salient set that intersects with the extension of 'student' to provide the domain of discourse for 'every'. It might be every student present, or every student enrolled in the speaker's class. Truth conditions for (8) might be explicitly represented by (8*),

(8*) Every student (i) must hand in homework.

where 'i' indexes not a single item (as in the last sort of contextual proposal) but rather a set¹³ that restricts the domain of 'every student'. Truth conditions for (8) are represented only slightly more rigorously by (8**).

¹³ Consider 'If there were more than 30 bottles, John would be happy' (Stanley and Szabo, 2000). This example shows we might need to invoke possible sets as well as actually ones, i.e., properties.

(8**) [Every x : x is a student & x is in i](x must hand in homework)

In short, in some contexts, a contextually determined domain of quantification might be the set of students enrolled in that class; in another it might include auditors. In uttering (8), a speaker succeeds in expressing something true, because her utterance succeeds in indexing a set of objects that serves to restrict the quantifier expression ‘every student’ to a narrower contextually defined domain. How a speaker conceptualizes this domain may or may not be relevant in determining which set gets selected—much the same way as on standard semantic accounts for indexicals it is irrelevant how a speaker conceptualizes himself with a use of ‘I’.

A bonus of appealing to context sensitivity in accounting for domain selection is its explanation for how distinct tokens of the same quantifier expression-type in a single sentence might range over distinct domains. Suppose there are two groups of sailors, one on deck and one on shore, and all the sailors on deck waved to all the sailors on shore.¹⁴ In such circumstances, one might conclude that (10) can be used to express something true.

(10) Every sailor waved to every sailor.

But relativizing quantification over a fixed set of sailors, an utterance of (10) asserts that every sailor in this set waved to every sailor in this same set, including each to him or herself. Instead, an utterance of the two quantifiers in (10) might reasonably be taken as restricted to distinct domains, expressing in its context of utterance what (11) would in that same context of utterance.

(11) Every sailor (here) waved to every sailor (there).

Accordingly, truth conditions for (10) might be better represented as (12),

(12) Every sailor (i) waved to every sailor (k),

where ‘ i ’ and ‘ k ’ can index disjoint sets, so that ‘ i ’ together with the first token of ‘sailor’ can pick out the group of sailors on deck, and ‘ k ’ together with the second token of ‘sailor’ can pick out the group on shore, thus rendering (10) to express a truth.

Assuming this second contextual proposal explains how a truth can be expressed with underspecified quantifiers, can it be extended to the article ‘the’ to explain the phenomenon of incomplete descriptions? A number of authors have claimed it can (Sainsbury, 1979, 114; Davies, 1981, 160; Lycan, 1984, 190; Recanati, 1986, 68; Higginbotham, 1988, 39; Neale, 1990, 94–5; Salmon, 1991, 89f, 95; Stanley and Williamson, 1995; Stanley and Szabo, 2000).¹⁵ Extended to the article ‘the’, the truth conditions for (1) or (2) might be more explicitly represented as (1+) and (2+):

(1+) [The x : table (x) & i (x)](x is covered with books)

(2+) [The x : murderer (x) & i (x)](x is insane)

¹⁴ Adapted from Stanley and Williamson, 1995.

¹⁵ See Larson and Segal, 1995, 333–4 for reasons not to assimilate incompleteness to quantifier underspecification.

where ‘*i*’ attaches to the nominal expressions ‘table’ and ‘murderer’; and in a context *C*, whatever set (property) is picked out by this indexical can restrict a domain of quantification of ‘the’. If the intersection of the extension of ‘table’ with whatever set is indexed has one table in it and it is covered with books, then (1) is true in that context; if it is not covered with books or if there is no unique table in that set, then (1) is false in that context.

2.6 Problems for context-determining domain selection

Does this proposal for domain specification rule out someone who utters (1) *in a home* with three tables from expressing a truth? Or, suppose he utters it *in a single room* in which three tables all lined up in a row, can he still express something true? Any notion of context that excludes two of these tables demands clarification (Reimer, 1992, 355–7; 1998a, 99–100; Larson and Segal, 1995, 332). We want to know by which means context can be so fine-grained. But this is less a problem about meaning, and more a metaphysical problem about how to individuate contexts, and so, I will defer.

There are also general worries about contextual treatments of domain selection. For example, requiring nominals attached to quantifiers to be contextually sensitive induces odd logical consequences, or at least odd to my ears (see, however, Recanati, 1986, 72). (13) looks logically false, and (14) logically true, but neither is on the current strategy.

- (13) Some man died and it is not the case that some man died.
 (14) If a man arrives, then a man arrives.

(13) and (14) are true in a context *C* on this account just in case (13*) and (14*) are true in *C* as well.

- (13*) Some man (*i*) died and it is false that some man (*k*) died.
 (14*) If a man (*i*) arrives, then a man (*k*) arrives.

Since nothing in the semantics guarantees distinct indexes are co-referential in every context, (13) turns out *not* to be contradictory and (14) *not* logically true. It’s not clear one can devise logical truths or falsehoods in sentences with iterated quantifiers expressions on this account.

This account of domain specification also spurns certain standard equivalences, e.g., definite descriptions as equivalent to an existential and universal. On a standard Russellian account, (15) is supposed to paraphrase (16).

- (15) [The *x*: *x* is a man](*x* is tall)
 (16) $(\exists x)(x \text{ is a man} \ \& \ (\forall y)(y \text{ is a man} \supset x = y) \ \& \ x \text{ is tall})$

But since distinct indexes attach to each quantifier in its expansion, logical equivalence is not guaranteed.¹⁶

¹⁶ Stanley and Williamson (1995, 294) embrace this conclusion that a Russellian expansion is not equivalent to its definite description counterpart, because the former is more context-dependent, as it allows more slots for contextual effects than are intuitively possible.

Finally, evaluating inferences can get hairy. Consider an inference from (12) and (17) to (18).

- (12) Every sailor waved to every sailor.
- (17) Popeye and Bluto are sailors.
- (18) So, Popeye waved to Bluto.

According to the current proposal, this inference is invalid, since nothing in the argument guarantees *which* restricted domain Popeye and Bluto are in.

There is a standard worry about evaluating arguments with context-sensitive elements. So, consider the argument (19)–(20).

- (19) That horse is brown and heavy.
- (20) Therefore, that horse is brown.

In order to show that this argument is valid requires fixing co-demonstrations across all contexts in every model. But even after analogous fixing for (12), thereby determining restricted domains for both quantifier tokens the verdict about the inference from (12) and (17) to (18) is still obscured. Needless to say, this will make for messy quantificational rules.¹⁷

These complaints against the current account of domain specification are pleas for elaboration. Instead of pursuing them, I want to end my discussion of contextual supplementation strategies by arguing that they all impose interpretations on incomplete definite descriptions (and underspecified quantifiers) that they do not have.

3. Pseudo-Ambiguity

The contributors do not all agree about how to effect context sensitivity. Some authors claim that in underlying linguistic representation, there must be an indexical. For them then everywhere I suggested that definite descriptions have an indexical in their truth conditions read that as: must have an indexical in their underlying syntactic form at that level of analysis which feeds into the interpretive component of the grammar. If this were correct, then the best syntactic theory would have to justify positing these indexical elements. Others believe context sensitivity requires no explicit representation of indexicality at any level of linguistic analysis. I want to sidestep these heady issues and so every criticism I will now raise against various semantic supplementation strategies applies to any account that aims to explain away the problem of unobtrusive uses of incomplete definite descriptions by invoking context sensitivity—regardless of how that sensitivity gets fleshed out.

My first of two chief complaints is that these strategies posit more ambiguity than is tolerable. I will begin with the strategy that says that ‘the table’ can denote because in a

¹⁷ I have no argument that these problems cannot be remedied, or that there might not be grounds for sacrificing relevant counter-intuitions, but a request for amplification is reasonable.

context of use an item gets singled out that anchors its descriptive content so as to render the description denoting. In different contexts a different item can get singled out, contingent on which item is more contextually salient vis-à-vis the use of that description. In uttering (1), either (1'), (1''), or something else gets expressed (see, however, Reimer, 1992, 353–4).¹⁸

- (1) The table is covered with books.
- (1') The table there is covered with books.
- (1'') Your table is covered with books.

The truth conditions for sentences (1') and (1'') are respectively, something like (21) and (22).

- (21) The table over there is covered with books.
- (22) The table of him is covered with books.

The expressions I want to highlight are the prepositions 'over' and 'of'. These particular choices needn't be correct or unique, but some such expression is needed to explain how the semantic value of 'table' combines with the contextually determined semantic values of 'him' or 'there' to determine a unique denotation for 'the table'; otherwise, we have a set of tables and place or a person without any relationship for connecting them. What's unclear is whether it's best to describe the needed relationship as composing the meaning of 'table' with an *adjunct* like 'over there' or 'of him' or whether it might not be more apt to describe it as positing two distinct relational nouns 'table of' and 'table over' with two locations—him and there (Neale, 1990, 101). It matters little how it's characterized to my main critical point.

How a contextually supplied item contributes to fixing a unique denotation can vary—sometimes the indexical locates a place (in *i*), sometimes a possessor (of *i*). The relation of the indexical to the wanted object may vary as distinct relational nouns or adjuncts are summoned, contingent on context. But this means that the account treats the noun phrase 'the table' or the simple noun 'table' as many ways ambiguous as there are distinct ways of combining an item with the meaning to 'table'. Suppose 'his table' is correctly interpreted as denoting *the table of him*; 'the table there' as *the table over there*; 'that table' as *the table identical to that* (see Lepore and Ludwig, 2000), and 'the topical table' as *the table under discussion*.¹⁹ Since from context to context any of these (or untold other) sorts of completions can be in play, contingent partly on which item is contextually salient, it follows that 'the table' or 'table' is multiply ambiguous. That's a

¹⁸ Underdetermination about which object is most contextually salient may result from normal epistemological underdetermination.

¹⁹ Recanati worries that even if an object implicitly indexed adds no descriptive content to the proposition expressed by an utterance, 'it's still totally indeterminate *which particular sentence expressing that content* the uttered sentence is elliptical for—it's indeterminate which particular 'relationally descriptive term' the incomplete description is elliptical for' (1996, 449–50). For polemical purposes, I'll assume it's determinate which mode of composition is invoked. I care little if it's elliptic, and assume only that it is determinate how an indexed object composes with the meaning of the nominal the indexical attaches to.

lot of ambiguity to have to swallow to protect ourselves from incompleteness.²⁰ Positing this much ambiguity is not only uneconomical, it also gets the semantics backwards. For it tells us we do not understand a use of a definite description until we figure out which item is being indexed.²¹ These circumstances are different from normal instances of ambiguity.

Ordinarily when we hear an ambiguous expression we appeal to various rules of thumb to disambiguate. The complex noun phrase ‘flying planes’ is ambiguous taken on one meaning, if ‘flying’ is a gerund and another if an adjective. Or because the simple noun ‘bank’ is ambiguous, in order to understand its tokens we look to context to help us determine which of its meanings is in play. But the way we look to context with ordinary ambiguous expressions is not by seeing which objects are semantically referenced and then returning to interpreting the sentence; rather, we have views about what sorts of things happen on the sides of rivers vs. what sorts of things happen in financial institutions or how someone’s flying a plane vs. how a plane’s flying overhead can be dangerous, and then we make our choice based on this background information. With incomplete descriptions, we are being asked to disambiguate (an unobvious ambiguous expression) by first seeing which sort of item is being indexed and then asking which sort of meaning can relate this sort of item to the extension of the nominal. That’s awfully bad news for supplementation strategies.

In order to accommodate within the semantics every alleged case of completion for incomplete definite descriptions, these strategies must be sufficiently liberal in what they say about what can be contextually relevant. But the more liberal they are the more implausible they become, until they wind up treating relatively simple expressions as potentially vastly ambiguous.²²

3.1 *Ambiguity criticism extended to domain selection strategy*

One’s first inclination might be to conclude that, since the interpretation rule in the domain specification strategy is uniformly set intersection, this last charge about manufactured ambiguities cannot be extended to it. However, because appeal to context in

²⁰ Notice too that this supplementation strategy in effect requires a partly indexical partly descriptive proposal, as in ‘the murderer of him’, where ‘of’ would be descriptive, ‘him’ indexical, and so this account inherits whatever problems the original descriptive accounts incur.

²¹ Thanks to Jason Stanley for helping me to see how variation in objects indexed together with variation in modes of composition, contingent on which object gets indexed, combine to generate rampant ambiguity. I’m not, however, committing him to endorsing this objection.

Indeed, Stanley believes this objection can be avoided by positing two indices in logical form, so that the logical form of (1) is ‘The table (R) (i) is covered with books’, where ‘R’ indexes a relation and ‘i’ an object. His suggestion fails. First of all, it permits indexing a genitive relation, as in ‘His table’, with ‘R’ while indexing a location, with ‘i’, so that, with an utterance of (1), a speaker could assert whatever ‘There table is covered with books’ means—something one cannot do. Also, it’s hard to see how the account doesn’t saddle itself with a sort of third-man argument, with both a salient relation and a salient item indexed but no statement of how they relate to each other or to the semantic value of the nominal.

²² One way to avoid ambiguity is to sacrifice compositionality: so meaning alone, even appropriately linguistically contextualized, fails to determine the proposition expressed. This route, I’ve argued elsewhere, is too expensive, and indeed, incoherent. Cf., Cappelen and Lepore (2000).

domain selection cannot succeed in discriminating, so I will argue, between distinct scoped readings of sentences with incomplete descriptions, a version of the same criticism extends to the domain specification strategy.

Sentences with more than one quantifier often are ambiguous, contingent on which quantifier has wider scope over the other. The sentence ‘Every man loves some woman’ might mean that there is some woman who is such that she loves every man. On this reading ‘some woman’ is assigned wider scope than ‘Every man’. Or the sentence could mean that every man is such that he loves some woman or other, perhaps a different woman for each man; in which case ‘every man’ is assigned wider scope over ‘some woman’. Likewise consider sentence (23).

(23) In every photo, the red sign stands out.

On the domain specification strategy presumably in uttering (23), a speaker might express that in every (contextually salient) photo the same (contextually salient) red sign stands out. That red sign is so brilliant it always stands out! Or it might mean that in every (contextually salient) photo, there is no more than one red sign, perhaps different in each photo, but whichever red sign that is, it stands out. On the first reading ‘the red sign’ is assigned wide scope over ‘every photo’; and on the second reading it is assigned small scope.²³

The article ‘the’ quantifies over a domain determined by the linguistic meaning of ‘red sign’ together with a contextually determined set; and ‘every’ quantifies over a domain determined by the linguistic meaning of ‘photo’ together with another contextually determined set. So on the wide scope reading of ‘the red sign’ an utterance of (23) is true just in case the unique red sign in the first domain stands out in every member of the second domain.

Actually, interpreting (23) is trickier than so far represented, since on the small scope reading of ‘the red sign’, there is what we might refer to as a *quantifier dependency* relationship between its two quantifiers in this sense: the range of the smaller scope quantifier ‘the red sign’ is partially determined by the range of the one with wider scope, i.e., ‘every photo’. It’s part of the meaning of (23), with this assignment of scopes, that we are only talking about red signs *in photos*. Therefore, it won’t do to interpret (23) along the lines of (24) (with ‘k’ and ‘i’ being distinct indices for each quantifier).

(24) [Every x: photo (k) x] [the y: red sign (i) y] (y stands out in x)

One problem with (24), *inter alia*, is that it fails to capture this small scope reading, because it fails to guarantee that the range of the interior quantifier depends on the range of the exterior one, and so, as is easy to prove, (24) is logically equivalent to the

²³ David Lewis suggests that the best treatment for a (incomplete) definite description ‘The F’ is always to assume they denote the most contextually salient F (Lewis, 1979). Sentences like (23) reveal a serious defect with this suggestion. In a context in which someone utters (23), there need not be any contextually salient F, and even if there were it couldn’t help to secure correct truth conditions for (23), where ‘the red sign’ takes wide scope.

wide scope reading of ‘the red sign’, i.e., it’s committed to there being just one red sign (at least just one salient one).

One clever suggestion²⁴ for fostering the dependency between the two quantifiers in (23) is provided by (25).

(25) [Every i: photo (k) i] [the y: red sign in (i) y] (y stands out in i)

Unlike (24), the definite description ‘[the y: red sign in (i) y]’ in (25) provably cannot be assigned wide scope, because of the restriction created by the first token of ‘in (i)’. So, (25), unlike (24), is distinct in meaning from the wide scope reading of ‘the red sign’ in (23).

According to (25), what is normally treated as a context-sensitive singular term, namely, the ‘i’ in ‘the red sign (i)’, ‘becomes’ in (25) a variable bound by its initial quantifier ‘every photo (k)’. Doing so forces the range of ‘the red sign (i)’ to be constrained by the wide scope quantifier ‘every photo (k)’. That we can effect a small scope reading of ‘the red sign (i)’ by positing a variable that restricts its range looks to provide independent support for positing indexed (quantifier) nominals in unembedded cases like (1) and (2).

In summary, where no higher quantifier binds an (implicit) indexical on a nominal, it behaves semantically like a context-sensitive singular term, as apparently with both ‘k’, in ‘every photo (k)’, and ‘i’, in ‘The red sign (i)’, on its wide scope reading in (23). But when there is quantifier dependency, as in the small scope reading of ‘the red sign’ in (23), the implicit ‘i’ is semantically interpreted as a bound variable, bound by a higher quantifier.²⁵

Obviously in order to have any real confidence in this account, a lot more detail is required. But even with only a sketch, we can show that it fails to secure different scope readings for (23), regardless of how the account winds up getting spelled out.

Suppose (25) accurately represents the truth conditions of (23), when its quantifier ‘the red sign’ takes small scope. Can we, then, represent its wide scope reading along the lines of (26)?

(26) [The y: red sign (i) y][every i: photo (k) i](y stands out in i)

(26) seems to get things right, but it’s not exactly parallel to its other scope read (25), since in swapping quantifiers in (25) to change relative scope, the relevant domain restriction, provided by ‘in’ in (25), has been conveniently dropped in (26). If we replenish this restriction, (26) becomes (27).

(27) *[[The y: red sign in (i) y][every i: photo (k) i](y stands out in i)

²⁴ Recommended by Stanley and Szabo (2000).

²⁵ The criticisms that follow I intend to apply to *any* account that tries to explain how the range of one quantifier can be (partly) determined by the range of another by exploiting the alleged context sensitivity of the dependent quantifier.

There are two ways of reading (27). Either its first token of ‘i’ is a free variable or an indexical singular term. If it’s the former, then (27) is ungrammatical. Every variable must be bound in a grammatical sentence. This token of ‘i’ cannot be bound by its lower quantifier (nor would we want it to be), since that would merely resuscitate the small scope reading of ‘the red sign’, and we are trying to fix its wide scope reading. Suppose, then, that this token of ‘i’ in (27) is an indexical singular term, whose referent is determined in a context of use. Isn’t it perverse to presume that the expression ‘red sign’ takes an indexical singular term whose referent is restricted to red signs that are *in photos*, or at least that are *in something*, in every context in which the nominal ‘red sign’ occurs? Granted, every red sign is in something or other, but surely not as a matter of meaning. Moreover, this sort of manoeuvre would require us to say that the expression ‘red sign’ contributes an entirely different sort of meaning in a sentence like ‘On every corner, the red sign stands out’. This sentence also has two relative quantifier scoped readings. Would it follow, then, that denotation of ‘the red sign’, on its wide scope reading, must be restricted to red signs *on corners* or red signs *on something*? And, then, there is the sentence, ‘Above every entrance, the red sign sits’. Must its wide scope reading for ‘the red sign’ denote, as a matter of meaning alone, red signs *above entrances* or red signs *above something*? And so on and so on. If we are talking about a unique red sign in each of these sentences, as we are alleged to be on a wide scope reading, we do *not* want its meaning so restricted.

The problem is this: once we settle on (25) as specifying the meaning for the small scope reading of ‘the red sign’ for (23), then dropping ‘in’, or ignoring its functional character, in interpreting (23), when ‘the red sign’ takes wide scope, is ad hoc.²⁶ What about sentences in which ‘the red sign’ is the only quantifier tokened? Does ‘the red sign’ contribute a different meaning, contingent on whether it is embedded and also contingent on how it is embedded? That’s a peculiar way to go about assigning meaning to expressions.

In short, on both supplementation strategies simple structures like ‘the red sign’ turn out to be multiply ambiguous, with their meaning (extension) in a linguistic context depending upon the meaning of expressions seemingly independent of it (syntactically or semantically). This meaning dependency is just the sort of semantic non-innocence, semantic theories must reject. The expression ‘red sign’ means *red sign* regardless of the context in which occurs; it doesn’t also mean *red sign on* or *red sign above* or *red sign in*.

This first serious worry I’ve presented for supplementation strategies has to do with their commitment to an open-endedness in determining which relational noun is in play in interpreting definite descriptions. I want now to argue that supplementation

²⁶ If my concern is legitimate, then (some) indexicalists about domain selection, *viz.*, Stanley and Szabo, 2000, who urge treating the implicit indexical either as a variable bound by a higher quantifier or, when free, as a singular term picking out a contextually salient set, would seem to be hoisted with their own petard. See, in particular, note 33 below.

strategies allow us to express too much, or much more than we are capable of expressing, with incomplete definite descriptions.²⁷

3.2 *Non-existent interpretations?*

On each proposal thus far considered, an utterance of (1) can be used to express a truth, because its context of utterance can effect a completion of its incomplete description:²⁸ either a use of (1) fastens on to a contextually salient item, or context restricts the range of its quantifier ‘the’.²⁹ Suppose that in using (1) a speaker indexes a contextually salient location, and thereby, expresses what she would have had she instead uttered (1’).

(1’) The table there is covered with books.

What, then, shall we conclude about (28)?

(28) The table is there.

(28), with its incomplete definite description, is as likely to be used to express a truth as (1). Since (1) can express what (1’) *can* in the same context, why can’t we infer that (28) can express in a context what (28’) would?

²⁷ Another strategy we will not discuss invokes a notion of an indeterminate statement. In addition to the intuition that incomplete definite descriptions can be used to denote, another widely shared intuition is that they can be used to make ‘a determinate assertion’ (Salmon, 1981, p.39; see, also, Recanati, 1986, 67; Neale, 1990, 94). Suppose we reject this latter intuition, and assume that in many cases no unique completion is achievable. A speaker who utters (2) is not saying what she would say with (2*) or any other sentence with a completed co-denoting description. Rather, she indeterminately (or sort of or vaguely) says them all, for every ‘one of each definite description that could be used to sharpen what [the speaker] vaguely meant’ (Schiffer, 1995, 377; Blackburn, 1988, 271). What’s expressed with an utterance of (2) is true only if everything indeterminately expressed is true.

In uttering (2), what is indeterminately expressed depends on what counts as an admissible sharpening. But what are these? Is the class of admissible sharpenings of, say, ‘the murderer’ determined by what a speaker has in mind at the time of utterance? Can a speaker deem both ‘the murderer of that victim’ and ‘the murderer of John Smith’ as equally good sharpenings? Must these completions be present in a speaker’s or auditor’s mind? Or, instead, must a speaker be prepared to ‘fall back on’ them if asked to identify which supplementation he meant (Blackburn, 1988, 271)? This suggestion faces a familiar worry: suppose an admissible sharpening is determined by false belief. Need that render what a speaker says with (2) without truth-value, or even false? It doesn’t seem so; suppose she based her utterance of (2) on the condition of an unknown victim. How, then, can false beliefs be excluded in determining admissible sharpenings? If they cannot, how do we explain how something true can be said with an incomplete description?

We might stipulate that ‘the murderer*’ counts as an admissible sharpening of ‘the murderer’ only if a speaker would not count her utterance of (2) as true unless the murderer* is insane. ‘the murderer*’ lacks criterial status if a speaker allows that what she asserted would be true even if it turned out that the murderer is not the murderer*. Since learning that the victim is not John Smith, she would cease holding true ‘The murderer of John Smith is insane’, does it follow that ‘the murderer of John Smith’ is not an admissible sharpening of her use of ‘the murderer’? What’s left? The problem of identifying acceptable sharpenings, I expect, is irresolvable, but I won’t exploit my expectation in what follows. (See Fodor and Lepore (1996) for further arguments against the coherence of an admissible sharpening.)

²⁸ These next few sub-sections have greatly benefited by invaluable input by Lou Goble, Tim Williamson, Herman Cappelen, and especially, Jeff King.

²⁹ I will run my argument against strategies of the first denomination, but it will be evident how to extend it to proposals of the second denomination as well.

(28') The table there is there.

These various supplementation strategies are *semantic accounts* about what can be said (or expressed) by sentences with (incomplete) definite descriptions and so complaining that (28') looks less informative than (28) needn't be a problem. After all, my present utterances of 'I am here now' and 'Ernie Lepore is at Rutgers University on 3 June 2000' aren't equally informative, but on at least one respectable semantic story, these sentences can be used to say (or express) the same proposition. If the same proposition can be expressed by non-synonymous sentences, where one, in some sense, is analytic and the other is not, then perhaps that sort of explanation extends to (28) and (28').

A genuine challenge to supplementation strategies must establish that, once appropriately contextually relativized, (28) and (28') do indeed express distinct propositions. Here are considerations intended to support this challenge.

Neither (28) nor (28') have true necessary readings, yet restricted to appropriate worlds, the latter, but not the former, seems to have one. This argument can be run on appropriately conditionalized versions of (28) and (28'); or by restricting evaluations to denoting worlds. I prefer to take the latter route.

In order to run this criticism, there is a restricted sort of necessity we want to consider. We will restrict ourselves to worlds where our definite descriptions denote. In all such denoting worlds, (28') is true. This is why it seems trivial because we often confine our attention to such worlds. We (and Kripke encouraged us to) do a similar thing in asking whether it is necessary that Bill Clinton is human or whether it is necessary that the teacher of Alexander taught Alexander. Kripke uses this sort of argument against a version of the descriptive of names when he writes: 'If "Aristotle" meant "the man who taught Alexander the Great", then "Aristotle was a teacher of Alexander the Great" would be a mere tautology. But it isn't' (Kripke, 1980, 30). Kripke is clearly expecting that we limit ourselves to considering worlds in which Alexander had a unique teacher (so that 'the teacher of Alexander' denotes). His is that 'The teacher of Alexander taught Alexander' has this kind of limited necessity, and 'Aristotle taught Alexander' doesn't.

We are making the same claim about (28) and (28'), namely, in any world in which the definite description in (28') denotes, it is true (thus, our feeling that it is trivial). But this is not so in any world in which the definite description in (28) denotes (thus, our not feeling that not any token of this sentence expresses something trivial). Minimally, it's worth pointing out that embracing a supplementation strategy requires accommodating such necessary interpretations.

Much along the same lines, but worse, consider the sentence 'The table isn't there.' It doesn't seem to have a reading under which it expresses a necessarily false proposition, but supplementation strategies predict that it should, since, assigning its negation narrow scope, assuming both tokens of 'there' co-index, any utterance of 'The table there isn't there' expresses a necessarily false proposition (again, with evaluations restricted to denoting worlds). Represented as '[The x: table at l (x)]NOT(x is at l)', to be true, there would have to be a unique table at l of which it is false that it is at l. (On a standard

Russellian account, the narrow scope reading could be true (e.g., if there is not a unique table at *l*), since it could be represented as, ‘NOT[The table at *l*: *x*](*x* is at *l*).’ This strengthens the critical point, because one would expect the description to take wide scope over negation (compare other examples), and so the most natural reading of ‘The table there isn’t there’ couldn’t be true; and its other reading is true only if there is no unique table at *l*.

Though not knockout punches, the accretion of these counter-intuitions succeeds in clarifying commitments inherited from adopting a supplementation strategy.³⁰ Each of these various cases appeals to speaker intuitions, but these intuitions seem no less firm than whatever other intuitions were supposed to have motivated supplementation in the first place.

3.3 Screening off

A proponent of one or another supplementation strategy might protest that these alleged troublesome interpretations are never available. Perhaps a policy of screening off is in place, thereby, excluding any contextually determined supplementation that is explicitly referenced in a predicate (or elsewhere). So, e.g., take a location *l*—if explicitly referenced by an utterance of (28), it is screened off as a candidate completer.

(28) The table is there.

This constraint is supposed to pre-empt contextual completions that would render some seemingly contingent sentences as expressing necessarily true or false propositions. Any item referenced, or expressed, explicitly is thereby rendered contextually irrelevant, at least for the purposes of supplementing a token of an incomplete description. If this ploy can be made to work, then what’s expressed with (28) cannot be what would be expressed with (28’), but, perhaps, instead what would be expressed with (29).

(28’) The table there is there.

(29) Your table is there.

Contextually supplementing a token of ‘the table’ by indexing an addressee secures uniqueness without rendering an utterance of (28), relative to the same assumptions, as expressing a necessity.

Whatever can be said in favour of screening off surely reflects nothing more than handy wisdom about the pragmatics of sound interpretation, and nothing about semantics. For how can semantics prohibit an explicitly referenced object from also being most contextually salient?³¹ With uses of potentially complete definite descriptions, as

³⁰ Of course, our intuition is also that we *cannot* express with an utterance of ‘The table is there’ a necessarily false proposition. But what prevents, on semantic grounds alone, a supplementation from issuing in a proposition that expresses what ‘The table here is there’ would, with distinct locations indexed. Or, an utterance of ‘The murderer didn’t murder him’ from expressing what ‘The murderer of him didn’t murder him’ would, with the same individual indexed? That is to say, what prevents the location of the utterance and the demonstrated victim from being the most contextually salient items?

³¹ Or worse, harkening back to the last footnote, the screening off strategy would somehow have to be made to work so as to prohibit internal incoherence, as with ‘The table here is there.’

in ‘your only table’ or ‘the table there,’ an explicitly referenced person or location is typically also contextually salient. Indeed, why can’t a speaker stipulate beforehand that she wants a certain location (and that location alone) to be the most contextually salient aspect of our conversation, and then proceed with (28)? Even in such circumstances, I presume, what she expresses is not, in any sense, rendered necessary.

Furthermore, even if some sort of screening off strategy could be made to work in the simple cases we have been discussing—though I don’t see how—it wouldn’t help to avoid counter-intuitive consequences for more complicated cases.

Contexts may or may not be accurately representable as sequences of items which context-sensitive expressions can take as semantic values (speaker, addressee, time, place, topic of discourse, perceptually salient objects, etc.),³² but I presume it’s not controversial that in any given context at most finitely many items are salient. Let C be an ordering $\langle a_1, a_2, \dots, a_n \rangle$ of every salient item, and then try to denote a table with a use of a sentence of form (30),

(30) The table $\phi(i)$,

where ‘ i ’ indexes a member of C , and ‘ ϕ ’ specifies the predicate of which ‘ i ’ is a constituent. ‘ $\phi(i)$ ’, e.g., might be ‘is there’, or ‘is here’ or ‘is that’, or ‘is yours’, etc. None of these envisaged instances of (30) seems to express a necessary truth (or falsehood) in C , yet counterparts of form (31) can,

(31) The table $(i_1) \phi_1(i_2)$.

(as would ‘The table there is there’, ‘The table here is there’, and so on). Suppose, because ‘ i_2 ’ occurs in the predicate in (30) it follows that *distinct* contextually salient items must be indexed by ‘ i_1 ’ and ‘ i_2 ’ in (31). But consider a new sentence with enough disjoint predicates such that each item in C can be picked out by a distinct index, as in (32).

(32) The table $\phi_1(i_2)$ or $\phi_2(i_3)$ or ... $\phi_m(i_n)$.

According to supplementation proposals, an utterance of (32) expresses in C what (32’) would.

(32’) The table (i_1) is $\phi_1(i_2)$ or $\phi_2(i_3)$ or ... $\phi_m(i_n)$.

But then (32), which expresses a seemingly contingent claim about the contextual salient setting of what, if anything, ‘the table’ denotes, has been transformed into a sentence which in that same contextual setting has a (nearly) necessary reading, without a possibility of further screening off. That we can devise such sentences might convince you something is fundamentally wrong with supplementation strategies.

By promiscuously forcing the *semantics* to acknowledge contextual aspects not explicitly represented as determinants of what is expressed on an occasion of use, supplementation strategies compromise a pragmatic, but semantically independent, alliance

³² Montague, 1974; Kaplan, 1989a; and Lewis, 1970.

between context and linguistic meaning. No one can (or should) deny that contextual salience is exploited by the speaker and auditors in their effort to converge on some point or other. But most of us thought we were being subtle and creative in so doing and not that we were obliged to as a matter of meaning alone.³³

4. Pragmatic Proposals

Confronted with what seem like insuperable (or, at least, a lot of troublesome) difficulties for supplementation strategies, one response is to resign oneself to every token of ‘the F’ as non-denoting if incomplete. As things are (every token of) (1) and (2) are not true. This does not mean one cannot communicate or convey something true using one, but in this case the proposition communicated or conveyed is not identical to whatever the uttered sentence expressed.

It might be enough to explain (away) the intuitions driving the search for a semantic solution, namely, that something true was said, by establishing that (1) and (2) can be used to convey *unsaid* truths. Mutual knowledge between speaker and auditor that there are many Fs triggers a search by auditors for what a speaker meant to convey but did not express. In uttering (1) or (2), a speaker gets across to auditors, perhaps, that the table there is covered with books, or his murderer is insane (Grice, 1969, 142–3; Bach, 1987, 103–4; Sainsbury, 1979, 115–16; Blackburn, 1984, 308–10; Reimer, 1998a, 100–1).

A speaker unrestrictedly expresses, or literally says, with an utterance of (1) that exactly one table exists and it is covered with books. Because any such utterance is obviously false, or lacks relevant specificity, given mutual knowledge that many tables exist, auditors will interpret a speaker as aiming to convey something else, perhaps, that the table over there is covered with books. There may be no unique proposition a speaker conveys but fails to express, if that requires something unique she has in mind to convey or a particular object she means to draw an auditor’s attention to. But indeterminacy or

³³ Reimer (1992, 360–2) also expresses scepticism (for different reasons) about supplementation strategies applied to incomplete definite descriptions, but nods with approval when they are extended to under-specified quantifier expressions. So, suppose when someone utters (8), she expresses what (8*) does,

(8) Every student must hand in homework.

(8*) Every student (i) must hand in homework.

where ‘i’ picks out a contextually salient set that restricts the domain of ‘every student’. Suppose ‘i’ picks out a set of objects in a certain classroom. Now consider (8**).

(a) Every student is among i.

If (8) can express a truth, why can’t (a)? If it can, then, since (8*) need not quantify over every student alive, shouldn’t the same be true for (a)? Suppose, for some utterance of (a), the set indexed by ‘i’ is the same as that allegedly restricting the domain of quantification for ‘every student’. Then (a) would have a reading under which it is necessary, but it doesn’t seem to, unless ‘i’ picks out the entire set of students.

Analogously, there is a reading of ‘Every student in this class is not in this class’ under which it expresses a necessarily false proposition, but not one for ‘Not every student is in this class’. We might try screening off, but consider a sentence like ‘All my brothers are Norwegian’, where there doesn’t seem to be a restriction. If supplementation strategists claim that in such cases what’s indexed can be identical to an explicit domain—how can they not?—then that’s incompatible with screening off as described.

falsity needn't hinder a pragmatic explanation, since speakers often convey more than what they express; and some of it may even be underspecified or false.

In short, on a pragmatic explanation, an utterance of (1) doesn't prohibit something true from being conveyed.³⁴ How speakers succeed in conveying unexpressed truths involves a non-linguistic story about speaker intentions, mutual beliefs, and so on—a familiar story about speaker meaning (Kripke, 1977), and so, about common-sense psychology.

4.1 Problems with pragmatic explanations

Some sceptics find pragmatic explanations of the phenomena of incompleteness intolerable. According to Neale, 'in view of the fact that context-dependence is such a ubiquitous feature of the use of natural language, it seems likely that the [various semantic strategies for yielding denotation] yield predictions more in accordance with our intuitive ascriptions of truth and falsity' (1990, 114–15; see, also, Soames, 1986, 300; Reimer, 1992, 361, 1998a, 102–3; Stanley and Szabo, 2000). Apparently, these sceptics' intuitions about what's true are stronger than the pragmatists'.

Larson and Segal find pragmatic explanations unappealing because 'one uses a clear falsehood to convey what one believes to be a truth only for some definite reason: out of politeness, coyness, a sense of drama, or, at the very least, for want of a better option. But in these cases there is no such motive: speakers are not being ironic or arch or coy' (1995, 329).³⁵ According to Larson and Segal, in order for it to be true that some proposition is speaker meant but not expressed, certain conditions must obtain—which seem unfulfilled with most acceptable uses of incomplete definite descriptions.

³⁴ The same strategy might be tried on quantifier underspecification. When a speaker utters (8), what she says is manifestly false, but in that conversational setting, speaker and auditor have mutual knowledge that this is so, and so auditor concludes that speaker is trying to convey a truth, e.g., that every student enrolled in class must hand in homework.

³⁵ According to a pragmatic explanation, sentence (a) can be used to express a true proposition just in case every bottle in the universe is on the only table in the universe.

(a) Every bottle is on the table.

An audience may use general principles to infer from information available in context that a speaker means a different proposition, maybe one with a more restricted domain.

Stanley and Szabo (2000) allege that the phenomenon of quantified contexts, which involves sentences with multiple quantified expressions whose intuitive readings can be captured only by assuming a variable representing the quantifier domain of a second quantifier expression bound by a first, refute pragmatic accounts. Since pragmatic explanations do *not* postulate syntactically represented quantifier domains, they cannot capture such readings. In particular, consider (b).

(b) In most of his classes, John fails exactly three Frenchmen.

The domain of the second quantifier varies with values of the first. (b) means '[For most x: John's class (x)][Exactly three y: Frenchmen in x (y)](John failed y)'. The indexical in 'exactly three Frenchmen' is bound by 'In most of John's classes'. Since pragmatic approaches posit *no* variable with 'exactly three Frenchmen', they fail to interpret (b) such that its second part is semantically related to its first. They cannot posit a restricted quantifier domain for the second quantifier in (b) because they believe it is illegitimate to postulate structure on semantic grounds alone. If (b) establishes bindable variables for quantifier domains, then arguing that quantifier restriction must be treated pragmatically is akin to arguing that any reading in which 'his' is bound in (c) should also be captured pragmatically.

(c) Every boy loves *his* mother.

Pragmatists would be unmoved by these attacks, for, I suppose, they doubt we have unimpeachable intuitions about what's true—either the propositions expressed or those conveyed (or speaker meant). In response to Larson and Segal, they would say we shouldn't be too narrow-minded about the circumstances under which we can convey truths that are not expressed. Still, there are serious problems with pragmatic accounts, ones they share with semantic accounts, namely, both appeal to a notion of what's literally said and both ignore that *what is said* often *exceeds* what is either expressed or conveyed. I want to challenge the coherence of their shared notion of what's literally said and also its utility.

5. What Is Said

According to Scott Soames, 'the fundamental task of a semantic theory is to tell us *what sentences say* in various contexts of utterance. On this view, the meaning of a *sentence* can be thought of as a function from contexts to what is said by the sentences in those contexts' (1989, 394, my emphasis). Soames' picture is commonplace among participants in this debate about incompleteness, both pragmatists and semanticists.³⁶ In what follows I will argue that this picture is wrong.

Briefly, here's how I see the relationship between sentences and their utterances: sentences have truth conditions, as determined by their meaning. Utterances are used to make statements (i.e., to say things), ask questions, and perform other kinds of speech acts. Utterances typically (but not always) are of sentences, but what's said, what's asked, etc., by an utterance can depend not only on the truth conditions of the sentence uttered, but also upon a number of other 'non-semantic' features of the context of utterance. In short:

- a. The truth conditions of a sentence S need not correspond to what's said or stated by an utterance of S.
- b. What's said by an utterance of S can be true, even though the truth conditions for S aren't satisfied (and *vice versa*).
- c. What's said by an utterance of S can be 'about' something (e.g., NYC), even though the truth conditions for S makes no reference (to that thing).
- d. Because of (a)–(c), intuitions about utterances of sentences can in no simple and direct way be used as guides to the truth conditions for those sentences.

Here I will briefly show how to support this general framework.³⁷

In Cappelen and Lepore (1997a and 1998a), we pointed out that reports of the form 'A said/asserted/stated that p' can be true even if A never uttered a sentence, appropriately

³⁶ The connection with pragmatic accounts is this: for anything to be pragmatically (or conversationally) implicated some specific p must be said, where the fact that p was said must be ascertainable and specifiable independent of and antecedently to fixing what the conversational implicatures were given that it is p which was said.

³⁷ See also the discussion in Chapter 2 of this volume.

contextualized, that is true just in case *p*. Since intuitions about what's said by an utterance are fixed by acceptable indirect reports, it follows that (a)–(d) are correct. I'll rehearse some data and arguments starting with rather obvious and trivial cases where there's a distinction between the truth conditions of the sentence uttered and what is said by that utterance.

Al, convinced that Stanley is Smith's murderer, says, looking at Stanley, 'Smith's murderer didn't comb his hair today'. Bill wants to report Al's utterance to Harriet, whom Bill knows is convinced of Stanley's innocence. Since Bill knows Harriet is unaware both of Al's contrary belief and of the context of Al's utterance (staring at Stanley), he might report Al to Harriet as having said that *Stanley didn't comb his hair today*. His so doing is appropriate, but is it correct? Even a philosopher whose deepest conviction is that definite descriptions aren't ambiguous between referential and attributive uses can accept Bill's report of Al's utterance as true.³⁸

Consider another sort of case, where Richard, after a particularly awful philosophy talk, remarks with heavy sarcasm, 'That was *really* good'. In most contexts it would be a mistake to report him as having said that the talk was really good—which is exactly what his words mean—but he can be reported correctly as having said that he didn't like it much. Here we are obliged to appeal to the negation of what his words express to correctly articulate what he said.³⁹

In both cases, pragmatic features of a reported utterance (a speaker meaning or an indirect speech act) get encoded into what was said. If correct reporting is a measure of what's been said—how could it not be?—how shall we account for the intuition that speakers can *say* something true with an incomplete description? Such talk, proponents of a more austere conception of what's said might argue, reflects nothing more than our practice of including into what's said more than whatever proposition was expressed. So, on this picture, what's said also, sometimes at least, includes what is meant but not expressed.

This straightforwardly rejects the identification of what's said with what's expressed, and so is incompatible with both pragmatist and semanticist strategies. Should we conflate what's said with what's meant? Doesn't the data recommend our doing so? Or, at least, doesn't it acknowledge that there is an ordinary notion of what's said that permits such conflation? But even accepting this conflation will not suffice to explain all of our intuitions about our uses of incomplete definite descriptions.

5.1 *How the context of a report can determine what's said*

That what's said with an utterance can be partly a function of context is old hat. For any utterance of 'That's a toy', what is thereby said is partly determined by a context of utterance,

³⁸ Though I won't argue for it here, it's easy to see that we can use a description attributively to report someone who used it referentially, and *vice versa*.

³⁹ Note that one cannot mimic the first speaker's sarcasm to get across that sarcasm. At best this indicates sarcasm of the reporter—or a direct report.

in particular, partly by which object gets demonstrated. So, what's said with utterances of this sentence can change even though conventional meaning does not. As we have seen, some semanticists hold that a speaker can utter (1), gesturing towards a table in front of him, and thereby say that the table there is covered with books. Appropriately contextually placed, a speaker can be correctly reported as having said truly that the table there is covered with books. On the basis of such considerations, some semanticists conclude that (1) must be contextually sensitive to the extent that its utterances can factor in contextual salience in rendering its incomplete description denoting. Some pragmatists conclude, instead, that, through speaker meaning, an utterance of (1) can convey a true proposition. But how can either semanticists or pragmatists explain the following sort of case?

Suppose that whatever table is under discussion currently sits comfortably in the speaker's father's office. Has a speaker thereby said or even meant with his earlier utterance of (1) that the table in his father's office is covered with books? By whatever means semanticists and pragmatists want to explain how a speaker with an utterance of (1) can say or mean that the table there is covered with books can their explanations extend to this sort of case? But how could either account factor being in his father's office into what was said or meant, since a speaker himself might be ignorant or misinformed about what is or isn't in his father's office—maybe that table got moved there after he spoke?

One reaction to this sort of case is that, though someone might *report* an utterance of (1) as saying that the table there is covered with books or that the table in his father's office is covered with books, since these reports disagree about what was said, at least one of them *must* be wrong. The soundness of this reaction requires holding that an utterance of a (unambiguous) sentence can be used to say no more than one thing—which is exactly what Soames' notion of what's said requires. So anyone who reacts this way to our alleged data must hold that:

The relationship between an utterance of a sentence and what's said by that utterance prohibits a speaker from saying both that *p* and that *q* with a single utterance, if *p* and *q* are distinct propositions.

But, then, how are we to explain that an utterance of 'Rudy loves New York and New Jersey' can say that Rudy loves New York, as well as that he loves New Jersey, as well as that he loves New York and New Jersey?⁴⁰

Amending this restriction on what can be said with a single utterance, as Soames suggests, with 'unless the proposition ("immediately") follows from whatever is said' fails to explain how Bush, in detailing a new economic programme, says that he will not cut taxes—even though none of his words express or imply this proposition. Is that what Bush really said? Why yes. But did he literally say it, or strictly speaking, is this

⁴⁰ Notice that if Rudy loves New Jersey but not New York, then, though the original utterance is false, at least one report of it attributes a truth, namely, Rudy said that he loves New Jersey.

what he said? Are you asking me for a direct quote? If not, then that's what he literally or strictly speaking said. Can you imagine his denying it; or, the press recanting, 'Well, yes, he sort of said it, but he didn't literally say it?'

Or, take Francois, in uttering 'Amethyst is Maria's favourite colour', not only says that Amethyst is Maria's favourite colour, but also that *the colour of that stone* is Maria's favourite colour. The latter clearly does not follow from what his words expressed. But if he hasn't said both, why, then, is it acceptable to report him as such? Indeed, it is *inappropriate* to report him as saying that amethyst is Maria's favourite colour, if you know your audience is unfamiliar with the colour word 'amethyst'.

To continue, consider a professor who, when asked whether Alice failed her exam, replies 'I failed no one'. Has he said that Alice passed? If not, why is it correct to report him as such, in which case what's said makes reference to Alice, though his words do not? Imagine telling poor Alice that her professor didn't say whether she failed or not. Or, that, strictly speaking, or literally, or actually he didn't say. That would be a lie.

A sceptic might conclude that actual reporting practices are irrelevant in determining what's said. But how could anyone reach this conclusion without letting theory override practice? Still, one might protest, there are limits to what can be said with an utterance. Who would deny that? These various examples establish no more than that delimiting a priori what those limits are is not only difficult but inadvisable. For example, suppose Frank utters (1) in a context with only one table present. Later, after another table is brought in, a question arises about what Frank said with his earlier utterance. On the sole table present when Frank spoke now sits a vase, and on the added one sits nothing. Anyone who reports Frank in this context as having said that the table is covered with books might grossly misrepresent what he said, but anyone who reports him as having said that the table with a vase on it is covered with books has gotten him just right. If you disagree, how would you in the context described usefully and correctly answer the question what did Frank say? Would you conclude that you cannot?

5.2 Diagnosis

According to Frege (or at least Fregeans), what's said by a single utterance is a single proposition. On this picture, a speaker says that *p* just in case he assertively utters a sentence that expresses the proposition that *p*. Accordingly, a speaker's words make reference to, or denote, an object just in case what he says with those words does as well. Therefore, any indirect report of his utterance according to which what he says makes reference to, or denotes, an object that neither he nor his words does, *misrepresents* what was said. However, what my various examples are supposed to have established is that correctly determining what's said by an utterance often requires attending to *non-interpretive, non-semantic* considerations.

I can imagine a pragmatist applauding my efforts but he shouldn't, since analogous critical points apply to him. He, unlike his semanticist opponent, is more generous about what can be got across (but not said) with a single utterance. But he's still too restrictive, since he too wants to ascribe a special status to what's literally said and, even

more seriously, he wants to limit what can be conveyed or said to whatever proposition is determined by a speaker's communicative intentions. The data adduced suggest that both the pragmatist and the semanticist are myopic. In addition to those cases already surveyed, consider these other particularly relevant sorts.

John utters, 'The new seven-foot centre for the New York Knickerbockers played last night'. In fact, no seven-footer plays for the Knicks, but a new centre has just joined the team and he is just under seven feet tall. Asked by someone else whether this new centre played last night, I correctly report John as having said that *the new centre for the New York Knickerbockers* did indeed play last night. If this report is right, then I can correctly represent John as having said something true, even though his original utterance is false. I also correctly represent him as denoting someone, even though his words do not.

Or, suppose, you tell me, 'I saw John wearing his blue jacket this morning.' Suppose further than John owns exactly one jacket and it's definitely not blue, though it's easy to see how someone not in the right lighting might think otherwise. If asked by someone else whether John was wearing his jacket this morning, I can correctly report you as having said that he was wearing *his jacket*. Again, context positions me to correctly report you as saying something true and also as denoting his jacket, though your exact words were false and none of them denoted John's jacket.

A distinct proposal for explaining speaker intuitions about acceptable uses of obviously incomplete definite descriptions and underspecified quantifiers emerges from this data. Incomplete definite descriptions are non-denoting, just as Russell taught us and nothing semantically restricts the range of so-called underspecified quantifiers. The long critical discussion in this chapter is supposed to have established all of this. Of course, nothing close to a proof was (or can be) on offer that no successful semantic account could ever be devised. But as evidence mounts, it does seem that any such effort will flagrantly violate other semantically innocent and theoretical intuitions we have. And, to boot, intuitions that underlie efforts to construct semantic accounts to render incomplete descriptions denoting, and underspecified quantifiers sufficiently restricted, can be *explained away* by appealing to overwhelming data that what's said with utterances of sentences involving such expressions (though neither expressed nor speaker meant) denote and are appropriately restricted. To see how this can be possible, we have to remind ourselves that when we try to represent or articulate what's said by an utterance we aim to characterize a speaker's act (that utterance). In so doing, our interests often are not in systematicity or generality, but rather our aim is to determine something about a particular act in a particular context C in order to pass it along onto to a particular audience situated in a (perhaps very) different (sort of) context C*. In effect, our practice of reporting others treats what's said as a four-place relation between a sentence and its context of utterance and a reporting sentence and its context of utterance. In determining what's said we obviously draw upon information about specific intentions, knowledge, and history of the speaker in C and, not so obviously, we can also draw upon *like features of C**, the context in which we report what's said.

That a context of a report can influence a correct assessment of what's said by an utterance makes it quite difficult to imagine how a pragmatic account, no matter how liberally construed, can explain what's said by an utterance. When our aim is to explain how what's said, with an incomplete (attribute) description, can be true or denoting, even though, according to the best semantic account available, nothing is denoted, it may be that both what his words mean and the pragmatics of their use fail us. Our reporting practices clarify that semantics should not a priori constrain what can and can't be said by an utterance. Competent speakers make such judgements all the time, often relying on information that exceeds anything expressed or meant. This competence consists, in part, in a capacity to judge whether a report about what's said is accurate or misleading. Theorists who try to systematically incorporate contextual cues into a semantic account of what's said seek to theorize about a practice that does not admit of it. There is no reason to believe that determining what's said will be simpler or more systematic than determining whether two items are similar.⁴¹

6. Conclusion

One intended consequence of my present effort has been to reclaim determinations of what's said from philosophers who pirated that notion as constraining the semantics for natural language. Philosophers as diverse as Kaplan and Grice have been telling us for so long that semantics aims to articulate what words say in contexts of use that it's become banal to agree with them. In their very different frameworks, what's said and the proposition expressed are interchangeable notions. The upshot of our discussion is that, though there may be a technical notion of *what's said* for which this identification is correct, if our actual reporting practices are to be any sort of clue to a determination of what's said by another's uttered sentence, then what's reported as said and uttered sentence reported often disagree in meaning and extension, right down to denotation and reference, and up to truth.

Readers might fear that should semantics not be supposed to be occupied with determining what's said by words used, then what's left for it to do? Fret not.

- Most semanticists were weaned on model theoretic notions, like validity, tautology, inconsistency, and nothing in this discussion challenges the connection between truth and meaning, and so, the connection between these logical notions and meaning.

⁴¹ Salmon (1991, 88–9) seems to endorse a similar view, though his use of 'literally saying' vs. 'the loose or popular sense [of 'say']' is *incompatible* with the facts. There is nothing loose about reporting what was said by an utterance of (1) as described above; indeed, each may be a literal report. Furthermore, in correctly reporting an utterance with a complement that does not express the proposition expressed, Salmon arbitrarily constrains what's acceptable, suggesting, wrongly I believe, that such departures are disguised *de re* reports (see, p. 88). See, also, Cappelen and Lepore, 1997a.

- Many semanticists first turned from formal languages (and their models) to natural language semantics, because of concerns about productivity and systematicity. It is (virtually) non-negotiable in the philosophy of language and linguistics today that no semantic theory can account for these aspects of natural language unless it is compositional, and so, one venerable project has been, and continues to be, to show how various intriguing fragments of natural language admit of a compositional semantics. Nothing in this discussion challenges the legitimacy of this project.

Indeed, as far as I can tell, no legitimate endeavour for semantics is impugned by denying that determining what's said requires a semantic explanation.

It has taken me all these pages to try to convince you that 'the table' means *the table* and that 'every student' means *every student* everywhere they occur. Don't believe anyone who tells you otherwise. I've also been trying to remind you that what happens to our words with their fixed meanings once they leave our mouths is often beyond our control.