

This, That, and The Other

*Stephen Neale**

1. Introduction

Debates over the correct way to accommodate so-called referential uses of definite (and to a lesser extent indefinite) descriptions, have aroused passions for half a century or more. The central controversy concerns the following question: does the existence of referential uses, which no-one seems to doubt, have any bearing on the meaning of the definite article and on what a speaker says by uttering a sentence of the form ‘the ϕ is ψ ’ on a particular occasion (or on what he *asks* by uttering a sentence of the form ‘is the ϕ ψ ?’). According to the Unitarian School, referential usage is a phenomenon readily explained without appeal to a referential semantics: what someone says by uttering a

* Longer and shorter versions of this chapter have been knocking around for some years, gathering strength, or at least steam, from the comments of readers or people subjected to this or that section in lectures, seminars, or conversation. Of necessity, the discussions of linguistic pragmatism and contexts in Sections 2 and 3 overlap considerably with the corresponding sections of Neale (forthcoming a). Particular thanks go to Kent Bach, Anne Bezuidenhout, Emma Borg, Ingar Brinck, Ray Buchanan, Herman Cappelen, Mark Crimmins, Michael Devitt, Stan Dubinsky, Kevan Edwards, Paul Elbourne, Jerry Fodor, Haidy Geismar, Owen Greenhall, Richard Hanley, Gilbert Harman, James Higginbotham, Jennifer Hornsby, Paul Horwich, Jerry Katz, Saul Kripke, Richard Larson, Barry Lee, Ernie Lepore, Paisley Livingstone, Colin McGinn, Anna-Sofia Maurin, Gary Ostertag, Angel Pinillos, François Récanati, Marga Reimer, Mark Sainsbury, Stephen Schiffer, Barry Smith, Jason Stanley, Kenneth Taylor, Dag Westerståhl, Matthew Whelpton, Deirdre Wilson, and Zsófia Zvolenszky for comments, questions, and advice. Talks at the following institutions led to numerous improvements, largely as a result of penetrating and persistent questions: Rutgers University; New York University; University College London; the School of Advanced Studies, University of London; Oxford University; the University of Maryland, College Park; the University of Arizona, Tucson; the University of Delaware; Tulane University; the University of Genoa, the University of San Marino, the University of California, Los Angeles, Lund University the University of Iceland; the Georg Brandes School, Institute for Nordic Philology, the University of Aalborg, and the University of Copenhagen. Special thanks go to Brian Loar, Colin McGinn, and Stephen Schiffer, with whom I have taught seminars on this material; to Joseph Almog, Tyler Burge, David Kaplan, Tony Martin, and Terry Parsons for comments during and after seminars at UCLA in which much of the material was beaten into a better shape; to Michael Devitt for reading through drafts, spotting all sorts of nonsense and untidiness, and making sure something actually got to press; and to Anne Bezuidenhout, Marga Reimer, Peter Momtchiloff and Rebecca Bryant for their patience and forgiveness. My indebtedness to the work of Dan Sperber and Deirdre Wilson, John Perry, Robyn Carston, François Récanati, and John Searle should be clear, and I make little attempt here to document this or that idea or piece of terminology. Finally, I gratefully acknowledge the generous support of Rutgers University, the University of Iceland, the Georg Brandes School, Institute for Nordic Philology, and the John Simon Guggenheim Foundation. The chapter is dedicated to the memory of Harry Snitcher QC.

sentence of the form ‘the ϕ is ψ ’ always has Russellian, quantificational truth conditions: the speaker is saying that (roughly) there is exactly one ϕ and every ϕ is ψ . As it is often put, the speaker is expressing a general proposition that might be represented in the notation of structured propositions as

- (1) $\langle\langle\text{THE}, \phi\rangle, \psi\rangle$.

The truth conditions of (1) might be specified by the following sentence of a formal language containing restricted quantifiers,

- (2) $[\textit{the } x: \phi(x)] \psi(x)$

assuming something like the following Tarski-style axiom for the quantifier *the*:

- (3) $[\textit{the}_k: \phi]\psi$ is true of a sequence s iff ψ is true of every sequence ϕ is true of differing from s at most in the k^{th} position, and there is exactly one such sequence.

The Ambiguity School, by contrast, maintains that the definite article ‘the’ is ambiguous according as it is used with its Russellian or its referential semantics: on the former, what is said has quantificational truth conditions as specified above. But if it is used with its referential semantics the speaker expresses a singular (particular) proposition that might be represented as

- (4) $\langle\alpha, \psi\rangle$

the truth conditions of which might be specified by the following sentence of the formal language:

- (5) $\psi(\alpha).$ ¹

Several questions need answering. (i) How does the Unitarian propose to explain the facts of referential usage? (ii) How does the Unitarian propose to explain the perfectly felicitous use of so-called *incomplete* descriptions like ‘the table’ or ‘the man’, the matrices of which are not uniquely satisfied? (iii) Why has the viability of the Russellian analysis engendered so much debate? (iv) Why does one group of distinguished philosophers seem so convinced that the unitary Russellian analysis of descriptions is correct (Simon Blackburn, William Blackburn, Hector-Neri Castañeda, Donald Davidson, Martin Davies, Gareth Evans, Peter Geach, Paul Grice, Stuart Hampshire, Saul

¹ The ambiguity in question is meant to be explicable or derivable in some way, an instance of *polysemy* rather than *homonymy*, more like the ambiguity in ‘horn’ (which is applied not only to the horns found on the heads of certain animals but also to certain musical instruments that are blown, early forms of which were made from animal horns) than the one in ‘bank’. Récanati (1986, 1989, 1993) and Bezuidenhout (1997) explicitly deny they are postulating an ambiguity in the definite article, claiming that on their accounts ‘the’ has a single lexical meaning that permits some utterances of ‘the ϕ is ψ ’ to express general propositions, and others to express singular propositions. I took issue with this sort of position in ch 3 of *Descriptions* (n. 36), and I am still inclined to think that, as stated, it is unacceptable; however, the desires that motivate it are highly instructive and I am inclined to think there is something valuable in it. Perhaps it just needs to be stated differently; see below.

Kripke, Stephen Neale, Mark Sainsbury, Nathan Salmon, John Searle, Scott Soames, David Wiggins)? (v) Why does another seem so convinced that a semantically distinct *referential* reading is also needed (Joseph Almog, Jon Barwise, Anne Bezuidenhout, Robyn Carston, Michael Devitt, Keith Donnellan, Jennifer Hornsby, David Kaplan, David Lewis, Chris Peacocke, John Perry, François Récanati, Marga Reimer, Bede Rundle, Stephen Schiffer, Robert Stalnaker, Howard Wettstein)?²

To these questions may be added another, voiced clearly by Devitt (1997a, 1997b, this volume), Devitt and Sterelny (1999), and Reimer (1998a), which the Unitarian needs to answer: (vi) Given that the referential use of descriptions is systematic, regular, conventional, and cross-linguistic, is it really plausible to maintain that no systematic facts about lexical meaning beyond those given by Russell's Theory of Descriptions—and seemingly accurate for attributive uses—are invoked where we have referential uses?

The question posed by Devitt and Reimer prompts three others: (vii) What is involved in *polysemy* or *systematic* ambiguity, and how does such a notion shape up with respect to Donnellan's talk of *pragmatic* ambiguity or ambiguity of *use*? (viii) Are there related phenomena in natural language that can be explained in terms of such notions? For example the behaviour of the third person pronoun 'his', which appears to admit of a use in (6) in which it is bound by 'John', and another in which it is merely co-referential with 'John':

(6) John loves his wife, and so does Paul.

Finally, (ix) What can we learn about the semantics and use of English definite, indefinite, and demonstrative descriptions, and English pronouns, by looking at the use of their counterparts in other languages, and what are the concepts driving the existence and use of such devices?

I shall sketch answers to most of these questions here, but it is impossible for me to talk about descriptions today without first discussing the philosophical background I assumed, without much comment, in *Descriptions* in 1990. It has become clearer to me with each

² Although I am going to focus on the conceptual pressures implicated in the descriptions debate, I am not entirely convinced that sociological factors have not helped shape it. Revolution and reaction are as commonplace in philosophy as elsewhere, and so are the personality types drawn to both. Some who have argued against a unitary Russellian account might have been drawn to the idea of giant-slaying, while some of those who have reacted to such assaults might have been drawn to preserving the logico-semantic *status quo*, taking comfort in the simplicity, transparency, and range of a theory proposed or endorsed by a giant. There appears to have been some geographical and genealogical clustering. Philosophers who teach or once taught in or near UCLA (for example, Almog, Donnellan, Kaplan, Lewis, Perry, Schiffer, Wettstein) tend to belong to the Ambiguity School; those who teach or once taught at Berkeley or Princeton (Grice, Davidson, Searle, Neale, Kripke, Soames, Salmon) tend to belong to the Unitarian School (the obvious exception being Lewis). The State of Maryland seems also to have once attracted Ambiguity Theorists (Devitt, Stich, and Wilson). The legacy of Carnap and formal language philosophy, and the influence of Kaplan may have played a part in the appeal of a referential interpretation around UCLA; and the influence of Grice and Kripke and the legacy of ordinary language philosophy may have contributed to the appeal of pragmatic explanations and the viability of a unitary Russellian interpretation at Berkeley and Princeton. No more speculation about personality and geography.

passing year that some philosophers and linguists writing about descriptions, pronouns, demonstratives, quantifiers, and names are involved in hybrid, formal exercises that have little to do with the project to which *Descriptions* was meant to be a small contribution. The nature of the overarching project is discussed in detail in forthcoming work, but I can and shall say just enough here to locate the ensuing discussion and render fully explicit certain syntactic, semantic, and pragmatic assumptions, particularly in connection with indexicals, syntax, and ellipsis, assumptions made by other like-minded philosophers and linguists, but which are not always announced or appreciated in the literature.

With the background out of the way I want to shed some pragmatist light on utterances of ‘this’, ‘that’, ‘the’, ‘a’, and pronouns (particularly in the third person). I shall suggest (i) that reflection on elliptical utterances leads us to a place from where the ambiguity debate between the Unitarian and Ambiguity Schools seems to lack real substance, and (ii) that there is much to be learned in reaching this position, in isolating the assumptions that originally engendered debate, and in understanding the various relationships that hold between descriptive devices and pronouns. In the course of all this, I want to concede an important point to those who have presented what I shall call the Argument from Convention for ambiguity, whilst rebutting new versions of the Argument from Incompleteness and the Argument from Anaphora, as well as something I shall call the Argument from Binding.³ The pragmatist synthesis on offer should, I think, be acceptable to Russellians and ambiguity theorists alike, especially those moved by the Argument from Convention. Finally, I want to look briefly at binding and lay bare an important pragmatist point that has repercussions for the interpretation of utterances containing pronouns and descriptive devices, thereby holding out the prospect of a unified theory of pronouns.

I am afraid much of the chapter is highly compressed and many things of relevance have been omitted to produce it, particularly in connection with abstraction, indexicality, the saying/implying distinction, speaker’s reference, and pronouns.

2. Linguistic Pragmatism

A central goal of the philosophy of language and linguistics is to explain how language works its magic or, more accurately, how *we* work our magic using language: how are we able to accomplish so much by making various noises or marks, for example how we are

³ I discussed four arguments in 1990, the Argument from Incompleteness (ch 3), Argument from Misdescription (ch 3), the Argument from Opacity (ch 4), and the Argument from Anaphora (Ch 5). My discussion of the Argument from Opacity inherits an anachronism upon which much of ch 4 was based (see Neale (forthcoming c).) The Argument from Binding was presented by Wilson (1991). I did not address the Argument from Convention, as I did not take it very seriously until I read Devitt’s and Reimer’s clear statements of its force. I imagine chapters in the present volume contain arguments to add to the present list of six. For an overview of the issues, see Ludlow and Neale (forthcoming).

able to express our thoughts, or how we are able to convey or request information about the world, and about our beliefs, wishes, feelings and so on so systematically and consistently?⁴ The noises and marks of particular interest are, of course, those belonging to, or at least governed by, systems called *languages*—which is not to say that examining non-linguistic acts of communication might not shed valuable light on the phenomenon of principal interest.

Several pre-emptive strikes against oversight were made by the phrasing of our question. The verbs ‘convey’ and ‘request’ were used as a gentle reminder that we use noises and marks not only to make statements but also to ask questions and do a good many other things besides—threaten, warn, baptize, marry, and so on. The word ‘feelings’ was bunged in as a reminder that we sometimes use language to convey things we are nonetheless inclined to say we ‘cannot express in words’.⁵ And the nouns ‘noises’ and ‘marks’ were used as a reminder that even if spoken language is viewed as primary or prior or dominant, the question must be answered it in such a way that the final theory extends to uses of language that do not involve speech.

Let us call a theory that aims to explain how hearers manage to identify what speakers are seeking to communicate a *theory of utterance interpretation*, or a *theory of interpretation* for short. To say that we are interested in providing a theory of interpretation is not to say we are prejudicing the issue against communication that does not involve speech or writing. It might simply turn out—it surely will—that a theory of interpretation will make reference to cognitive capacities involved in interpreting non-linguistic acts of communication, indeed non-linguistic acts more generally.⁶

Interpreting an utterance or inscription of a sentence involves substantially more than identifying and interpreting individual words and seeing how they have been put together to form that sentence. Probably there isn’t much it *doesn’t* involve, and it is hardly surprising that we have not yet succeeded in producing a theory of interpretation with much empirical clout. There have been successes in some of the *subtheories*—phonology and syntax, for example. But there is widespread suspicion that producing an overarching theory of interpretation will require nothing short of a complete theory of mind.⁷

⁴ It has been the hope of many philosophers that the attention paid to this question will pay dividends elsewhere, by clarifying statements of independent philosophical importance, for example, but I shall not be concerned with any of that here. Any account, however abstract, of how we are able to communicate will have to presuppose some sort of picture of what is involved in having a thought, a belief, or an intention to communicate. In principle, one’s position on the nature of psychological states may influence one’s final position on what is involved in communication. This is particularly true if one is attracted to something like Fodor’s (1975, 1992) idea of a ‘language of thought’.

⁵ Throughout, I use ‘word’ and ‘sentence’ where some people might be inclined to use ‘word-*type*’ and ‘sentence-*type*’ (similarly for ‘phrase’ and ‘expression’), as I am talking about abstract linguistic entities that may be ‘tokened’ through speech and writing at least. Thus I talk about the word ‘man’ and about particular utterances and inscriptions thereof.

⁶ See Grice (1989), Sperber and Wilson (1986, 1995), Carston (2002).

⁷ See, for example, Chomsky (2000), Davidson (1986), and Fodor (1983, 2001). For assessment, see Carston (2002).

The project of explaining interpretation has many components and involves people from several fields.⁸ Philosophers have two rôles, one in the board room, the other on the shop-floor, as it were. First, they will attempt to articulate clearly the nature of the project (distinguishing it carefully from various other projects with which it might be confused), distinguish clearly the various sub-projects, and distinguish and analyse the central concepts or at least the relations between them (for example, meaning, saying, implying, referring, and intending). At the same time, they will attempt to work alongside linguists whose expertise involves explaining how individual words are assembled into sentences and the extent to which communicatively relevant features of the sentences we use to say things depend upon features of the words out of which they are assembled and the mode of assembly itself. And alongside psychologists who can tell philosophers and linguists about cognition, in particular about the way we integrate information from different sources and channels in the process of identifying what someone is trying to communicate.

With distinct nods to the American Pragmatists, to Wittgenstein, Sellars, and Quine, and to Sperber and Wilson, I call the general outlook I have on the matter of interpretation, *linguistic pragmatism* (or *pragmatism* for short). It is an outlook that can be held by philosophers, linguists, psychologists, and no doubt others—no diplomas are checked at the door—who take themselves to be involved in the project of constructing a general theory of (utterance) interpretation, construed as an empirical theory and, as such, a contribution to cognitive psychology. It might be seen as a collection of theses that can emerge *only* in the context of attempting to articulate the outlines of such a theory, theses whose truth may well have repercussions elsewhere but which are not themselves motivated by the desire to bolster this or that philosophical or political doctrine. (The pragmatist outlook may well be implicated in various works of a ‘contextualist’ nature, but I am anxious to distance my own views from extant contextualist proposals in epistemology, metaphysics, ethics, and political philosophy, many of which seem to me rather suspect).

Some of the central tenets of linguistic pragmatism were accepted by a number of British philosophers in the 1950s, particularly J. L. Austin, P. F. Strawson, and (contrary to the claims of some pragmatists) H. P. Grice.⁹ But it was not until the late 1970s, by

⁸ *Descriptions* was (and still is) meant to form one small part of the project: to set out the relevant properties of definite descriptions and their components in a way that squares with and, in some sense, *explains* the various forms of empirical data whilst respecting vital conceptual distinctions, which may themselves be sharpened or refined in the process.

⁹ A somewhat simplistic picture of the relationship between the focus on ‘ordinary’ language and the use of ‘ideal’ or ‘formal’ languages appears to be accepted by many linguists and even some philosophers. The received view in linguistics appears to be that for some years there was a major philosophical conflict (between ‘formalists’ and ‘informalists’) which Grice somehow dissipated by distinguishing what a speaker said from what he ‘conversationally implicated’. (The picture is perhaps fostered by a naïve reading of the opening paragraphs of Grice’s ‘Logic and Conversation’ and by Strawson’s (1969) bizarre claims about a ‘Homeric struggle’ in his inaugural lecture ‘Meaning and Truth’.) Some of the people in the grip of this picture have been led to conclude that Grice was not actually a pragmatist. I know from conversations with

which time the Language of Thought hypothesis articulated by Fodor, the Chomskyan idea of LF as a level of linguistic representation, and important distinctions made by Grice and Searle had truly sunk in, that the conceptual resources were generally available to articulate the outlook clearly and in a form that made it relevant to more formal studies of language that were by that time blossoming in linguistics and philosophy departments in the United States. To the best of my knowledge, it was not until the work of Dan Sperber and Deirdre Wilson began to appear in print in the early 1980s that pragmatists made sustained efforts to render explicit the basic tenets of their work. Indeed, without Sperber and Wilson's work, and the work of Chomsky, Fodor, Grice, and Searle upon which it drew, philosophy and linguistics might still lack the distinctions and resources needed to say anything more substantive than the ramblings about 'contextual meanings' and 'relative meanings' that issue periodically from the darker areas of philosophy and linguistics departments (not to mention departments or 'programs' housing people unaccountably known as 'theorists' or 'philosophers').

It would be a mistake, I think, to attempt a *definition* of linguistic pragmatism as it is essentially an outlook that engenders a very practical approach to interpretation. I cannot go into the sort of detail I go into in *Linguistic Pragmatism* here, so I have produced twenty-four numbered and labelled paragraphs to give the general flavour of linguistic pragmatism as I see it. Some of the points are quite general or intuitive, others are very specific or theory-laden. Some are less central than others and could be withdrawn without upsetting the whole too much, but I am strongly inclined to go along with the whole lot, and there is no doubt that many gain strength through association with others. A few are held by some philosophers and linguists I would call anti-pragmatists; a few are rejected by people with a pragmatist outlook; and a few are conspicuous here by their absence elsewhere in the literature.¹⁰

him (a) that he saw the problem of providing an accurate account of what the speaker says when using an incomplete description as providing powerful *evidence* for pragmatism, and (b) that he never intended to be seen as denying pragmatism.

¹⁰ Putting aside differences of terminology and philosophical temperament, as well as apparent disagreements about particular analyses, I am inclined to view all of the following as operating in a broadly pragmatist spirit—although my own brand of linguistic pragmatism is, I suspect, rather too ascetic, inferential, beholden to ordinary language strictures, and driven by underlying concerns about practicality and the concepts of society and regulation for many of them: Austin (1962), Bezuidenhout (1997), S. Blackburn (1984), W. Blackburn (1987), Blakemore (1987, 2002), Barwise and Perry (1983), Carston (1988, 1993, 2002), Chomsky (1976, 1986, 1995, 2000), Crimmins (1992), Crimmins and Perry (1989), Evans (1982, 1985), Fodor (1987, 2001), Grayling (1995), Grice (1989), Neale (1990, 1993), Quine (1940, 1960), Papafragou (1998*a,b*), Perry (1986, 1993, 1998, 2001), Récanati (1987, 1989, 1993, 2001), Rouchota (1992, 1994), Searle (1969, 1979), Sellars (1954), Sperber and Wilson (1986, 1995), and Strawson (1950, 1952). I suspect pragmatism is also taken for granted by many others who take truth to be property of what a speaker says or expresses by uttering a sentence *X* on a specific occasion, and not of what *X* itself says or expresses relative to a context. I shall point to or draw liberally from the work of other pragmatists as I go. (In the case of Sperber and Wilson I shall not bother to cite: every page of the present essay can be viewed as containing an aphonic footnote that could have been rendered, as 'See the work of Sperber and Wilson, esp. their book *Relevance*'.) By saying I am a pragmatist, I do not mean to be saying that I endorse the details of Récanati's (1989, 1993) or Rouchota's (1992) or Sperber and Wilson's (1986, 1995) or Bezuidenhout's (1997) or

1. *Cooperation*

For the most part speakers (writers) want to be understood, and hearers (readers) seek to understand. To this extent they are involved in a cooperative exercise. *Ceteris paribus*, both parties tacitly assume they are using words with shared meanings, combining these words in accordance with a shared syntax, and operating in accordance with shared and very general, rational principles of *interpretation*.

2. *Meaning*

A theory of interpretation should explain how hearers (and readers) manage to integrate linguistic and non-linguistic information to identify what a speaker (or writer) *meant* on a given occasion by uttering (or inscribing) a linguistic expression *X*.¹¹ Valuable information can be gleaned from examining situations in which we report on speech acts using sentences of the form,

by uttering (or writing) *X*, *A* meant that *p*

where the reporter is using the expression replacing '*A*' to pick out an agent and the expression replacing '*X*' to pick out a linguistic expression, and where the expression replacing '*p*' is a declarative sentence. Examples:

by uttering, 'I'm tired', John meant that he was tired

by uttering, 'I'm tired,' John meant that he wanted us to leave.

We should be suspicious of locutions of the form '*X* means that *p*', where the expression replacing '*X*' is being used to pick out a sentence (e.g. 'the sentence 'snow is white' means that snow is white').

3. *Explanation*

To interpret is to provide an *explanation*, and the concept of interpretation makes no sense in the absence of a *problem* to be solved. We reflexively generate *hypotheses* about the things we perceive. Nowhere is this more in evidence than when we perceive one another's *actions*. We act out of *reasons*. To interpret an action is to form a hypothesis about the *intentions* behind it, the intentions that *explain* it. Interpreting a speech act is a

Carston's (2002) pragmatist analyses of definite descriptions and the attributive-referential distinction, all of which I find problematic. I shall propose an account that should be attractive not only to them but to Russellians and ambiguity theorists alike.

¹¹ Linguistic pragmatism does not necessarily assume there is much chance of ever producing an empirically interesting theory of interpretation. At least two pragmatists, Chomsky (2000) and Fodor (1983, 1987, 2001), have argued that asking for a theory of interpretation is tantamount to asking for a 'theory of everything', a complete cognitive psychology, because virtually anything can impinge upon the holistic process of interpretation. For more optimistic pragmatist outlooks, see Sperber and Wilson (1995, 1996, 2002) and Carston (2002). The present essay assumes neither outlook, but it brings into sharp relief the need for a clear picture of what the more tractable sub-theories of a theory of interpretation are supposed to do and how they must come together.

special case of this. The use of language is one form of rational activity, and the principles at work in the interpretation of linguistic behaviour are intimately related to those at work in interpreting intentional *non*-linguistic behaviour. What makes interpreting a speech act special is that a proprietary body of information, knowledge of language, is accessed immediately in the interpretation process. The hearer's or reader's goal is to identify what the speaker or writer meant. When this has been done, the interpretive problem has been solved.¹²

4. *Asymmetry*

The epistemic situations of the speaker and hearer are fundamentally asymmetric: the speaker *knows* what he means whereas the hearer has to *work it out*. If you want to find out whether I'm hungry (or in pain) you will have to watch me, see what I do, or ask me. I don't have to do that. I have 'privileged access' to that information.¹³ Similarly if you want to know whether I am worried about missing my flight, where on an aeroplane I prefer to sit, or whether I think Norway is a member of the European Union. And similarly where we have speech. Unlike you, I have privileged access to what I mean when I utter *X* on a given occasion. We can characterize a typical speech situation as follows. Person *A* intends to communicate something to some other person *B*. He selects a form of words *X* that he thinks will, in the circumstances, get across his point (and, perhaps, also get it across in some particular way or other. *A* knows what he means by uttering, 'That's his bank,' for example. He knows *which thing* he meant by 'that', *who* and *what relation* he meant by 'his' and *what* he meant by 'bank'.

B's situation is quite different: *B* is trying to work out what *A* meant and he must use anything he can get his hands on to get the job done since he has no direct access to *A*'s communicative intentions. The words *A* uses constitute partial evidence for what *A* meant. Other evidence may come from the physical environment, from *B*'s take on the conversation up to that point (if any), from *B*'s beliefs about *A*, and a whole lot more besides. The epistemic asymmetry of speaker and hearer underscores (i) the need to separate the *metaphysical* question concerning what *determines* (or fixes) what *A* means and the *epistemological* question concerning what is used to *identify* what *A* means, and (ii) the need to scrutinize simplistic appeals to *contexts*, *maxims of conversation*, *salience*, and *pragmatic factors*, which are frequently and mistakenly introduced together with

¹² I mean this to apply equally to interpretations in literary theory. The idea of textual interpretation makes no sense if there is no problem (about, for example, a word or phrase, a character, a plot, a work, or even a whole genre) to which an interpretation constitutes a possible solution. It is embarrassing that some 'theorists' who also call themselves 'philosophers' are unable to see this.

¹³ There are philosophers concerned to deny this idea today, but no coherent case has been made (indeed could be made) for a total failure of asymmetry. The idea of 'privileged access' to a state is often introduced with the idea of 'incorrigibility' (the idea that I cannot be mistaken about whether I am in the state in question). Whilst the case for denying this might be more promising, the fact that incorrigibility is at least *arguable* (and has been argued) in the first-person case, and is not is the least arguable (and has not been argued except, perhaps, by the deluded) in the third-person case, is enough to distinguish the cases here.

intentions in contemporary discussions as if these things conspire to bridge certain interpretive gaps. Scanning the context of utterance for salient objects and bringing to bear pragmatic principles (e.g. Grice's conversational maxims) is not going to provide *A* with any information that will help him identify what he meant. From *A*'s perspective, context and pragmatic principles have already played their rôles: *A*'s perception of the context—whatever a context turns out to be—his perception of *B*'s perception of the context, the assumption that *B* is operating in accordance with the same pragmatic principles as *A*, and *A*'s estimation of *B*'s ability to work things out (and probably a whole lot more besides) have already impinged upon whatever processes led *A* to use the particular form of words he used with the intentions with which he used them.

5. *Reciprocity*

Despite the epistemic asymmetry, the perspectives of *A* and *B* are not independent. The asymmetry is *reciprocal* or *complementary* as in adjoining pieces of a jig-saw puzzle. In producing his utterance, *A* relies on what he takes to be *B*'s capacity to identify what he intends to convey; *B* assumes that *A* is so relying. And, possibly, so on. The ways in which *A* and *B* operate form a *dovetail joint* and are *mutually sustaining*. And to this extent, there is simply no possibility of making sense of *B*'s capacity to interpret *A* without making sense of *A*'s capacity to exploit that capacity, and vice versa. So the project of constructing a theory of interpretation may be approached from either of two complementary perspectives, and an adequate answer must make sense of both.

6. *Intention*. What *A* meant by uttering *X* on a particular occasion is determined by, and only by, certain very specific *interpreter-directed intentions* *A* had in uttering *X*. The precise content of a psychological state such as a belief or intention may be determined, in part, by something external to *A* and beyond *A*'s control ('externalism'). Furthermore, the formation of genuine intentions is severely constrained by beliefs. I cannot intend to become a prime number, intend to digest my food through my lungs on alternate Tuesdays, or swim from New York to Sydney because (roughly) I cannot intend *what I believe to be impossible*. (There is no need to get into the exact force of the modal or the exact formulation of the constraint here. It is enough to recognize, as Grice (1971) does, that it is severe.) If, as Grice suggests, what *A* meant by uttering *X* on a given occasion is determined by certain interpreter-directed intentions, then assuming he is being co-operative *A* cannot mean that *p* by uttering some sentence *X* if he believes it is impossible for his audience *B* (or at least any rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as meaning that *p*. Among the things constraining *A*'s communicative intentions are *A*'s beliefs about the world, his (tacit) beliefs about the sorts of interpretive principles *B* will be employing, and his (tacit) estimation of *B*'s capacity to work certain things out (the list is not meant to be anywhere near exhaustive). So without some stage-setting *A* cannot mean that Jones is no good at philosophy by

producing the sentence ‘Jones has excellent handwriting and is always punctual’, for example, or by reproducing the mating call of some exotic bird.

7. *Factorization*

What *A* meant by uttering *X* may be factored into what *A* *said* (or *asked*) by uttering *X* and what *A* only *implied*.¹⁴ Thus, again following Grice, what *A* said and what *A* implied are determined by, and only by, certain very specific interpreter-directed intentions *A* had in uttering *X*.¹⁵ Nonetheless, although it would be perverse to insist upon a distinction between what *A* meant and what *A* *intended* to mean (and for good reason if Grice is right), a distinction between what *A* said and what *A* intended to say is not one obviously lacking a point. So, in the first instance, we should separate (i) *what A intended to say by uttering X on a given occasion*, and (ii) *what a rational, reasonably well-informed interpreter in B's shoes would think A intended to say by uttering X on that occasion* (which is not to say there are not problems with the idea of a *rational, reasonably well-informed interpreter in B's shoes*). In cases where (i) = (ii), we can talk freely about *what the speaker said*. (In cases where (i) ≠ (ii), certainly we *could* argue about which of (i) or (ii) or some third thing has the ‘right’ to be called *what is said*, but what would be the point? First, what third thing distinct from (i) and (ii) could be of any significance to a theory of interpretation? There is simply no rôle for a transcendent notion of what is said upon which (i) and (ii) converge when all goes well. Second, why is a choice between (i) and (ii) even needed in cases where (i) ≠ (ii)? Conceptually they are distinct, and they are both needed in a theory of interpretation. When all goes well, they coincide, and it’s just too bad they don’t always do so. Surely there is no philosophical payoff in bestowing the honorific ‘what was said’ on one rather than the other when they diverge.) Saying-intentions are constrained by belief and knowledge. *A* cannot (intend to) *say* that *p* by uttering some sentence *X* if he believes it is impossible for his audience *B* (or at least any rational, reasonably well-informed interpreter in *B*’s shoes) to construe him as saying that *p*.

Among the things that constrain the formation of *A*’s *saying-intentions* are *A*’s knowledge of the meanings of the words he is using and his (tacit) knowledge of the syntax of the language he is using. Thus *A* cannot (intend to) say that snow is white by uttering the sentence ‘grass is green’. And he cannot (intend to) say that John asked his

¹⁴ In my view it is vital to distinguish between (i) what *A* implied *by uttering X*, and (ii) what *A* implied *by saying what he said*, as this distinction gets to the heart of Grice’s distinction between conventional and conversational implicature, and to provide the framework within which to solve problems concerning the former (cf. Frege’s notion of tone or colouring), non-detachability, dictiveness and formality, and central vs. non-central speech acts.

¹⁵ If what *A* said is determined by, and only by, certain very specific *interpreter-directed intentions* *A* had in uttering *X*, then at least some contemporary talk of ‘contexts’ ‘fixing’ or ‘determining’ aspects of what *A* said—for example the references of indexical expressions—must involve some form of confusion. See below.

brother to shave him by uttering ‘John asked his brother to shave himself’. More generally, he cannot (intend to) say that p by uttering X if he believes it is impossible for his audience B (or at least any rational, reasonably well-informed interpreter in B ’s shoes) to construe him as intending to say that p .

8. *Speakers*

Saying and implying are things people do. (Similarly, communicating.) Following ordinary usage, *the speaker* is taken to be the understood subject, so to speak, of the verbs ‘say’ and ‘imply’ the verbs in talk about ‘what is said’ and ‘what is implied.’ (Similarly, with verbs such as ‘communicate’, ‘convey’, and ‘get across’.) We should be initially suspicious of talk about what *uses of sentences* say (imply, communicate, etc.) and talk about what *sentences-relative-to-contexts* say (imply, communicate, etc.), unless such talk is taken to be straightforwardly translatable into talk of things that speakers are doing. And we should deplore the unannounced slipping and sliding back and forth between different subjects of ‘say’ (‘imply’, ‘communicate’, etc.). At the same time, we should be open to the idea that new, *technical* uses of the verbs ‘say’ (‘imply’, ‘communicate’, etc.) may need to be defined, or at least developed, in the course of our inquiries, such uses earning their keep because of ineliminable theoretical work they do.

9. *Truth*

What A says and implies are the sorts of things that are true or false. (Perhaps A may say things that are neither true nor false, but it might prove useful to start out sceptical about this.) It does not follow that when A utters a sentence he says only one thing or that he implies only one thing. Nor does this talk of truth mean that in order to produce a theory of interpretation we shall have to construct a theory that recursively assigns *truth conditions* to sentences relativized to contexts of utterance (a *semantic* theory, in *one* sense of ‘semantic’), or construct a theory that assigns things in the world to linguistic expressions relativized to contexts of utterance (a theory of *reference*, in *one* sense of ‘reference’).

10. *Judgment*

Our intuitive judgments about what A meant, said, and implied, and judgments about whether what A said was true or false in specified situations constitute the primary data for a theory of interpretation, the data it is the theory’s business to explain. (Since *no-one* has intuitive judgments about *what is said by a sentence X* relative to a context C or about the *semantic content* of X relative to C (these being philosophers’ notions), several distinct mistakes would be involved in the claim that linguistic pragmatism aims to show that our intuitive judgments about what a speaker said may be ‘unreliable guides to semantic content.’ If talk of the ‘semantic content’ of a sentence X relative to a context C is just a snazzy way of talking about what the speaker said by uttering X on a particular

occasion—the occasion that *C* is being used to partially model—then of course we can accept its empirical significance. If it is not, then its empirical significance must be justified in some other way, from *within* the theory of interpretation by reference to some empirical rôle it is required to play in an explanation of what a speaker says and implies by uttering *X* on a given occasion, in much the same way that notions such as binding are motivated from within.

11. Reference

Saying typically involves *referring* and *saying of*. That is, saying something typically involves *referring to* something and *saying* something *of* it. Saying that London is pretty, for example, involves referring to London and saying of it that it is pretty. One way of doing this is to use ‘London’ to refer to London and ‘is pretty’ to say of it that it is pretty. Following ordinary usage, the *speaker* is taken as the understood subject, so to speak, of ‘refer to’ and ‘say of’. Initially, we should deplore the unannounced slipping and sliding, back and forth, between different subjects of ‘refer to’ and ‘say of’, and we should be suspicious of talk about what *uses of words* refer to and say of things, and of talk about what *words-relative-to-contexts* refer to and say of things—unless such talk is taken to be straightforwardly translatable into talk about things that speakers are doing. But we should be open to the idea that new, *technical* uses of ‘refer to’ and ‘say of’ may emerge in the course of our inquiries. Who or what *A* is referring to by uttering some expression *X* is determined by *A*’s *referential intentions* in uttering *X*. (Talk of ‘contexts’ ‘fixing’ or ‘determining’ the references of expressions—for example the references of indexical expressions—must involve some form of confusion.)

Nonetheless, a distinction between what *A* referred to and what *A intended* to refer to is not one obviously lacking a point. So, in the first instance we should separate (i) *who or what A intended to refer to by an expression X on a given occasion*, and (ii) *who or what a rational, reasonably well-informed interpreter in B’s shoes thinks A intended to refer to by X on that occasion*. In cases where (i) = (ii), we can talk freely about *what the speaker referred to*. (In cases where (i) ≠ (ii), we *could* argue about which of (i) or (ii) or some third thing has the ‘right’ to be called *the person of thing referred to*, but what would be the point? First, what third thing distinct from (i) and (ii) could be of any significance to a theory of interpretation? There is simply no rôle for a transcendent notion of what was referred to upon which (i) and (ii) converge when all goes well. Second, why is a choice between (i) and (ii) even needed in cases where (i) ≠ (ii)? Conceptually they are distinct, and they are both needed in a theory of interpretation. When all goes well, they coincide, and it’s just too bad they don’t always do so. Surely there is no philosophical payoff in bestowing the honorific ‘what was referred to’ on one rather than the other when they diverge.) Referential intentions are constrained by belief and knowledge. Assuming he is being co-operative, *A* cannot (intend to) refer to some particular person α by uttering some expression *X* on a given occasion if he believes it is impossible for his audience *B*

(or at least any rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as referring to α .¹⁶ Among the things that constrain the formation of *A*'s referential intentions are *A*'s knowledge of the meanings of the referring expressions he is using and his (tacit) knowledge of the syntax of the language he is using (which may bear on the matter of co-reference or binding where pronouns are concerned). Thus *A* cannot (intend to) refer to some particular individual α by *X* if he believes it is impossible for his audience *B* (or at least any rational, reasonably well-informed interpreter in *B*'s shoes) to construe him as (intending to) refer to α by *X*. Similar points can be made in connection with *A*'s *predicative* intentions.

12. *Aphonicity*

It is now time to get more theoretical. A distinction between PF ('Phonetic Form') and LF ('Logical Form') in something like Chomsky's sense, is almost certain to play a key rôle in a theory of interpretation, where a sentence's PF is (roughly) a representation that expresses its phonology, and its LF a representation that expresses *all* syntactic properties relevant to interpretation.¹⁷ This distinction brings with it the possibility of revealing in the LF of a sentence *X* syntactic objects that have no counterparts in *X*'s PF. Such 'aphonic' ('phonologically null') expressions are as much in need of interpretation when *X* is uttered as any other elements in *X*'s LF. (If a sentence's LF expresses *only* syntactic properties relevant to interpretation, as current theory dictates, this becomes a matter of definition.)

13. *Indexicality*

Identifying the LF of a sentence *X* does not constitute identifying what *A* says on a given occasion by uttering *X*. For one thing, *X*'s LF may contain an indexical expression like 'I' or 'he' or 'that'. So identifying *X*'s LF still leaves *B* some interpretive work to do, work that will involve accessing and integrating all sorts of information not carried or revealed by the LF itself.¹⁸

¹⁶ I mention *co-operation* because in certain circumstances one may seek to disguise one's intended referent from others, for example in cryptic poetry, diary entries, or dramatic irony. In such cases, there is either no intended audience distinct from oneself or some individual distinct from oneself with whom one is engaged but with whom one is being less than fully co-operative in the sense discussed earlier. The issues here are intimately connected to Grice's (1989) discussions of communicative intentions in the absence of an audience and to the issue of whether one's future self constitutes an audience.

¹⁷ See Chomsky (1986, 1995, 2000).

¹⁸ This does not mean that the pragmatist cannot, for certain expository or investigative purposes, operate *as if* a description of a sentence's LF gives us a description of what *A* said *relative to certain heuristic stipulations*. Formal 'contexts' or 'indices' are used in logic to anchor or co-anchor indexical elements in order to cancel or pair their effects across similar structures. Without commitment to the view that formal contexts play any sort of rôle in a theory of interpretation, the linguistic pragmatist may sometimes borrow this technique in order that a particular investigation may focus on particular non-indexical properties of LFs that are relevant to a theory of utterance interpretation.

14. Anchoring.

Idealization and abstraction from the details of particular speech situations or contexts are unavoidable if work is to proceed. To this extent, we may temporarily avail ourselves of the formal ‘indices’ or ‘contexts’ of indexical logics in order to anchor or co-anchor the interpretations of indexical or anaphoric expressions that are not of primary concern at a certain point of investigation. We should not take formal indices themselves particularly seriously, however. They are useful transitory tools, methodological or heuristic devices, not serious posits in a theory of utterance interpretation.

15. Mongrels

Since LFs may contain aphonics and may contain indexicals, we should be open to the possibility that they may contain aphonic indexicals. (At the same time, we should no more take seriously the idea that pairing a formal ‘context’ with an aphonic indexical in a sentence *X eo ipso* solves a genuine problem about the *interpretation* of an utterance of *X* than we should take seriously the idea that pairing a formal ‘context’ with a *phonic* indexical does so.) Aphonic indexicals are not the only possible mongrels. Since LFs may contain aphonics and may contain bindable variables, we should be open to the possibility that they may contain aphonic bindable variables. (Compare ‘everyone wants John to leave’ and ‘everyone wants to leave’. Perhaps the subject of ‘to leave’ in the latter is an aphonic variable bound by ‘everyone’. Certainly it is not an aphonic copy of ‘everyone’.) And why not aphonic, indexical, bindable variables?

16. Isomorphism

It is at least methodologically useful to say that identifying what a speaker said by uttering a sentence *X* on a given occasion involves entertaining a ‘sentence’ of Mentalese; or that it involves entertaining a structured proposition. (Perhaps entertaining a sentence of Mentalese ultimately amounts to entertaining a structured proposition because it involves entertaining a representation whose rôle in our mental life can be explained only in terms of it having a certain ‘content’ that a structured proposition supplies. Who knows?) With Sperber and Wilson, let us work for the moment with Mentalese. Since LFs are not full-blown representations of Mentalese (and so do not express propositions), but only ‘blueprints’, ‘schemas’, ‘skeletons’, or ‘templates’ for such, in advance of serious empirical investigation we cannot rule out the possibility of a *failure of isomorphism* in the mapping between the LF of a sentence *X* and a Mentalese representation the entertaining of which constitutes understanding what *A* said by uttering *X* on a particular occasion (or in the mapping between *X*’s LF and a structured proposition, the entertaining of which constitutes understanding what *A* said by uttering *X* on a particular occasion). That is, we cannot rule out atoms of the Mentalese representation (or atoms of the structured proposition) to which no element of *X*’s LF corresponds. (*A* might utter the sentence ‘The embassy is closed’ on a particular occasion

and *B* may be required to entertain the Mentalese sentence all too conveniently rendered as THE U.S. EMBASSY IN LONDON IS CLOSED in order to grasp what *A* said.¹⁹ *A* may utter, ‘the ham sandwich wants extra pickles’, and *B* may be required to entertain the Mentalese sentence THE MAN WHO JUST ORDERED A HAM SANDWICH WANTS EXTRA PICKLES. *A* may utter, ‘the hostages landed back on American soil today’ and *B* may be required to entertain the Mentalese sentence THE FORMER AMERICAN HOSTAGES AT THE U.S. EMBASSY IN TEHRAN LANDED BACK ON AMERICAN SOIL TODAY. It would seem that *A* may even utter less than a whole sentence—for example, ‘no thank you’ or ‘a cappuccino, please’—and thereby say something.)

We cannot rule out the possibility, however, that future work in syntax will indicate that we are closer to isomorphism than superficial appearances suggest, for all sorts of aphonics in LF may be revealed. Presumably, Mentalese representations will have to contain elements that function as (or at least do the work done by) bound variables, so we may well have to consider the possibility that interpreting a particular utterance of a sentence *X* may involve entertaining a Mentalese representation that contains a mental variable with no counterpart in *X*’s LF. On the other hand, syntactic evidence might be found for the existence of an aphonic variable in *X*’s LF. We cannot dogmatically assume that there *must* be isomorphism, and we should recoil from the unargued goal of attaining isomorphism by freely adding aphonics to LFs as if adorning some garish Christmas tree with a new light wherever it seems to dark. (We shouldn’t get hooked on aphonics.)

17. *Ellipsis*

Corresponding to the sentence-utterance distinction impressed upon us so forcefully by Austin, Grice, and Strawson, we must take seriously two important and distinct uses of the words ‘ellipsis’ and ‘elliptical.’ The first is a strict linguistic (or grammatical) notion found in talk of elliptical *sentences* in generative linguistics, a notion sometimes called *deletion* and which involves erasing elements in the generation of PF representations.²⁰ Linguistic ellipsis concerns the superficial incompleteness of *structures*, and as such is subject to a stringent condition on the constancy of form and interpretation that has been investigated by linguists under the rubric of *recoverability*. (*A* can use the sentence ‘I can tango but Mary can’t’ to say that he can tango but Mary can’t *tango*, but not to say that he can tango but Mary can’t *sing*. This is because it is elliptical for the complete sentence ‘I can tango but Mary can’t tango.’) The second notion of ellipsis is a pragmatic (or speech act) notion, found in talk of elliptical *utterances* of (elliptical on non-elliptical)

¹⁹ If you know English and your shift-key works, Mentalese is a cinch (Mentalease?). Structured-Propositionese is a little harder: you need good angled brackets.

²⁰ See, for example, Heim and Kratzer (1998), May (2001), Sag (1976), Williams (1977). May provides a particularly clear and user-friendly discussion of linguistic ellipsis. The grammatical notion clearly has its roots in talk of ellipsis and elision in roots in some traditional grammars.

sentences.²¹ Pragmatic ellipsis concerns the incompleteness of *interpretations*, and as such is governed only by general pragmatic principles governing interpretation. (A can use the sentence 'I'm going to a party at the embassy' to say that he's going to a party at the British embassy in Athens, for example, or to say that he's going to a party at the US embassy in London because there is no particular complete sentence that the sentence A uttered is an incomplete or elliptical version of.) We must accept that people often speak elliptically without much (if any) conscious effort and that hearers interpret elliptical utterances without much (if any) conscious effort. In such situations typically the speaker and hearer can both readily expand upon the sentence uttered in such a way that explains the ellipsis.

18. Competence

Three major components of a theory of interpretation are a *syntactic* theory, a *semantic theory* and a *pragmatic theory*. Certain preconceptions about the labels 'syntactic', 'semantic', and 'pragmatic' need to be put aside if the pragmatist position is to be understood, for these words are used in very precise ways. (Self-serving edicts from those who claim to have isolated the 'correct' way of making the semantics-pragmatics distinction or the 'correct' uses of the terms 'semantics' and 'pragmatics' should be ignored.) A *syntactic* theory for a person *A* who speaks a language *L* is an abstract description of *A*'s syntactic *competence* (in Chomsky's sense), *A*'s tacit knowledge of the syntax of *L*. This only becomes interesting when we are clear about what counts as a syntactic fact or phenomenon. Are binding and scope syntactic phenomena? A *semantic theory* is an abstract description of *A*'s semantic *competence*, his knowledge, tacit or otherwise, of the semantics of *L*. This only becomes interesting when we are clear about what counts as a semantic fact or phenomenon. Are binding and scope semantic phenomena? Binding shows that a sharp division between syntax and semantics is illusory (which is why Chomsky is prepared to use the label 'syntactic' in connection with much of what many philosophers and linguists label 'semantic'). Drawing a sharp line between semantics and pragmatics is straightforward. A *pragmatic* theory transcends individual speakers and particular languages. It is an abstract description of the mechanisms that make it possible for interpreters to identify what a speaker means by uttering a sentence (or sentence fragment) *X* on a given occasion given (at most) what a semantic theory has to say about *X*. As such, a pragmatic theory is a description of an intentional and richly inferential system, our common pragmatic *competence*. There is no assumption here, nor is there any antecedent reason to suspect that this semantics-pragmatics distinction just drawn will be co-ordinate with the saying-implying distinction.

²¹ See, for example, Quine (1940), Sellars (1954), Brinton (1977), Bach (1981), Salmon (1982), and Neale (1990).

19. *Semantics*

Words and the ways in which they can be combined have properties that enter into an explanation of why speakers use the particular combinations they do, and why hearers interpret speakers using these combinations in the ways they do. The two most obvious properties are the *meanings* of words and *syntax*. Qua description of semantic competence, a semantic theory for a language will explain how the syntactic structure of a sentence (or sentence fragment) *X* and the meanings of the individual words in *X* conspire to constrain what speakers can *say* using *X*. Flushing out the modal: a semantic theory for a language *L* will provide, for each sentence *X* of *L*, a *blueprint* for (a *template*, a *schematic* or *skeletal* representation of) what someone will be taken to be saying when using *X* to say something. The blueprint associated with *X* is its *semantics*, and the set of such blueprints, one for every sentence of a language *L*, is the *semantics* for *L*. (The study of these blueprints is also called *semantics*. The study of the rôle of word meanings is called *lexical semantics*; the study of the rôle of syntax is called *compositional semantics*.) Semantic competence comprises at least (i) knowledge of the meanings of individual words and (ii) knowledge of syntax (syntactic competence). It is a matter of debate whether it involves more. On the one hand, if *A* claims not to understand a sentence *X*, then it would seem that either the meaning of some word in *X* eludes him or else some aspect of *X*'s structure (ultimately *X*'s LF) does. On the other, having a model aeroplane kit, a foolproof set of instructions, excellent glue, plenty of space, good lighting, and the fingers of a heart surgeon is not the same thing as having the model aeroplane (that's why they write 'kit' on the box). Settling this debate involves settling (among other things) what syntactic competence amounts to and how bad the model aeroplane analogy is. (Certainly knowledge (in the requisite sense) of syntax does not amount to the propositional representation of a set of syntactic rules).

20. *Pragmatics*

Whereas each language (perhaps even each idiolect) has its own syntax and its own semantics—which is not to say that vital syntactic and (hence) semantic properties are not shared across languages as a result of our common biological endowment—there is, so to speak, only *one* pragmatics. Qua description of our shared pragmatic competence, a pragmatic theory will explain how interpreters identify what a speaker means by uttering a sentence (or sentence fragment) *X* on a given occasion given (at most) what a semantic theory has to say about *X*.

The semantics-pragmatics distinction, thus construed, is not co-ordinate with the saying-implying distinction. What *A* means by uttering *X* on a given occasion comprises what *A* *said* and what *A* *implied*. So a pragmatic theory will explain how interpreters identify what *A* said and implied by uttering *X* on that occasion given (at most) what a semantic theory has to say about *X*. If a *pragmatic* theory explained only how interpreters identify what *A* *implied* given (at most) what the speaker *said* as 'input', a gaping hole in

our taxonomy of *theories* would appear. A *semantic* theory specifies the constraints that word meanings and syntax place on what *A* can say by uttering *X*, a blueprint for *X*. What would we call a theory that explains how interpreters identify what *A* said on that occasion? Not a *semantic* theory, for that specifies only a blueprint for what *A* said, i.e. the *sort of thing* he said. Clearly, a pragmatic theory has *two* rôles in a theory of interpretation. Even if an utterance of a sentence *X* always wore on its sleeve an unambiguous representation of its syntactic structure with no ambiguous elements, a semantic theory could still fail to identify fully what *A* said by uttering *X* on a particular occasion. For one thing, *X* may contain /he/, /this/, /here/ or /John/, in which case the interpreter needs to identify who or what *A* is referring to.²² Since these words are not

²² My wording should make it clear that I am putting aside, for now, the fact, stressed by many pragmatists, that a pragmatic theory may have to be invoked in order to identify what *sentence* *A* uttered because of ambiguities at PF. Among the things a hearer or reader has to do in order to identify what *A* is saying on a given occasion, is identify which *words* *A* is using. /Bank/ is the superficial form of either a single, ambiguous word of English or else of two distinct unambiguous words, and I do not want one's position on this matter to impinge upon one's understanding of IT. (If /bank/ is the superficial form of a single, ambiguous word, then identifying what *A* is saying when he utters, 'I'm going to the bank' involves identifying which *meaning* *A* has in mind for /bank/; if /bank/ is the superficial form of two distinct, unambiguous words then identifying what *A* is saying when he utters, /I'm going to the bank/ involves identifying which of the two *words* *A* is using. The latter view seems more useful in theorizing about language. 'Word' and 'sentence' are quasi-technical terms, there are no ambiguous words or sentences, and (following Chomsky) every sentence comprises a superficial form PF and an underlying form LF, the former being what is relevant to speech perception, the latter what is relevant to speech comprehension. When I wish to talk explicitly about an expression's PF or about the sound common to two expressions, and when I wish to avoid commitment one way or the other as to whether I am talking about one expression or two, I shall borrow the old slash notation of phonology (but with standard orthography rather than a phonological representation enclosed, as in /bank/) to individuate coarsely in terms of phonological properties. (On one use, then, /bank/ is what Perry (1998) calls a *vocable*.) Thus I sometimes use /he/, /him/, and /his/ in my discussions of pronouns because it is arguable that each corresponds to two distinct words in English, one that is bound and another that is not. Similarly, I sometimes use /the/ so as not to prejudge the issue on the matter of a purported ambiguity in the definite article(s). Of course if /he/ and /the/ really are ambiguous, the ambiguity in question will have to be more systematic than the sort found with /pen/ or /bank/. It is easy enough to cause trouble for my use of the slash notation. Almost certainly we want to distinguish the phonologically identical but orthographically distinct 'so', 'sew' and 'sow^a' (as in *seeds*), distinguish the orthographically identical but phonologically distinct 'sow^a' (as in *seeds*) and 'sow^b' (as in *pig*), and distinguish the phonologically and orthographically identical 'pen^a' (as in *writing instrument*) and 'pen^b' (as in *enclosure*); but probably we shall not need not bother distinguishing the (merely) orthographically distinct 'judgment' and 'judgement' or the (merely) phonologically distinct 'côntroversy' and 'contrôversy.' (Actually, I'm not so sure I should have said that: perhaps there are worries here not entirely unconnected to those Kripke (1979) brings up in connection with 'Paderewski'.) So when I want to individuate coarsely in terms of (roughly) phonology, I use the slash notation. Thus /pen/ and /sew/ (and, unfortunately, /sow^a/ and /sow^b/). If a single word can have two distinct orthographies ('judgment' and 'judgement') and a single word can have two distinct phonologies ('côntroversy' and 'contrôversy'), should we explore the idea that a single word can have two distinct orthographies *and* two distinct phonologies? Or should we say that something is a single word only if it is grounded in a single phonology or a single orthography? Or should some intermediate position be explored that invokes etymology or the relative similarity of distinct orthographies and phonologies, a position according to which /doctor/ and /physician/ would be too far apart to qualify, orthographically, phonologically, and etymologically? (Notice how the notation just exploded.) Many Greek villages or islands still have two names, and the reason we talk this way is because the names seem too far apart to count as a single name (e.g. 'Thíra' and 'Santoríni'). But what about 'Aperáthou' and 'Apeíranthos' (a village on the island of Náxos)?

ambiguous in the way /pen/ or /bank/ are said to be, and since each is not merely the unambiguous surface form of a context-insensitive definite description, something other than a semantic theory must be invoked.²³

The slack is taken up by a *pragmatic* theory: identifying what *A* said involves the exercise of cognitive capacities that integrate the semantic information carried by the sentence uttered and all sorts of ‘pragmatic’ or ‘contextual’ information including, but not limited to, information obtained by perception from the physical environment, information about the interpretation of prior utterances in the conversation (if any), information in memory, and information about how people typically behave, particularly in communicative exchanges. That is, identifying what *A* said involves processing not only the semantic information encoded in a sentence’s form, but accessing and processing information that must be picked up by listening, watching, remembering, hypothesizing and inferring, essentially the capacities exercised in identifying what *A* implied. To this extent, then, identifying what is said is a *pragmatic* as well as a semantic matter. It involves *pragmatic inference* as well as *linguistic decoding*. Identifying what a speaker implied is something explained by a pragmatic theory, typically taking into account what *A* said; but identifying what the speaker said is also something explained by a pragmatic theory, taking into account (in a big way, to be sure) a sentence’s blueprint, which is explained by a semantic theory. Underpinning the difference between identifying what *A* said and what he implied is a distinction in the type of typical *input*: to identify what *A* said on given occasion by uttering *X* the pragmatic system typically takes as its primary input the output of the semantic system (the semantic information encoded in *X*); to identify what a speaker implied it takes as its primary input what the speaker said. This leaves many questions open: (1) To what extent is pragmatic processing deductive? (2) To what extent does it take place unconsciously? (3) What sorts of things affect its speed? (4) To what extent is it task-specific or modular?

21. Underdetermination.

It is now possible to bring together several points. The rôle of a pragmatic theory in identifying what *A* said by uttering *X* on a given occasion is not restricted to identifying who or what *A* is referring to by any referential expressions in *X*. Saying involves referring and predicating; and just as identifying what *A* is up to with any referential devices in *X* involves more than consulting a mental lexicon, so does identifying what he is up to with any predicative devices in *X*. It may, for example, require the ‘saturation’ of an ‘implicit argument’ as in ‘It’s raining’ or ‘I’ve finished’. (Some implicit arguments may be mandated by syntax as well as semantics.) Or it may require ‘enriching’ a predicate in some way that is reasonably obvious and presumably acceptable to *A*. The

²³ Three points must be separated here: (i) failure of lexical ambiguity, (ii) context sensitivity, and (iii) rigidity.

sentence ‘every woman has a job’ might be used to say that every woman in Flint has a job, or that every woman in Woodside has a job, etc.

It will pay to separate and rejoin two points here, one epistemological, the other metaphysical, both intimately connected to the points made earlier about *Asymmetry* and *Reciprocity*. The epistemological point concerns *insufficiency*, the metaphysical point *underdetermination*. From the hearer’s perspective, we can talk first about the fact that knowledge of the syntax of *X* and knowledge of the meanings of all the words in *X* do not suffice for identifying what *A* is saying by uttering *X* (even where the superficial form evinces no lexical or structural ambiguity). At most they yield a blueprint. Now we can bring in the *A* himself. What *A* says is wholly determined by certain specific intentions he had in speaking, intentions massively constrained by his knowledge of syntax and word meaning (and a whole lot more). *A* tacitly knows that *B*’s knowledge of word meaning and syntax will not suffice to furnish *B* with a complete account of what he has said. We can now introduce some theoretical shorthand to obviate the need to keep talking about speakers’ and hearers’ knowledge, tacit or otherwise. Let us say that syntax and word meaning together *underdetermine* what is said (all the time remembering this is shorthand). But we are not yet where we need to be. The first thesis we need is this:

- (IT) *The Insufficiency Thesis*: Identifying what a speaker or writer, *A*, is *saying* by uttering an unambiguous, declarative sentence *X* on given occasion goes well beyond recovering *X*’s underlying syntax, knowing the meanings of all of the words in *X*, and identifying who or what *A* is referring to by any referential expressions in *X*.

This goes beyond the insufficiency just mentioned because it entails that even when *B* has identified who *A* is referring to by any referential expressions in *X* (/John/, /he/, /here/, /that/, /I/, /you/, and so on), *B* still doesn’t have everything he needs to identify what *A* said. Since *A* tacitly knows that *B*’s knowledge of word meaning, knowledge of syntax *and* knowledge of who or what *A* is referring to by any referring expressions in *X* will not suffice to furnish *B* with a complete account of what he has said, we can formulate the shorthand we really want:

- (UT) *The Underdetermination Thesis*: What *A* *says* by uttering an unambiguous, declarative sentence *X* on given occasion is underdetermined by *X*’s syntax, the meanings the words (and any other morphemes) in *X* (and the meanings, if any, of prosodic features of *X*), and the assignment of references to any referring expressions in *X*.²⁴

²⁴ The Underdetermination Thesis is regularly stressed by linguistic pragmatists, and some view it as a cornerstone of the general outlook. This use of ‘underdetermination’ is found in the work of Sperber and Wilson (1986, 1995) and borrowed by many of those they have influenced, including Bezuidenhout, Blakemore, Carston, Papafragou, Récanati, Rouchota, and me (*Descriptions*, 1990: 114, n. 46). In the language of Perry (1986, 1993, 1998, 2001), talk of underdetermination is roughly equivalent to talk of constituents of propositions expressed that are ‘unarticulated’, i.e. constituents corresponding to no constituents of the sentence uttered. (As Perry sometimes puts it, we don’t always articulate things when it’s

As far as constructing and evaluating a theory of interpretation are concerned, we must make sure we separate talk of (a) the *interpretive target* (stipulated in advance by the theorist in particular cases on the basis of intuitive judgment, modulo reflective equilibrium with the best theory up to that point), and talk of (b) the *knowledge and mechanisms* in play (under investigation and hypothesized by the theorist). And in talk of knowledge and mechanisms we must be careful separate (i) the role of syntax; (ii) the role of word meaning; and (iii) the actual pragmatic mechanisms.

An interpretive target is a characterization of what our intuitive judgments reveal the speaker to have said and which an adequate theory of interpretation should deliver. It is common to specify interpretive targets using more language: By uttering ‘she has a job’ A is saying that Margaret Thatcher has a job; by uttering ‘every woman has a job’, A is saying that every woman living in Woodside has a job. There is, of course, something a bit funny about this, for surely we can now ask for a characterization of *what the theorist said* when he uttered the sentence ‘By uttering ‘every woman has a job’, A is saying that every woman living in Woodside has a job’. Nonetheless, this is what we do, and when pressed often we wheel out some set theory: By uttering ‘every woman has a job’, A is saying something that is true iff at the time of utterance t , the intersection of the set of things that are women at t and the set of things living in Woodside at t is a subset of the set of things that have jobs at t . While this may provide us with the conditions under which what A said is true, it falls short of specifying *what* A said for familiar reasons. (First, A is sure to deny it. Second, surely A would have said something different had he uttered, ‘every woman has a job and $19^2 = 361$ ’). So when pressed again we wheel out something like a *situation* or a *structured proposition*.

Where *knowledge* of syntax, word meaning, and the theory of blueprints are concerned, there is much work to be done by philosophers of language and linguists together. As far as *mechanisms* are concerned, we are squarely in the realm of psychology, and some philosophers of language and linguists may well opt out. The psychological part of the overall project may certainly be informed by philosophical reflections, such as Grice’s, on the nature of rational, purposive behaviour, but ultimately it is a wholly empirical

clear from context what they are.) I don’t know why underdetermination is rejected so vehemently by some philosophers of language, but I have a suspicion two related factors may be implicated, both involving fear and philosophical temperament. The first is a simple unwillingness to concede apparently hard-earned territory, a reluctance to accept that some of the traditional problems involved in so-called ‘compositional semantics’ are actually the products of specious questions in the philosophy of *language* with genuine and important counterparts in the philosophy of *mind*, mostly about *inference* and *the composition of thought* (a reluctance to accept that as far as natural language is concerned, trying build pure *content* is as futile as trying to build pure *character*). The reluctance, it seems to me, amounts to little more than obstinacy or fear of a philosophical pink-slip. The second fear might be viewed as an extension of the first: the fear that if there is not at least one component (what is said) of what a speaker means that can be nailed down precisely and completely without taking into account too many ‘pragmatic considerations’, then systematic semantics as typically understood is doomed, and with it any chance of producing a serious theory of language. This fear seems to me entirely unwarranted. Natural language semantics may not be quite as straightforward or far reaching as many have thought, but there is plenty of systematic semantics for all of us (and more) to do for longer than we will ever have to do it.

enterprise the aim of which is to identify the cognitive mechanisms whereby the hearer effects the relevant identifications on a given occasion.

22. Indeterminacy

What a speaker says and implies may be indeterminate in at least the following sense: in any vocabulary in which what someone says or implies can be usefully specified, there will be alternative and strictly distinct specifications between which no principled choice can be made. We should not worry that indeterminacy of this sort presents problems for particular semantic proposals. (For example, we should not regard traditional accounts of descriptions as damaged in any way by the indeterminacy attaching to intuitive ‘completions’ of those that are said to be ‘incomplete’ relative to a particular occasion of utterance.)

23. Convergence

A univocal saying-implying distinction is *empirical*, *ordinary*, and *beneficial (practical)*. The distinction is empirical insofar as it assumes that, typically, representations corresponding to what A said and implied are the outputs of cognitive mechanisms involved in the interpretation process. It corresponds to something entrenched in ordinary talk (despite the fact that we may disagree in particular cases). And it underpins the very idea of codifying principles meant to regulate societies and the behaviour of their members (e.g. laws, contracts, and commitments) by virtue of being a distinction one side of which (saying) is about as *objective* as anything can be, a fact itself guaranteed by the empirical and ordinary nature of the distinction. (To say this is not to say there cannot be disputes about what was said, changes of opinion after discussion of problematic cases, or specialists (or at least professionals) in societies to whom tough cases are referred when the issue needs forcing. Rather it is to say that there is enough overlap in judgment to render regulation, commitment, and so on meaningful notions.

24. Formalism.

Advances in our thinking about language have come out of developments in logic and formal philosophy, particularly by way of the construction and use of various types of broadly mathematical theories, systems or analyses—the predicate calculus, model theory, modal logics, set theory, recursion theory, and generalized quantifier theory to name the most obvious. But it does not follow that associating utterances with models, possible worlds, structured propositions, indices, functions or even favoured formalisms *ipso facto* constitutes part of a theory of utterance interpretation. Rigorous formalism almost certainly has its place; but a favoured mathematical idea and an associated formalism must not so dominate our inquiry that the questions motivating it in the first place become obscured or transmogrified to the point of demanding purely technical answers. We should strive to use our formalisms judiciously, sparingly, only where they

were needed to effect a useful idealization or abstraction, forestall a potential ambiguity, capture a generalization, facilitate a transition, or usefully abbreviate something. Appeals to, say, higher-order functions or set-theoretic entities (and the use of corresponding notations) are ultimately dispensable in the theory of interpretation, sets and functions being no more than occasional, transitory tools of no intrinsic interest outside mathematics proper and the philosophy thereof. Such entities are not the objects of semantic investigation themselves, and it should not be a goal of any branch of philosophy to drag them into investigations whenever the opportunity presents itself.

3. The Irrelevance of Contexts

Following Grice, let us talk about what a speaker *A* means on a given occasion by uttering some sentence *X*, factoring this into what *A* says and what he (merely) implies. As Grice notes, *identifying* what the speaker is saying is not simply a matter of identifying *X* and recovering its linguistic meaning (blueprint), if only because of the existence of pronouns. Unlike some of his critics, Grice is careful not to run together epistemological and metaphysical points here, despite their evident interconnectedness. The important metaphysical question is: what *determines* what a speaker said on a given occasion? And the Gricean answer is: certain specific *intentions* he had in producing his utterance, intentions that are severely constrained by his beliefs about the meanings of the words he uses, about his audience, about the context, about the topic of conversation, and probably a whole lot more.²⁵ The important epistemological question is: what knowledge or information does a hearer use in *identifying* what the speaker said? And the Gricean answer is: knowledge of the linguistic meaning of the sentence uttered, pragmatic knowledge about the way rational, co-operative beings operate, knowledge about the speaker, knowledge of context, and just about anything else he can get his hands on. Let us take a concrete example. If I say something by uttering a sentence *X* that contains the personal pronoun ‘he’ and the demonstrative pronoun ‘this’, then (in the simplest case, at any rate) my referential intentions determine who I mean by ‘he’ and what I meant by ‘this’. (Similarly, my lexical intentions determine what I meant by /bank/ if *X* contains one of the words we write that way.) Your job as hearer is to identify what I meant by uttering *X*, and very likely you will not succeed unless you identify who I meant by ‘he’ and what I meant by ‘this’. (Similarly, what I meant by /bank/.)

Sperber and Wilson (1986, 1995) point out that pronouns are just the tip of a pragmatic iceberg for traditional, titanic accounts of what is said: quite generally, what a speaker says is underdetermined by the meaning of the sentence uttered, even relative to reference assignment. Now one can perfectly consistently accept Sperber and Wilson’s under-

²⁵ See Grice (1957, 1969, 1989), Neale (1992, forthcoming *b*). Very roughly, one cannot intend what one believes to be impossible.

determination thesis without rushing to embrace the details of their Relevance Theory. For that theory is meant to provide an account of the mechanics of utterance interpretation, of the richly inferential processes providing the basis of an empirically satisfying account of how interpreters (i) identify which sentence a speaker has produced on a given occasion in cases where identification of phonological form fails to yield a unique result; (ii) identify what the speaker *said* by uttering *X* on a given occasion in cases where identification of the meaning of *X* falls short; and (iii) identify what the speaker implied by uttering *X* on that occasion. Relevance Theory goes well beyond accepting that the meaning of a sentence *X* may underdetermine what a speaker says by uttering it on a given occasion and well beyond the vague Gricean idea that quite general principles governing the way we reason about the behaviour of others lie at the heart of an explanation of how we communicate.

In principle just about any information could be relevant or brought to bear on interpretation, and one of the main problems involved in constructing a pragmatic theory is explaining how the information that actually *is* brought to bear is delimited.²⁶ The second problem concerns *how* it is brought to bear. Linguistic pragmatism finds little sense to the idea that two quite distinct sets of information-gathering and inferential mechanisms are at work when a hearer tries to identify what a speaker means, one set that works on sentence meanings and yields what the speaker said, and another set that works on what the speaker said and yields what he meant but did not say (i.e. what he implies).

It is odd that some philosophers write as if (or even claim that) two quite distinct sets of cognitive mechanisms must be at work. I detect two related ideas lurking behind this assumption: (i) something to do with ‘simplicity’ (or ‘degree of difficulty’) or ‘systematicity’ (or ‘range of possibilities’); (ii) the influence of ‘indexical logics.’

(i) Whilst it is true that identifying who or what a speaker intends to be referring to on a given occasion by some particular referring expression *X* is constrained by the linguistic conventions governing the use of *X*, this does not necessarily make matters particularly straightforward or reduce the number of hypotheses that could, in principle, be investigated and assessed. Consider the interpretation of an utterance of the pronoun ‘it’. As Sperber and Wilson (1986: 187) note, all that the linguistic conventions governing the pronoun ‘it’ insist upon, in any context, is that the object should be non-human, giving every hearer in every context an indefinitely large choice of possible referents. And surely the same general considerations about, say, relevance, truthfulness, informativeness or whatever, that are invoked in identifying what a speaker is implying on a given occasion will be invoked in identifying who or what a speaker is referring to by ‘it’ on a given occasion.

²⁶ There is a division here between the optimists and the pessimists. Unlike Sperber and Wilson, Blakemore (1987), Carston (2002), and other Relevance theorists, Fodor (1983, 2001), Chomsky (2000), and Davidson (1986) suspect that producing an overarching theory of interpretation will require nothing short of a complete theory of mind.

The point can also be made in connection with an incomplete description like ‘the table’ or some other incomplete quantifier expression like ‘every man’ or ‘no one’. It is sometimes said that identifying what *A* says by uttering a sentence containing such an expression involves either (a) coming up with an appropriate domain of quantification *implicit* in the utterance (the ‘implicit’ approach), or (b) coming up with an appropriately ‘richer’ nominal *A* could have used to make his meaning more *explicit* (the ‘explicit’ approach). Whatever the final merits of such suggestions, one thing is quite clear: the same general considerations about, say, relevance, truthfulness, informativeness or whatever, that are invoked in identifying what a speaker is implying on a given occasion will have to be invoked in identifying an appropriate completing domain or an appropriate completing expression.²⁷

(ii) Many philosophers write as if (or even argue that) understanding what a speaker *A* said on a given occasion by uttering a sentence *X* with its conventional meaning is a matter determined by the meaning of that sentence and a ‘context’, in a sense of this frequently invoked word that is meant to make it more than simply a label for whatever it is that ‘bridges the gap’ between the meaning of *X* and what *A* said by uttering *X* on that occasion. For example, it is frequently claimed that all one needs to bridge the gap is some sort of formal object, an ‘index’ or ‘context’ in the form of an ordered *n*-tuple that secures the references of a few annoying ‘indexical’ pronouns (‘I’, ‘you’, and ‘he’, for example) and one or two other ‘indexical’ words that have a somewhat pronominal nature (‘here’ and ‘now’, for example).²⁸

This idea is rightly spurned in Sperber and Wilson’s *Relevance*, spurned in Evans’s *The Varieties of Reference* and *Collected Papers*, and spurned in Chomsky’s *Reflections on Language*—indeed spurned in *every* book or article that Chomsky has ever written in which the interpretation of pronouns is discussed. For there is an implicit recognition in these works, and in many others that bear their influence, that whilst formal contexts may have a useful *methodological* rôle from time to time, they are strictly irrelevant to a proper theory of utterance interpretation.

For various semantic and syntactic purposes, it is often desirable—if not mandatory—to abstract or idealize away from facts to do with particular speech situations—‘pragmatic’ or ‘contextual’ factors, as they are sometimes called—in order to get on with a particular piece of work. And as long as caution is exercised there is no harm in this.

²⁷ Postulating an aphonic indexical, *domain* variable in underlying syntax makes no more of a contribution to explaining *how* hearers interpret utterances than does postulating an aphonic indexical *assertion* variable in underlying syntax (or an aphonic indexical *irony* variable). That is, the interpretive task facing the hearer is made no easier by the existence of an aphonic contextual variable, even when, as in the case of the supposed assertion or irony variables there are just two possible values to choose from. And of course, giving phonetic form to such operators—‘asserting-or-not, it’s Tuesday’ or ‘being-ironical-or-not, it’s a lovely day’—doesn’t help the hearer either.

²⁸ The words ‘indexical’ is itself part of the problem, suggesting as it does that interpreting such devices involves merely looking something up in an ‘index’. People can be more influenced by labels than they sometimes realise.

For example, *with certain restricted purposes in mind*—and without any sort of absurd commitment to the idea that such entities play a role in utterance interpretation—formal ‘indices’ can be introduced to serve as ‘contexts’ with which sentences can be paired in order to ‘anchor’ or ‘co-anchor’ the interpretations of certain indexical expressions. The usual idea is to treat such expressions as free variables and treat indices as sequences or functions that assign them values. Famously, this idea has been used to capture model-theoretically the validity of inferences whose premises and conclusions are stated using indexical sentences:²⁹

A: If the next left is not Bank Street, that man gave you the wrong directions.

B: It’s not Bank Street; so he gave me the wrong directions.³⁰

It is paramount in such work to keep things tightly under control in the following sense: the logician wants a mechanism that can (a) scan a set of sentences for occurrences of symbols on some pre-existing list of devices that do not carry their values with them, then (b) use an index to assign a value to each occurrence of such a symbol. If this goes well, logical deductions can proceed (assuming a semantics for items of a pre-selected ‘logical’ vocabulary of course). If there is still slippage after the index has made its assignments, on standard assumptions there is only one solution: posit further indexical symbols in the sentences involved, symbols which are invisible in surface syntax yet revealed by an analysis of their ‘logical forms’, then try again.³¹

In the philosophy of language, indices have a methodological rôle for they can be used to anchor or co-anchor indexical and anaphoric expressions and so allow work to proceed more easily on *other* expressions and on what people say (and imply for that matter) by

²⁹ See (e.g.) Kaplan (1989), Lewis (1975), Montague (1974). A lot in this area turns on one’s conception of logic, and my wording evinces a particular stance, though not one I want to insist on: logical relations hold among what is expressed by sentences not among sentences themselves. (Various issues about the notion of formal validity and inference rule must be faced (but usually are not) by people who hold this view of logic.) The point I am making in the text is not dependent upon this stance. Cf. discussions of the difference between the logical form of a proposition and the logical form of a particular sentence used to express that proposition.

³⁰ A related point might be made in connection with anaphora: (i) every man loves his mother; (ii) John is a man; therefore (iii) John loves his mother. For discussion, see Neale (forthcoming *a*)

³¹ It is, perhaps, tacit recognition of this fact that has led some philosophers to conclude that there is no hope of producing a theory of *utterance interpretation* without positing all sorts of phonetically null, indexical elements in the underlying syntax of natural language sentences. We may use anything we like to throw light on the syntax of natural language, but we must never lose sight of the fact that discerning the syntactic structures of our sentences is an empirical exercise. Certainly the idea of aphonic elements in syntax is not objectionable in itself. On the assumption that syntax relates sound and meaning, we must certainly allow for the possibility of elements that have sound but no meaning (‘it’ in ‘it’s raining’?), or meaning but no sound (the understood subject of ‘leave’ in ‘Tom wants to leave’?). And there can be little doubt today that great advances in our understanding of syntax have been made by those such as Chomsky who have not shied away from the idea of aphonic items in syntax and argued for their existence and explanatory value. But we cannot simply *assume* that whenever we encounter some feature of what is said that does not appear to correspond to any element or feature of the sentence uttered it means there is some element in underlying syntax waiting to be exposed.

uttering them on given occasions.³² However, there is an idea that has emerged from work on indexical logics for which we can have little sympathy. This is the idea that sentence meanings and contexts can be paired to provide something of *empirical* significance: what a *sentence X says relative to a context C*.³³ We must not lose sight of certain facts. First, as far as utterance interpretation is concerned, such ‘contexts’ are strictly irrelevant. Utterances do not come with such devices attached that anchor or co-anchor indexical, demonstrative, or anaphoric pronouns. The hearer has plenty of *pragmatic* work to do, much of it rightly called inferential, albeit inferential in a way that is steered by the meanings of individual words. Evans (1982, 1985) summarizes the situation well:

All that the conventions governing the referring expression ‘he’ insist upon, in any given context, is that the object referred to should be male. (1982: 312) There is no linguistic rule which determines that a ‘he’ or a ‘that man’ refers to *x* rather than *y* in the vicinity, or that it refers to someone who has just left rather than someone who has been recently mentioned (1985: 230–1). ‘This’ and ‘that’ are even less specific, contributing merely the vaguest suggestion of a contrast between nearer and further (in some generalised sense). . . [Footnote: Often the *predicate* does more to narrow down the range of possible interpretations of the referring expression than does the referring expression itself . . .] (1982: 312). Let me take another example: the expression ‘you’: If a speaker addresses a remark to someone, saying, ‘You are a crook’, it is surely clear that an identification is called for on the part of the audience: in order to understand the remark, it is not enough to know that there is one, and only one, person whom the speaker is addressing, and that the speaker is saying of that person that he is a crook . . . a quite specific *kind* of identification is called for; the person addressed has not understood the remark unless he realizes that the speaker is saying that *he* is a crook. . . understanding the remark requires the hearer to know *of* an individual that he is being addressed. (1982: 314).

³² See *Descriptions*, Ch 3, for example.

³³ I am putting aside here some very real concerns about talk of *sentences* saying things relative to contexts. I am sceptical about the value or relevance of the use of the verb ‘say’ assumed in this way of talking to the project of constructing a theory of utterance interpretation, unless it is understood as a stylistic variant of talk of *speakers* saying things by uttering sentences on given occasions. Judgments about what a *speaker* said, and about whether what he said was true or false in specified situations, constitute the primary data for a theory of interpretation, the data it is the business of such a theory to explain. What a speaker *says* and what he *implies* (e.g. conversationally implicates) on a given occasion are the things that together constitute what the speaker *means*, and a theory of interpretation is meant to explain the role of linguistic meaning and inference in the hearer’s identification of what the speaker meant. No-one has intuitions about *what is said by a sentence X relative to a context C* or about the truth or falsity of *X* relative to *C* unless this is just a formal way of talking about what the speaker said by uttering *X* on a particular occasion—the occasion that *C* is being used to partially model. If such talk is straightforwardly transposable into talk about what the speaker said then we can accept its empirical significance. If it is not so transposable, then its empirical significance must be justified in some other way, from *within* the theory of interpretation by reference to some empirical role it is required to play in an explanation of what a speaker says and implies by uttering *X* on a given occasion, in much the same way that notions such as LF (‘Logical Form’), scope, and binding are motivated from within. If some such motivation is forthcoming, we should be only too happy to listen. I suspect it will not be forthcoming because the notion of what a *sentence* says relative to a context is going to be too thin and overly-detached from speakers’ communicative intentions to carry any empirical weight. Nonetheless, I adopt a wait-and-see approach. We are involved in an empirical enterprise after all.

Nothing about the meaning of the word ‘you’ tells you that you are being addressed.³⁴

We need, then, to distinguish two ideas, one sensible, the other silly. The silly idea is that utterances come with pre-packaged ‘contexts’ that provide values for indexical expressions. The sensible idea is what I call *methodological anchoring* (*anchoring* for short). For various pragmatic, semantic and syntactic purposes, it is often helpful, perhaps even mandatory, for a theorist to abstract from certain ‘contextual effects’ or ‘pragmatic factors’ in order to get on with a piece of work, and so it is sometimes useful to use an ‘index’ as a way of anchoring the interpretations of indexical expressions that are not, at that moment, the objects of primary concern, *even though the theorist knows the interpretation of these indexicals is not as straightforward as invoking an index might suggest*. If one is working on definite descriptions, for example, one might want to prescind, as much as possible, from the effects of, say, indexical pronouns occurring inside nominals; and if one is working on ‘and’, for example, one might want to prescind, as much as possible, from the effects of, say, indexical pronouns occurring inside conjuncts:³⁵

- (1) *He* drove home and *he* drank those six beers *you* bought *him*
- (2) *He* drank those six beers *you* bought *him* and *he* drove home.

To this end, we might use an index to anchor or co-anchor these expressions, to keep their special features and the complexities they introduce out of the picture as it were.³⁶

A certain amount of care is needed in the use of the word ‘semantic’ when indices are used to anchor (or co-anchor) indexical expressions. To the extent that we are investigating the conventions governing a word whose rôle cannot be set out clearly without taking into account the conventions governing other expression(s) with which it combines to form larger expressions, we may find it convenient to talk about the (derived) conventions governing the larger phrases with respect to a particular index. For example, if the semantics of ‘the’, is being investigated, it may be useful, even mandatory, to anchor indexicals so that *other* contextual effects may be monitored. And

³⁴ As soon as we introduce anaphoric pronouns—those that are linked in some interpretive fashion to other expressions (their ‘antecedents’)—matters become more complicated. The reflexive ‘himself’ *must be* interpreted via an antecedent; the non-reflexives ‘he’, ‘him’, and ‘his’ *can be* so interpreted (under certain conditions). Very roughly, reflexives cannot be ‘too far away’ from their antecedents, and non-reflexives cannot be ‘too close’ to them, putting the two in virtual complementary distribution as far as interpretive dependence is concerned, as suggested by Chomsky’s (1981, 1986) Binding Theory. For discussion within the present framework, see Neale (forthcoming *a*).

³⁵ The following examples are due to Deirdre Wilson.

³⁶ Carston (1988, 1993, 2002) implicitly anchors in her examinations of ‘pragmatic enrichments’ in connection with utterances of conjunctions (indeed, it is what she implicitly does throughout). Similarly, Evans implicitly anchors in *The Varieties of Reference* (and elsewhere), Sperber and Wilson do it throughout *Relevance* (and elsewhere), and I do it explicitly in ch 3 of *Descriptions* in connection with the effects of indexicals appearing in definite descriptions such as ‘the first person *I* saw this morning’, ‘my mother’, ‘the present king of France’, and ‘the girl who made *this*.’

Chomsky also does something analogous to anchoring in every work in which he discusses pronouns. (I say ‘analogous’ because of Chomsky’s concerns about reference).

although we may want to talk about the ‘linguistic meaning’ of, the ‘semantics’ of, the ‘conventions governing’ an indexical or any other expression, we may also wish to talk about its ‘semantic value’ relative to a particular index, the object conveniently assigned to it by an index in order that work on pressing matters is not held up needlessly.³⁷ There is no harm in such talk as long as everyone is clear about what is going on. ‘Semantic values’ in this sense, are just *stipulated interpretations*, and the anchoring it involves is quite consistent with the idea that the interpretation of indexical expressions is basically a pragmatic matter only steered by semantic constraints.

Although I have not seen the point discussed explicitly in the literature, I get the impression some ‘anti-pragmatist’ sentiment may have as its underlying source the worry that the sort of pragmatism inspired by Sperber and Wilson involves self-suspension, a willingness to abstract from contextual effects in ways that are self-defeating or paradoxical. In reality, the situation is not that different from Neurath’s. The pragmatist certainly has to tread carefully, all the while monitoring for and then abstracting from aspects of what is said that are not fixed by syntax and word meaning, and as a matter of working practice, subtle and silent measures are usually taken to prevent things becoming unmanageable, measures that certainly narrow the pragmatist’s options on the vexed matter of the relation between linguistic structure and the structure of thought.³⁸ Pragmatist abstractions from context are always going to be juggling acts, the artistry of which is rather like that involved in solving for several variables at once whilst looking for an unknown number of others that are not yet in the equation and cannot be located without extremely good approximate values for those that are. On the basis of perceived use, intuition, discussions with friends, books we have read and who knows what else, we isolate what we take to be the ‘linguistic meaning’ or ‘semantics’ of an expression α —its invariant role in determining what someone says by uttering sentences containing it—not unreasonably confident that certain things speakers mean when they use α in their speech and writing are explicable in very general terms as things they only imply, and are not of a gravity sufficient to make us question the meaning we think we have isolated. Holding the meaning of α constant, we go on to investigate and isolate the meaning of β and find we can make some headway by appealing to ‘facts’ about the meaning of α . And so on, until we get further up the alphabet and find our best attempts at meaning isolation force us seriously to question whether we were really right about all of α , β , γ , etc. and to re-examine our methods of abstraction and idealization, which may have led us to oversimplify in ways we now worry about.

³⁷ This convenience is employed time and again in ch 3 of *Descriptions*.

³⁸ If there is a worry about pragmatism, it surely resides here and not in minutiae like relativization (implicit binding).

This is basically the position Grice and others have found themselves in with the natural language counterparts of the logical particles, and certainly it takes some skill to keep all of the balls in the air in a stable configuration.³⁹

4. Elliptical Speech

Among the numerous overlapping and interconnected uses of the words ‘elliptical’ and ‘ellipsis’ we can, I think, isolate and make precise three of central concern to the philosophy of language, one corresponding to each of (a) what *A* implied on a given occasion, (b) what *A* said, and (c) the sentence *A* used. We can call these (a) *conversational* ellipsis, (b) *utterance* ellipsis, and (c) *sentence* ellipsis, respectively.⁴⁰

(a) *Conversational Ellipsis (Elliptical Remarks)*. When asked by a colleague for his opinion of one of his students, a professor, says, ‘Smith has wonderful handwriting and is always punctual.’ The example, due to Grice (1961), is meant to illustrate the difference between saying something and merely implying it (‘conversationally implicating’ it, as he puts it in later work). People unfamiliar with philosophical discussion in this area might say that Smith was ‘damned with faint praise’ or that the professor chose a rather ‘indirect’ or ‘elliptical’ or ‘circuitous’ or ‘roundabout’ way of letting his colleague know what he thought of Smith. The thought behind these forms of words is that the professor has *avoided* using certain words, and the beauty of the word ‘elliptical’ here is the connotation not only of something omitted or avoided but also of a path that curves, a path that is ‘circuitous’ or ‘roundabout’, and hence longer and less direct. (Sometimes speech or writing is also said to be elliptical when it so concise or compressed as to be difficult to understand, an idea clearly related to omission again.) Conversational ellipsis is very common in everyday speech and writing, and only a full-fledged pragmatic theory will throw any light on how we manage to get away with it.

There is a modal dimension to conversational ellipsis: it is usually possible to find some alternative form of words the speaker or writer *could have used* to make his point more directly or less elliptically. Indeed the relevant form of words is often used, after the fact, to specify what the speaker or writer was implying: By saying ‘he has wonderful handwriting and is always punctual,’ the professor means that *Smith is no good at philosophy*.⁴¹ But as Grice stressed, there may be a considerable degree of

³⁹ See, in particular, Carston’s (1988, 1993, 2002) successively refined analyses of the meaning and use of ‘and’.

⁴⁰ Cf. Gk. *elleipsis* lit. a condition of falling short, from the verb *elleipein*: to leave in, leave out, or fall short.

⁴¹ As I have argued elsewhere, Grice’s idea that a conversational implicature is ‘cancellable’ is best described modally (in the same situation, the speaker could have made an utterance that is a continuation of the utterance he did make, by which he would not have conversationally implicated what he did in fact implicate (e.g. ‘Smith has wonderful handwriting and is always punctual; what’s more, he’s erudite, one of the best philosophers of his generation’)).

indeterminacy—did the professor really mean that Smith is no good at philosophy rather than, say, Smith is not particularly bright or Smith does not deserve a fellowship for next year?

(b) *Utterance Ellipsis (Elliptical Utterances)*. When discussing definite descriptions in his 1940 book *Mathematical Logic*, W. V. Quine says,

Everyday use of descriptions is often elliptical, essential parts of the condition ‘...x...’ being left understood; thus we may say simply ‘the yellow house’ . . . when what is to be understood is rather ‘the yellow house in the third block of Lee Street, Tulsa’ (1940: 146).

Obviously Quine is *not* here claiming that someone who uses ‘the yellow house’ is not using a whole noun phrase; and he is not claiming that someone who uses

(1) the yellow house is on fire

is not using a whole sentence. It would be absurdly uncharitable to attribute such views to anyone, and here they would, in any event, fly in the face of the text: Quine, ever picky about his words, is careful to say that everyday *use* of descriptions is elliptical, and not that everyday descriptions (the expressions themselves) are elliptical. What Quine seeks to get across can be captured, without remainder, by saying that someone may speak elliptically using a perfectly well-formed description.⁴² The point seems obvious, and I suspect Quine saw himself as expressing something of a platitude: the linguistic expression ‘the yellow house’ is not itself elliptical, but everyday uses are: since there is more than one yellow house, understanding an utterance of (1) will involve recovering more than the lexical meanings of ‘the’, ‘yellow’, ‘house’, ‘is’, ‘on’ and ‘fire’ and projecting these in some way in accordance with the way they are here syntactically combined. The basic point seems undeniable.

In a similar vein, Wilfrid Sellars in his 1954 article, ‘Presupposing’ says that,

a given utterance of [‘The table is large’] is elliptical and states what would be nonelliptically stated, for example, by ‘The table *over here* is large’ . . . the context

⁴² The noun ‘use’ is deployed in several distinct but related ways in the philosophy of language. Sometimes it is used in the sense of ‘type of use’ or ‘way of using’, as, for example, when there is talk of the *bound* and *free uses* of pronouns, of the *attributive* and *referential uses* of definite descriptions, of Strawson’s *uniquely referring* use and of the *generic* (‘the whale is a mammal’) *use* of descriptions. At other times it is used in the sense of ‘instance of a way of using’, as, for example, when there is talk of a particular bound *use* of a pronoun, a particular referential *use* of a description, a particular uniquely referring *use* of a description. (Perhaps these are not different uses in any very interesting sense, just examples of polysemy. Can we usefully think of particular elliptical uses of descriptions as instances of the *elliptical* use of descriptions? Surely there is nothing to gain by so-doing. Some have argued there is nothing to gain by analogously typing particular *referential* uses of descriptions either, at least as far as *semantic* theory is concerned. Strawson himself takes advantage of another use of ‘use’: if two men utter the sentence ‘the king of France is wise’, one in the reign of Louis XV the other in the reign of Louis XIV, ‘each made a different use of the same sentence’; if two men utter the sentence simultaneously in the reign of, say, Louis XIV, ‘[they] made the same use of the same sentence’ (1950: 325-6). For discussion, see Neale (forthcoming c)

functions to give the statement the force, for example, of ‘The table *over here* is large’ (1954: 199).⁴³

Obviously Sellars is *not* here claiming that someone who uses ‘the table’ is not using a whole noun phrase; and he is not claiming that someone who uses

(2) the table is large

is not using a whole sentence. It would be absurdly uncharitable to attribute such views to anyone, and here they would, in any event, fly in the face of the text: Sellars, ever picky about his words, is careful to say that a given *utterance* of (2) may be elliptical, not that the sentence (2) itself is. Sellars is surely making the same point as Quine: although the linguistic expression ‘the table’ is not itself elliptical, one may use it elliptically in the sense that understanding an utterance of a sentence that contains it will typically involve recovering more the lexical meanings of ‘the’, ‘table’, ‘is’ and large, and projecting these in some way in accordance with the way they are here syntactically combined. Again, the basic point seems incontrovertible. And in stating the point Quine and Sellars very rightly steer away from talking about elliptical expressions *per se* and introduce their respective speech act words: ‘use’ and ‘utterance’.⁴⁴ These guys were no slouches.

There is no difference of any importance to any point I shall make here between Quine’s use of ‘use’ and Sellars’s use of ‘utterance’ in the passages quoted above; in what follows I shall use Sellars’s terminology (as I did in *Descriptions*). So when Quine says that ‘Everyday use of descriptions is often elliptical,’ we can transpose this into Sellars’s terminology by construing Quine as saying, ‘Everyday utterances of descriptions are often elliptical.’⁴⁵

⁴³ See also Reichenbach (1947). Talking of the description ‘the train’ as it occurs in ‘the train will arrive at 7 P.M.’ Reichenbach says, ‘The necessary addition then is understood. It usually consists in a reference to a preceding utterance; for instance, it may be assumed in the form “the train of which we spoke”.’ (1947: 258).

⁴⁴ See also Bach (1981, 1987, 1994). The distinction was particularly important to Sellars, who was responding, in part, to Strawson’s (1950) complaint that Russell failed to distinguish linguistic expressions and particular dated utterances (or uses) of those expressions. (See below.)

⁴⁵ We may call a particular dated *utterance* of an expression a particular dated *use* of the expression or a particular dated *tokening* of it without too much harm—which is not to say we might not want to use ‘use’ and ‘utterance’ differently elsewhere. See Strawson (1950) and note 43. For convenience I will stick to talk of particular dated *utterances*. The word ‘utterance’ (also ‘assertion’, ‘statement’, ‘remark’, ‘question’, ‘inscription’, ‘expression’, ‘interjection’ and others) has a convenient act-object (or process-product) ambiguity that can be nicely exploited as long as great care is taken (as it is by Sellars (1954) and also by Grice (1989), for example). It would be quite an undertaking to sort out the logical grammar of such devices, and I have no intention of getting bogged down in it here (although I am strongly inclined to say that someone who utters ‘Tom’s utterance was loud but true’ or ‘his question was well-timed and very relevant’ or ‘his outburst was justified and broke a window’ is making a category mistake (perhaps in order to make a joke)). When ‘utterance’ is used for the object/product in, say, ‘my utterance of X’, at least some of the time it is replaceable by ‘what I said by uttering X’ or by ‘what I said by my utterance of X’ in which ‘utterance’ is used in the sense of act/process. And it is for this reason, I think, that many of us are drawn to talk of true or false utterances. When I talk of utterances being true or false this is what I shall have in mind in what follows. There are many other adjectives that we put alongside ‘utterance’: ‘reassuring’, ‘frightening’, ‘intelligent’, ‘hurtful’, ‘courageous’, ‘bold’, ‘rude’, ‘sentimental’, ‘loud’, ‘ironic’, ‘metaphorical’, for example. We must exercise caution here: in many cases we are predicating something of what the speaker said, of the fact that something was said, of the speaker’s act of saying what he said, or of the particular

Again, there is a modal dimension, which Sellars makes explicit in the passage just quoted: the utterance states ‘what would be nonelliptically stated’ by an utterance of some longer sentence. In *Descriptions*, I called this the *explicit* response to the problem of so-called incomplete descriptions (it is sometimes called the *ellipsis* response in the literature). The connotation I had in mind was the possibility of producing a lengthier utterance that was more *explicit*, i.e. one that more *explicitly* specified what the speaker was saying.⁴⁶ According to the explicit response, utterances of, say, ‘the table’ (or ‘every table’) are elliptical for utterances of richer (‘more explicit’) descriptions, as suggested by Quine and Sellars. (I took Sellars’s discussion, with its explicitly modal characterization, as the locus classicus of the explicit response, although I did also mention Quine’s earlier discussion). No thesis about *syntax* is implied by Quine and Sellars, by my giving a label to the general response their remarks typify, or by anything else in my discussion.

It would be a mistake to think that Quine and Sellars are making a point they take to be specific to utterances of *descriptions* (or even utterances of quantificational expressions generally); and it would be a mistake to see them as proposing a *theory* in any interesting sense or positing some sort of psychological *process* that will form part of a theory. They are just making an elementary observation about a widespread phenomenon—context-sensitivity—in the use of language, and pointing out that the phenomenon itself constitutes no refutation of Russell’s Theory of Descriptions (which is not to deny that it may be possible to find specific cases of incompleteness that are more problematic). As Bach (1981) observes in a discussion of elliptical uses of descriptions,

if I say ‘I drink only Scotch’, I would be stating not that I drink nothing but scotch but merely that the only liquor I drink is Scotch . . . The phenomenon of elliptical speech is commonplace, indeed, it often seems stilted not to suppress words that can easily be inferred . . . Using incomplete descriptions elliptically . . . is just another case of this familiar phenomenon (1981: 238).

Like Quine and Sellars, I was not rash enough to offer a theory about how we manage to *pull off* the interpretation of incomplete utterances; I posited no psychological *process* to form part of a theory, and I held out no promissory note. I was defending a particular *semantic* proposal, a proposal about the linguistic meaning of the word ‘the’ as it occurs in expressions of the form ‘the ϕ ’, a proposal that seems to be threatened by incomplete utterances. It was no part of my task to provide a theory that explains how we manage to identify what a speaker is saying and implying on the basis of the meagre evidence we

words used in the act. (If Bill, who has been knocked down and seems close to death, manages to utter the words ‘I am dying’, his utterance may be reassuring in one sense but not in another). All sorts of dangerous ambiguities lurk here because of issues about the applicability of adjectives such as those listed above to the sorts of things that ‘statement’, ‘utterance’, ‘inscription’, ‘remark’, ‘question’ and so on are applied to. There is clearly an enormous amount of work to be done here, but I think I can get across the main points of the present article without getting bogged down in it. This use of ‘elliptical’ in connection with incomplete descriptions has been borrowed by many people. See below.

⁴⁶ I do not recall whether I intended a nod in the direction of Sperber and Wilson’s (1986) talk of ‘explicit content’ in my choice of ‘explicit’ over ‘overt’, as Zsófia Zvolensky has suggested. It is quite possible.

get from knowledge of syntax and word meaning (a theory of *interpretation* in the sense described earlier). That is a task for cognitive science. Sperber and Wilson (1986, 1995) appear to have made some headway here, and the observation made by Quine and Sellars provides grist for their mill. As I said in *Descriptions*:

As (e.g.) Sellars (1954) Sperber and Wilson (1986) have stressed, in many cases the linguistic meaning of the sentence—or sentence fragment—uttered radically underdetermines the proposition it is used to express on a given occasion. We have already considered the sort of contextual supplementation that that is required where an utterance contains overtly indexical or demonstrative components; but context-sensitivity does not end there (1990: 114, n. 46).

In short, I was quite explicit about the larger problem of the underdetermination of what is said, that Quine, Sellars, and Sperber and Wilson provide examples of. And I took the standard philosophy cop-out: that is, like practically every other philosopher who discusses this matter, I did the philosophical equivalent of taking the fifth: when we wish to avoid discussing genuine psychological processes (being philosophers not psychologists) but also wish to emphasize that we think pragmatic, inferential processes must be at work, we gesture in the direction of Grice or at least in the direction of ‘Gricean considerations’. *Useless* to cognitive science, of course, but our way of acknowledging that we are out of our depth and that we realize there is serious work to be done by cognitive psychology.

Our interpretive abilities are so good that we can reasonably expect our addressees (and often those who overhear) to identify the thoughts we seek to express even when we use expressions whose linguistic meanings fall short of serving up the precise concepts involved in the thought. The point is easily seen with predication (indeed, it might be argued that the point is *about* predication). Often a speaker will use a simple predicate, even if a richer or more complex predicate might be used, one that could, in principle, leave the hearer with less inferential work to do. (To say this is not to say that use of the richer predicate will speed up communication—it could slow it down). The broadly pragmatist literature brims with examples which might be borrowed to make this point (the parenthetical expressions making more explicit what was left only implicit in a particular conversational setting):⁴⁷

I haven’t had breakfast (this morning)
 It’s snowing (in Reykjavik)
 Maria wants to get married (to Fred)
 Maria and Fred want to get married (to one another)
 Maria and Fred pushed the car to the garage (together)
 I hadn’t noticed (that Mike was limping)
 Maria wants to leave (this party)
 Maria is ready (to leave (this party))

⁴⁷ For dozens more examples, see Bach (2001a), from which a few of those above are also borrowed.

I haven't seen Maria (here tonight)

You are not going to die (from that injury).

What is true of the predication inherent in using verb phrases (as in the examples above) is true of the predication inherent in using noun phrases:

everyone (at Ragga's party last night) had a great time

no-one (from the U.S. embassy) has arrived yet

the car (Tom bought this morning) broke down on the way (to Tom's) home

the Russian (judge) voted for the Russian (skater)

the (former) hostages were greeted at the White House

(the man on) table six wants to change his order

the (man who ordered a) ham sandwich on (table) twelve wants pickles

every farmer (in my village) owns exactly one donkey and feeds

the donkey (he owns) at night.

Something similar appears to be going on with the following where an overall 'point of evaluation' may be suppressed:

(In the Sherlock Holmes stories) Moriarty was Holmes's arch-enemy

(In Homer) Penelope is the wife of Odysseus

(In Greek mythology) Ares is the god of war.

There is nothing of great theoretical significance in the groupings, and certainly no syntactic thesis is intended. The first group merely illustrates cases in which a richer sentence of the sort required may be provided by beefing up the predicate that constitutes the VP; those in the second group merely illustrate cases in which a richer sentence of the sort required may be provided by beefing up the predicate inside the subject noun phrase; and the third merely illustrates cases in which a richer sentence of the sort required may be provided by adding some sort of sentential operator that can be used to indicate an overall 'point of evaluation.'

Talking involves a trade off. More detail can lead to greater precision and reduce misunderstanding, but it can also sound stilted, as Bach said, or bore or confuse an audience. We all know people who use too few or too many words for our own tastes, and we have all been in situations where we have used too few or too many ourselves—sometimes we have suffered accordingly. All of this is pretty obvious really, and it is somewhat surprising that philosophers have tied themselves in knots over examples, which are so *linguistically* trivial. (As far as examples involving incomplete descriptions are concerned, Quine (1940) and Sellars (1954) realized just how trivial half a century ago.) What is *not* trivial, of course, is the articulation of a general *pragmatic* theory that answers the following question: how it is that we manage to identify what a speaker is saying—let alone implying—by uttering a sentence *X* on a given occasion given the precious little information we obtain by virtue of our knowledge of the meanings of the words in *X* and our knowledge of *X*'s syntactic structure. It is an empirical theory of

utterance interpretation in this sense that people like Sperber and Wilson (1986) are trying to construct. Obviously Quine and Sellars were not presenting any sort of theory of that sort: they were just making commonsense observations.

(c) *Sentence Ellipsis (Elliptical Sentences)*. Bach (1987) and Ostertag (1998, 1999) explicitly warn against confusing the speech act notion of utterance ellipsis just discussed with a quite different syntactic notion that emerged from talk in generative grammar of deletion transformations. Here is Ostertag:

the notion of ellipsis appealed to is not the one familiar from syntactic theory ... whatever the relationship is that holds between [a sentence containing an incomplete description] and [a sentence containing a completed one] it is not syntactic. (1998: 20).

There seems no reason to suppose that the relation between, say, ['the table is covered with books'] and its completion is anything like that between: 'John left and Bill did too' and 'John left and Bill left'. (1999: 144, n. 3).

In a recent encyclopædia article on ellipsis, Robert May (2002) provides a succinct description of the notion of ellipsis familiar from syntactic theory:

A linguistic ellipsis, most generally expressed, is a truncated or partial linguistic form. This partiality is measured relative to a complete sentence; an elliptical sentence is one in which some of the constituent parts of a 'full' sentence are missing (2002: 1094)

The following forms of words are elliptical in this sense:

- (3) Peter went to Paris, Bob to Brussels
- (4) Peter has learned to drive, but Bob hasn't.

Do the expressions 'Bob to Brussels' and 'Bob hasn't' constitute 'whole sentences'? Ultimately we will want to defer to empirical linguistics here, and given the current state of play this will probably mean positing (a) various types of aphonic expressions, and (b) syntactic structures that resist the sort of description that can be produced by imposing constituent structure (using pairs of labelled brackets, for example) on the word sequences of our usual, non-technical orthography. For concreteness, then, let us follow May (and common practice in linguistics) in saying that we *are* dealing with sentences: the linguistic forms are 'partial' or 'incomplete', relative to complete sentences, because they are

'linguistic form[s] in which constituents normally occurring in a sentence are superficially absent, licenced by structurally prior antecedents (2002: 1094)

Certainly I *speak elliptically*, in the sense of the second grade of ellipsis, when I utter (3) or (4): that is, my *utterance* is elliptical. But surely there is something more going on here, something that even traditional grammars recognize. Let alone generative linguistics. The fact that my utterance is elliptical can be explained in *linguistic* terms. As we might put it, the elliptical nature of my utterance is attributable to the fact that the *sentence* I utter is elliptical for some longer sentence that can be straightforwardly 'recovered':

- (3') Peter went to Paris, Bob went to Brussels
 (4') Peter has learned to drive, but Bob hasn't learned to drive.

The modal dimension in sentence ellipsis is clear: there is some longer surface form the speaker could have uttered for which the surface form actually uttered is elliptical, a sentence that can be restored from the syntactic context.

Putting things rather simplistically for a moment, we might say that as a matter of *grammar* (i) we can here *recover* the ellipsed elements by *copying* them over from elsewhere (subject to certain specifiable constraints), and (ii) we can create an elliptical sentence by *deleting* elements that are *duplicated* elsewhere (again subject to certain specifiable constraints). This is very rough, but it's good enough to set the scene.⁴⁸ In this respect, sentences (3) and (4) are unlike sentences (1) and (2), and this is surely because the (1) and (2) are not elliptical sentences.⁴⁹

So back to our question: are (3) and (4) whole sentences? Let generative linguistics decide; we'll call them 'elliptical sentences' and if generative grammar says they are sentences, no harm is done; equally, if generative grammar says they are not sentences then 'elliptical' in this context is rather like 'plastic' in 'plastic flowers' or 'counterfeit' in 'counterfeit banknote', and still no harm is done. So, given, the current state of play, as presented by May (2002), let's call them elliptical sentences.

5. The Incompleteness Question

The label 'incomplete description' is misleading. But we need to begin somewhere, so let us have some preliminary definitions. Let us say for the moment that a description is proper if, and only if, its nominal—or its superficial matrix in some standard system of representation—is true of exactly one thing, and *improper* otherwise. And let us say that an *improper* description is *empty* if it is true of nothing, and *incomplete* if it is true of more than one thing. (As the need arises, we can tolerate loose but intelligible talk of the matrix of an English description being incomplete, and this will allow us to move back and forth between talk of the matrix *table x*, for example, and (loose) talk of the matrix 'table'. No confusion should arise.)

Incomplete descriptions are supposed to be interesting because of a *question* they force the Russellian to answer, one simple variant of which might be put thus: How are we to explain the incontrovertible fact that *A* can use a description 'the ϕ ' in an utterance of the simple form 'the ϕ is ψ ' and thereby perform a perfectly felicitous speech act, indeed *say*

⁴⁸ Issues to do with (e.g.) binding make a final statement more complex. See below.

⁴⁹ Of course there may be cases in which it is easy enough to create a longer sentence by drawing explicitly upon linguistic material elsewhere (as in 'the yellow house in the third block of Lee Street, Tulsa is on fire. The house was built in 1865'). But that is not the point.

something true, even though *A* and *B* (the hearer) both know that $\phi(x)$ is true of more than one thing? The question generalizes: How are we to explain the fact that (roughly) for a range of determiners, *D*, *A* can use '*D* ϕ ' in an utterance of the simple form '*D* ϕ is ψ ' and thereby perform a perfectly felicitous speech act, indeed *say something true*, even though *A* and *B* both know that $\phi(x)$ is true of some things that are not relevant to the truth or falsity of what *A* said?⁵⁰

Many philosophers appear to think the answer to the question the Russellian must answer is obvious. 'There's always an *implicit background restriction* on the domain over which a quantifier expression ranges,' is one old reply. Another is, 'An utterance of 'the ϕ is ψ ' is sometimes *elliptical* for an utterance of 'the ϕ that ζ is ψ ', where ζ is something the speaker could have made explicit but didn't.' Call these the *implicit* reply and *explicit* replies, respectively, based only on the appearance of the words 'implicit' and 'explicit' in the quoted remarks. Quine (1940), Sellars (1954), Sainsbury (1979), Davies (1981), Evans (1982), and many others have replied in one or both of these ways, and when I was writing *Descriptions* I thought their remarks were sensible enough, the sort of general remarks that set a *target* a pragmatic theory should reach, explaining as it should how we integrate information linguistic and non-linguistic information to work out what the speaker is saying.

I spent very little time comparing or even discussing the implicit and explicit responses because (for better or worse) I saw incomplete descriptions as posing a pretty spurious threat to the account of definite descriptions I was defending. I argued that methodological considerations of the sort advanced by Grice and Kripke strongly favoured a unitary Russellian analysis of descriptions, and that the usual arguments for semantically distinct referential readings—including the Argument from Incompleteness—failed to demonstrate the desired conclusion, indications to the contrary being largely false impressions engendered by inattention to the distinction between what a speaker *says* and what he *means*. The discussion in *Descriptions* was, I think, rather lazy, and I now believe only half of it. To be precise, I still think the Russellian analysis is basically correct for both attributive and referential uses of descriptions; but I no longer think the difference between *saying* and *meaning* lies at the heart of a characterization of referential usage, and I want to provide a more satisfying (less lazy) account of the distinction.

What prompted me to seek something more satisfying? Interesting problems for the unitary Russellian position I favoured in 1990 posed by Bezuidenhout (1997), Carston (2002), Devitt (1997a, 1997b), Devitt and Sterelny (1999), Larson and Segal (1995), Récanati (1993, 1996), Ramchandran (1993, 1995), Reimer (1992, 1998a), Rouchota (1992), Schiffer (1995), Wilson (1991), and Zvolenszky (2000), problems that either take the form of counterexamples exploiting apparent weaknesses in one or other of the

⁵⁰ The words 'roughly', 'range', and 'simple' appear here so that I can say something straightforward without bringing up issues to do with negation, monotonicity, persistence, and binding. See Neale (forthcoming c).

explicit and implicit approaches, or else build on foundational or methodological worries revolving around the fact that referential uses of descriptions are common, standard, systematic, and cross-linguistic.⁵¹ It is heartening to see one's views presented clearly, improved upon, criticized cogently, even replaced in print, to see ingenious counterexamples and insightful methodological criticisms in work that advances our understanding whilst displaying mastery of the subject matter and great sensitivity to the text (and the intentions behind it) and. The works just mentioned all do one or more of these things and have advanced my understanding of the issues considerably. Not only were they instrumental in leading me to change my mind on one or two of them, more importantly they led me to see certain things in new ways and to believe the central debate in this area is the product of a powerful illusion. Explaining that illusion is a central task in what follows.

Before getting to that illusion, however, certain others need to be dispelled. Our question, recall is, 'How are we to explain the incontrovertible fact that a speaker can use a description 'the ϕ ' in an utterance of the simple form 'the ϕ is ψ ' and thereby perform a perfectly felicitous speech act, indeed *say something true*, even though he and his hearer both know that $\phi(x)$ is true of more than one thing?' By deigning to provide even a vague answer to *this* particular question—rather than rejecting it or answering a question with which it might be confused—one has, in effect, already accepted a central tenet of linguistic pragmatism, the *underdetermination thesis*, found in embryonic form in Quine (1940) and Sellars (1954) in their discussions of incomplete descriptions, lurking in the work of Austin (1962), and articulated clearly by Sperber and Wilson (1986). It is a point familiar from discussions of *context-sensitive* or *indexical* expressions such as 'that', 'he', 'I', and 'you' that knowing what sentence has been uttered on a given occasion and knowing its linguistic meaning may be insufficient for identifying what the speaker said. The interpretation of utterances containing incomplete descriptions extends the point, for even when the references of any overt indexicals and other referring expressions have been identified, there may be further context-sensitivity to be resolved. This was something I tried to make clear in *Descriptions*, where I explicitly accepted the

⁵¹ Bezuidenhout, Devitt, Ramchandran, Récanati, are primarily interested in motivating a supplementary, semantically distinct, referential interpretation. Reimer is more concerned with explaining why the explicit/ellipsis approach is inferior to the implicit/domain restriction approach, a novel variation of which she motivates and seeks to distance from anything that could be construed as a notational variant of the explicit approach. Reimer's discussion displays a keen appreciation of the virtues of the implicit approach, but ultimately I think the attempt to separate it from the explicit approach fails (see below). Wilson and Ramchandran aim is to undermine the unitary Russellian account by producing examples apparently requiring a semantically distinct referential interpretation to supplement the Russellian interpretation. In this way, Wilson attempts to motivate his own 'pronominal' theory of referentially used descriptions. Schiffer's examples are presented in a broader context, as part of an elaborate argument designed to produce a dilemma for the Russellian who also holds a direct reference theory of indexicals, itself part of a larger argument designed to produce a dilemma for the hidden-indexical theory of attitude reports. Larson and Segal are rightly concerned with producing a theory that is satisfactory from a syntactic as well as a semantic perspective, and see the way in which syntax and semantics come together as strongly suggesting the need for a supplementary referential interpretation for certain descriptions used referentially.

underdetermination thesis, as already noted. Underdetermination that goes beyond that induced by ‘overtly indexical or demonstrative’ or other referring expressions in a sentence *X* may be of two types, distinguishable in respect of their syntactic commitments. (i) It may be due to Sellarsian utterance ellipsis, a notion that has no syntactic dimension. Or (ii) it may be due to the existence of covert, i.e. aphonetic indexical or demonstrative elements in *X*’s syntax, elements just as much in need of interpretation as occurrences of the phonic elements ‘he’ or ‘that’.

Elementary reflections on utterance interpretation and empirical work on syntactic structure suggest very strongly that we should accept the existence of both utterance ellipsis and aphonetic items in syntax. Nonetheless, there are those who seek to deny the existence of one or the other, or to reduce all cases of one to cases of the other, motivated it would seem by either a prejudice against pragmatism or a prejudice against Chomskyan aphonics.⁵² On the assumption that we are involved in an empirical exercise that looks as though many of the tenets of pragmatism and Chomskyan syntax will loom large, I see no good reason to think either form of reduction will be successful, so I will assume the existence of both utterance ellipsis and aphonetic elements in syntax until a good argument against one or the other comes along.

The postulation of a crucial aphonetic in every sentence containing a description does not really involve rejecting the original question the Russellian has to face: for aphonics are just as much in need of interpretation as phonics. Rather positing an aphonetic constitutes a tottering first step towards providing an *answer* to the question, and does not change the basic point: knowing the blueprint for ‘the table is brown’ (by virtue of knowing its syntax and the linguistic meanings of the words ‘the’, ‘table’, ‘is’ and ‘brown’) does not suffice for grasping what *A* has said by uttering it on a given occasion.⁵³ And the mere act

⁵² The former prejudice appears in the work of Stanley and Szabó (2000a,b), Stanley (2000, 2002a,b,c), the latter in the work of Barwise and Perry (1983).

⁵³ To the best of my knowledge the only sustained attempt to deny this has come from Bach (1981, 1987, 1994, 2000), who argues that if there is more than one ϕ , the speaker actually *said* something false by uttering ‘the ϕ is ψ ,’ but nonetheless *meant* something true. In *Descriptions* I was pretty short with this idea on the grounds that it is inconsistent with a basic tenet of linguistic pragmatism: it clashes with our intuitive judgments of truth and falsity by virtue of clashing with our intuitive judgments about what the speaker is saying, which I take to be the principal data a theory of interpretation is meant to explain. (Soames (1986), Reimer (1998), and Stanley and Szabo (2000a,b) also reject this type of approach.) Following Cohen (1971), the standard way of making trouble for theories that propose a thin notion of what is said in a particular type of example is to bury it inside a conditional. Thus Cohen objected to Grice’s thin (truth-functional) account of ‘and’ by contrasting ‘if ϕ and ψ then ζ ’ and ‘if ψ and ϕ then ζ ’. (For detailed discussion of this tactic, see Carston (2002).) So one is naturally drawn to examine examples such as the following in connection with Bach’s claim about descriptions: (i) If the table is dirty, then the waitress will be dismissed. I take it the existence of a dirty table at Famous Ray’s Pizza could not be cited by a waitress at Tom’s Diner in her defence when faced with dismissal (from Tom’s Diner). In order to have a framework within which to discuss the relevant issues, I am assuming, as I did in *Descriptions*, that we can get by with just the two most intuitive components of what a speaker, *A*, meant, at least for the purposes of discussing definite descriptions, viz. what *A* said (the proposition(s) *A* expressed) and what *A* implied (the propositions *A* implied). Within *that* framework, giving a direct answer to our original question means accepting that incomplete descriptions introduce a measure of insufficiency, and rejecting that question means denying they do. This is not to say, of course, that one cannot construct a more sophisticated framework within which one distinguishes what *A* says

of positing a context-sensitive aphonic lurking in the sentence's LF that gets interpreted in the right way thereby making everything work out just right does not itself constitute an explanation of *how* it gets interpreted in the right way (if only producing a theory of interpretation were so easy!). One still needs to explain the *semantics* of this new element, explain the sorts of values the speaker can intend it to have on different occasions, and then explain how hearers go about identifying these values. The last of these is something a *pragmatic* theory must explain, and the explanation is going to be very similar to the one given for words like 'he', 'she', 'it', 'this', 'that', and 'here'. A semantic theory explains the sorts of values the speaker can intend, say, 'he' to have on different occasions, and a pragmatic theory explains how hearers go about identifying these values on particular occasions.

Only a full-fledged pragmatic theory can explain how the speaker is able to decide how thin a description he can get away with, being reasonably certain the hearer will identify what he intends to be saying. And only a full-fledged pragmatic theory can explain how hearers do in fact identify what speakers are saying when they are using incomplete descriptions. A cognitive theory of the sort being constructed by Sperber and Wilson is what is needed here—not the sort of Mickey Mouse Gricean theory I was using in *Descriptions* to illustrate the intuitive and, one would imagine, theoretically significant, distinction between what a speaker says and what he means but does not say. I make no excuse for not having addressed the big cognitive questions.

What I could have said more about, however, were the labels 'incomplete', 'implicit', and 'explicit'. The first had been used for many years in connection with descriptions. For better or worse, the second and third were picked up from *Descriptions* by philosophers and linguists attempting to set out the virtues and vices of one or both of the general strategies. Labels are only tags, of course, but often we select this or that one because of what it connotes; and this can, on occasion, increase rather than decrease the risk of misunderstanding. The connotations I had in mind for 'explicit' and 'implicit'—I contemplated 'overt' and 'covert'—were these: the possibility of producing a windier utterance that *explicitly* (*overtly*) specifies what an incomplete quantifier, in context, is taken to express; the existence of an *implicit* (*covert*) limitation on the domain of quantification. I shall say more about 'explicit' and 'implicit' in the next section. First I need to say something about 'incomplete'.

The label 'incomplete' is misleading in its application to a type of linguistic expression such as a definite description (similarly 'empty' and 'improper').⁵⁴ The label—which

and A states, for example, as Bach and Harnish (1981) and Bach (1987, 1994) do; or distinguish what A says and what A's *utterance* says. It may well be the case that the use of incomplete descriptions is precisely the sort of thing that *motivates* making such distinctions, and perhaps I am remiss in not making all the distinctions I mentioned at the beginning of ch 2 of *Descriptions*.

⁵⁴ By talking of the definite description as a type of expression, I do not mean to be suggesting that the concept of a definite description is basic enough to appear in any final syntactic or semantic theory, unlike say the concept of a determiner.

appears to have beaten out such rivals as ‘imperfect’, ‘indefinite’, and ‘elliptical’ — seems to have emerged from Sellars’s discussion, and it may seem strained at first because of an unfortunate connotation. The word ‘incomplete’ carries a suggestion of something *missing* and of something *falling short*, something that *needs completing*. And so care is required, for it is easy to slip back and forth between different objects that seem to be in need of completion. Of course no-one seriously involved in the debate over incomplete descriptions over the past sixty years or so construes it as a debate about points of *grammar*. That is, no-one takes it to be a debate about grammatically incomplete or grammatically deficient expressions in need of additional linguistic material in order to be turned into grammatically well-formed descriptions; no-one actually thinks that ‘the table’, ‘the house’ and so on fall short of being complete English noun phrases, or that the ‘the table is large’ or ‘the house is on fire’ fall short of being complete English sentences; and no-one thinks that the stock explicit and implicit responses to the original question involve methods for turning grammatically incomplete or defective expressions into ones that are not grammatically incomplete or defective. Grammar is simply not the issue.

So what sorts of things have we really been attributing *incompleteness* to for the past sixty years? The remarks by Quine and Sellars quoted above suggest we have been talking all along about incomplete *uses or utterances* of descriptions. Recall that they brought the suggestive word ‘elliptical’ into the debate in the course of sketching their own answers to the question the Russellian must answer. They talk of elliptical ‘uses’ (Quine) or elliptical ‘utterances’ (Sellars) of descriptions, and not of descriptions *per se* being elliptical. According to Sellars, an utterance of ‘the table’ will typically be elliptical for an utterance the speaker could have made of a richer description such as ‘the table over here’ or ‘the table beside me’ (1954: 200). The connection between ellipsis and incompleteness in Sellars’s thinking manifests itself when he says (i) that ‘in *ellipsis* the context completes the utterance and enables it to say something which it otherwise would not, different contexts enabling it to say different things,’ (ii) that some ‘utterances ... are *not* complete and are only made complete by the context in which they are uttered,’ and (iii) that ‘statements which are non-elliptical ... do not depend on their contexts for their *completion*’ (1954: 200). Drawing upon these early discussions, we might talk of incomplete ‘utterances’ of descriptions. On one occasion, the speaker and hearer may both know there is exactly one ϕ , in which case an utterance of ‘the ϕ ’ will *not* be incomplete. On a later occasion they may both know there is more than one ϕ , in which case an utterance of ‘the ϕ ’ *will* be incomplete. Of course a time the speaker refers to (or describes) *in* the utterance (rather than the time *of* the utterance) might be the relevant one. If speaker and hearer know that in 1950 there was exactly one ϕ and that in 1999 there was more than one A , an utterance of ‘the ϕ ’ occurring as part of an utterance of ‘In 1950 the ϕ is ψ ’ (in which ‘in 1950’ has larger scope) will *not* be incomplete; but an utterance of ‘the ϕ ’ occurring as part of a simultaneous utterance of ‘In 1999 the ϕ is ψ ’ (in which ‘in 1999’ has larger scope) *will* be incomplete. (Further issues are raised, of

course, because of sequences of tense and mood operators used to signal connections with one another or with a time frame already being assumed in the discourse.)

So when we use the expression ‘incomplete description’, we should probably construe this as shorthand for the rather ugly expression ‘incomplete utterance-occurrence of a description’ or as shorthand for the less ugly but very long-winded ‘description that is incomplete relative to a particular utterance-occurrence’.⁵⁵ No wonder we use ‘incomplete description’! (*Mutatis mutandis*, for talk of an ‘empty description’ and an ‘improper description’.)

Another, perhaps preferable, way of legitimizing ‘incomplete’ as it applies to descriptions emerges if we set the discussion against the sort of background against which it should be set: a theory of utterance interpretation in the sense discussed earlier. The leading idea is that a *complete interpretation* of an utterance of a sentence $X(\delta)$ containing a description δ cannot always be extracted from the linguistic form of $X(\delta)$ alone; in particular, a complete interpretation of the sub-utterance of δ cannot always be extracted from the linguistic form of δ alone: contextual considerations have to be exploited by the hearer in order to identify the speaker’s intentions.⁵⁶

This way of stating things turns out to be rather illuminating and avoids prejudicing the issue in favour of either the Unitarian or the Ambiguity Theorist. The central disagreement between these theorists will be about the *form* of complete interpretations in cases of referential usage, the Unitarian arguing that they are to be expressed in terms of *general* propositions, the Ambiguity theorist arguing that they are to be expressed in terms of *singular* propositions. Attacks on either position will take the form of attacks either on the *truth conditions* of the favoured proposition or on the favoured *method of completion*.

Whether or not they are correct, the truth conditions supplied by Russell’s theory for utterances of sentences containing (sub-utterances of) improper descriptions are at least clear, and they flow directly from the theory without further ado (in this respect the theory differs markedly from many other theories). Nonetheless, incomplete descriptions appear to present interesting challenges. The (optimistic) line I took in *Descriptions* was basically that incomplete descriptions were an argumentative dead-end: no matter how clever or spare the matrix, how fiendish the context, it was always possible, I suggested, to generate essentially the same incompleteness problems using occurrences of descriptions that were uncontroversially *non-referential*, or occurrences of other noun

⁵⁵ I add ‘occurrence’ only to highlight the fact that a description δ should not be regarded as (in)complete with respect to the whole utterance in which it occurs but with respect a particular *occurrence* within the sentence uttered. In principle, one utterance-occurrence of δ may be complete and another occurrence incomplete in the very same utterance. Convincing cases are, perhaps, not that easy to produce, but it would be rash not to guard against the possibility of such cases.

⁵⁶ As I said in ch 5, when summarizing part of ch 3, ‘even the descriptive content of an *overt* description is not always fixed by purely linguistic factors’ (p. 201). The word ‘even’ appears here and the word ‘overt’ was italicized because I was, in that passage, commenting first on the fact the descriptive content of descriptive (i.e. D-type) pronouns is not always determined by purely linguistic factors.

phrases that were uncontroversially assigned quantificational interpretations.⁵⁷ Rather than committing resources to a futile ideological debate—for that is what I thought it had become—I thought we should try to understand the general problem of what I called *quantifier incompleteness*: how is it that we can legitimately use, say, ‘every horse’, ‘no horse’, ‘the horses’, or ‘the horse’ without being understood as making claims about every horse in existence, or use ‘the horse’ without being understood as claiming (in part) that there exists only one horse? On the assumption that (much of) the semantic power of (many) noun phrases in natural language can be understood in terms of restricted quantifiers of the form $[Dx: \phi(x)]$, where D is a quantificational determiner, the problem of matrix interpretation is to explain how it is that $\phi(x)$ is often understood, in context, as true of fewer objects than it is true of when taken at face-value.⁵⁸ Let us now examine this matter in more detail.

6. LF and Aphonicity

Following Chomsky, it is common in much work in syntax to assume some sort of distinction between a sentence’s superficial and underlying form. The details may have changed over the decades, labels may have come and gone, and various theory-internal commitments may have changed or evolved, but the idea of some sort of distinction between superficial and underlying form is still with us, except that today it is, in many ways, more natural than ever in so far as it is connected more transparently to sound and meaning. With Chomsky (1995), let us identify a sentence with a pair $\langle \pi, \lambda \rangle$ where π is a PF (‘Phonetic Form’) to be read by the sound system, and λ an LF (‘Logical Form’) to be read by the intentional system). As Chomsky has stressed for a quarter of a century, LFs are not full-fledged *intentional* representations, as rich in content as the those involved in beliefs, intentions or expectations, objects with truth conditions (or the analogues thereof), inferential rôles and so on. LFs are simply *the grammar’s contribution* to the generation of such representations by the intentional system—which receives inputs from

⁵⁷ That incomplete descriptions may be used attributively was noted by Donnellan (1968). That any adequate account of incomplete descriptions must handle such uses as well as referential uses is stressed by Peacocke (1976), Evans (1982), Soames (1986), Bach (1987) and by me *Descriptions*.

⁵⁸ I said that quantificational noun phrases *can be understood in terms of* restricted quantifiers of the form $[Dx: \phi(x)]$, not that the noun phrases *are, or are represented as*, such quantifiers. Strictly speaking the English quantificational noun phrases ‘every man’ and ‘the man’ are not even of the same syntactic category as the restricted quantifiers $[every\ x: man\ x]$ and $[the\ x: man\ x]$: the former combine with a verb phrase to form a sentence, the latter combine with a formula ψ to form a formula ζ , and ψ must contain an overt variable x if ζ is to count as an adequate representation of a sentence of English. To avoid wordiness, I shall henceforth pass over this detail and allow a certain amount of intelligible and readily corrigible loose talk. It will be useful on many occasion to move back and forth between English descriptions such as ‘the man’ and *RQ* ‘descriptions’ such as $[the\ x: man\ x]$, and on many such occasions I will use ‘quantificational noun phrase’ and ‘restricted quantifier’ interchangeably, despite the strict inaccuracy of such talk. I shall also move back and forth between sentences of English and formulae of the formal language. No confusion should arise.

various cognitive faculties. The basic idea behind the concept of LF representations has remained robust since its inception: An LF incorporates, ‘whatever features of sentence structure (1) enter into the semantic interpretation of sentences and (2) are strictly determined by properties of sentence grammar’ (Chomsky, 1976b: 305). The only difference today is what is meant by ‘strictly determined by properties of sentence grammar’. With the emergence of the minimalist outlook, this phrase may be usefully understood as ‘strictly determined by the exigencies of connecting sound and meaning.’

Running through Chomsky’s (1995, 2002) recent work on syntax is an argument from ‘virtual conceptual necessity’: complexity and stipulation are to be avoided as, all else being equal, language will employ only those devices needed to link sound and meaning. On the assumption that there is a component of the mind/brain dedicated to language, the human language faculty, one consequence of Chomsky’s ‘minimalist’ outlook is that all properties of sentences relevant to sound and meaning should be derivable from quite general considerations about the way the language faculty must engage with two other cognitive systems, one dealing with the articulation of sounds and their perception (henceforth *the sound system*), the other trading in intentional/conceptual representations (henceforth *the intentional system*). A particular language can be seen as an instantiation of the language faculty (with certain options specified), something that can provide ‘instructions’ to be interpreted by the sound system, on the one hand, and the intentional system, on the other. More specifically, a language is a computational system that generates pairs $\langle \pi, \lambda \rangle$ of representations.

If LFs are not full-fledged *intentional* representations, what are they? They are, I said, the grammar’s contribution to the generation of such representations by the intentional system. But what does this amount to? To get the flavour it is helpful to think of the intentional system as receiving inputs from various cognitive faculties and trading in representations something like the sentences of the ‘language of thought’ in Fodor’s (1975, 1983) sense, a modality-neutral symbolic system of representation in which thought takes place and into whose sentences utterances of natural language sentences must be mapped if understanding is to take place. LFs exhaust the grammar’s contributions to this system; an LF is not an interpreted object with an intentional content: it is simply a syntactic representation, determined by sentence grammar, of those features of grammatical structure that enter into the interpretation of utterances of that sentence (for example, relations of scope and binding). Not only do LFs fail, for example, to specify references for referentially independent occurrences of pronouns, they fall short of being full-blown intentional representations in all sorts of other ways.⁵⁹

⁵⁹ See Chomsky (1976a, 1976b, 1986, 1995, 2002), Sperber and Wilson (1986, 1995), and Carston (2002). In a Davidsonian spirit, Higginbotham and May (1981), Higginbotham (1983a, 1983b, 1985), Larson and Segal (1995), Ludlow (1989, 2002), Neale (1994) and others have treated LFs as objects which (relative to assignments of values to referential elements, some of which may be of an indexical nature) have recursively specifiable truth conditions, an idea Chomsky rejects. There is, I think, *something* of value in the alternative conception of LF that needs to be super-imposed upon the official Chomskyan conception to give it bite, but it is difficult to make it precise. In trying to effect the superimposition, my own earlier discussions of LF

Reference to the theory of phrase structure is virtually eliminated in the minimalist framework, there being no phrase structure rules in the traditional sense, and strictly speaking there is no constituent structure to a PF representation. So let us borrow the old slash notation from phonology when talking about PFs, but with standard orthography rather than phonological symbols inside the slashes:

- (1) /John said he was at the bank/
 (2) /every man loves his mother/

(This notational convention allows us to talk of /bank/, /he/, and /his/ without prejudging the issue of ambiguity.) We can still allow ourselves the convenience of using a phrase structure trees or a labelled bracketing to explicate superficial structure, so to speak, even if what we write down is strictly an unholy hodgepodge of PF and LF. Thus we might use (3') to explicate the PF (3):

- (3) /the king snores/
 (3')
- | | | | |
|--|-----|------|---------|
| | | S | |
| | DP | | VP |
| | D | NP | V |
| | | N | |
| | the | king | snores. |
- [_S[_{DP}[_Dthe] [_{NP}[_Nking]]] [_{VP}[_Vsnores]]]

evinced deeply worrying ambiguities. In 'Events and LF' (1988) and *Descriptions* (1990) the discussion is very Chomskyan: I treat LFs as no more than syntactic objects encoding those aspects of syntax relevant to interpretation, and I am careful to distinguish LFs themselves from the formulae in a system of restricted quantification I use to represent the truth conditions of utterances of sentences, formulae that could depart significantly in structure from the LFs of the sentences uttered (for example, in the discussions of perceptual reports, incomplete descriptions, and descriptive pronouns). Various people tried to convince me I needed to embrace, or at least move closer to the truth-evaluable conception of LFs, but the interpretation of *utterances* containing incomplete descriptions and other underspecified DPs, as well as the interpretation of D-type pronouns held me back as it seemed, and still seems, preposterous to treat LFs as containing all sorts of unrecoverable *predicative* devices not present in surface syntax. At the same time, it has always seemed to me that if the Chomskyan LF of a sentence *X* is meant to lay bare the contribution made to the process of *interpretation* by *X*'s syntax, it must provide genuine constraints on what someone uttering *X* can be *saying*, something for which the alternative truth-conditional conception is tailor-made. In 'Logical Form and LF' (1994), I attempted a reconciliation of sorts—originally with Larson—by abstracting as much as possible from *pragmatically* determined aspects of the truth-conditions of utterances. The attempt was ultimately unsuccessful, I think, because I strayed so far from the Chomskyan conception of LF. The truth-evaluable notion of LF appears to have a simpler time with inference. A classic case of 'logical form' revealing inferentially vital elements is Davidson's (1967) account of action sentences: 'John left quickly' is meant to entail 'John left' because each is an existential quantification over events and first-order logic guarantees the validity of the following: (i) $(\exists x)(\text{leave}(\text{John}, x) \cdot \text{quick}(x))$; therefore (ii) $(\exists x)\text{leave}(\text{John}, x)$. It is in this spirit that Higginbotham (1983b, 1985) posits variables corresponding to those Davidson sees in the 'logical forms' of English sentences in their LFs, again assuming that LFs are ripe with truth conditions (relative to assignments). On the assumption that the LF of 'Mary saw John leave' also contains an event variable corresponding to the one in the premise of the inference below, Higginbotham is able to capture further inferences by appealing to LFs: (i) $(\exists x)(\text{leave}(\text{John}, x) \cdot \text{saw}(\text{Mary}, x))$; therefore (ii) $(\exists x)\text{leave}(\text{John}, x)$.

Following Abney (1987), today it is common to call what used to be called an NP (noun phrase) a DP (determiner phrase) today to respect the idea that its head is the D (determiner) not the N (noun). On this usage, NP is the label for the nominal expression, simple or complex, with which a D merges to form a DP:

- (4) $[_{DP}[_{D}the] [_{NP}[_{N}king]]]$
 (5) $[_{D}[_{D}the] [_{NP}[_{A}French] [_{N}king]]]$.

The LF corresponding to the PF (3) will look something like (6):

- (6)
- | | | | | |
|--|--|-----------------|--------------------------------|---------|
| | | S | | |
| | | DP ¹ | | S |
| | D | NP | DP | VP |
| | | N | | V |
| | the | king | x_1 | snores. |
| | $[_{S}[_{DP}[_{D}the] [_{NP}[_{N}king]]}]^1$ | | $[_{S}x_1[_{VP}[_{V}snores]]]$ | |

In (6) the quantifier expression ‘the king’ has been extracted—indeed forced by general principles of morphosyntax—from its original position (discernible in the hodgepodge (3')) and merged with the original S node to form another (it has been ‘Chomsky-adjoined’ to S). From this ‘new’ position it binds the variable x that now occupies its ‘original’ position, that position now being within its scope. ‘the king’ and x are *co-indexed*: the numerical *subscript* on x is an *index* indicating that it is to be read as bound by the quantifier expression ‘the king’, which bears the same index as *superscript*.⁶⁰ This talk of variables and binding amounts to a description of an important part of the *interpretive* information carried by (6), precisely the sort of thing Chomsky has always ascribed to LFs.⁶¹

One question that will have to be addressed is whether quantifier expressions are the only DPs that are raised at LF or whether the phenomenon is fully general, involving

⁶⁰ Using only subscripts to co-index would obscure the fact that binding is an asymmetric relation, unlike, say, co-reference; this will be important later. I am deliberately simplifying here. For example, I have ignored the fact that the superscripted index on the DP ‘the king’ has been projected upwards from the index on the D ‘the’.

⁶¹ Variables in syntactic theory are syntactic objects, but as Heim and Kratzer (1998) stress they are not merely syntactic objects. The concept that variable-binding attempts to formalize is an interpretive one. Variables are expressions *interpreted* in a certain way, expressions whose values *vary* with something or other, as the etymology suggests. (In order to generalize the notion of variable-binding, we will however want to include the case of zero variation, for example if names are treated as variable-binders. See below.)

names, pronouns, and possessives for example (DPs whose structures we will look at in a moment):⁶²

- (7) $[_S[_{DP} Anne]^1[_S[_{DP} the\ king]^2[_S x_1[_{VP} respects\ x_2]]]]]$
 (8) $[_S[_{DP} she]^1[_S[_{DP} the\ king]^2[_S x_1[_{VP} respects\ x_2]]]]]$
 (9) $[_S[_{DP} Marga]^1[_S[_{DP} her_1\ mayor]^2[_S x_1[_{VP} respects\ x_2]]]]]$.

My own suspicion is that raising will have to be fully general, and I shall give my reasons later. (To jump ahead: notice ‘her’ in (9) is co-indexed with ‘Marga’, and notice the order of the raised DPs.)

Consider (10) and (11):

- (10) John promised Ann to sing
 (11) John asked Ann to sing.

Traditional grammars talk about the ‘understood subject’ of the verb ‘sing’ in such examples: in (10) it is ‘John’, in (11) it is ‘Ann’. In syntactic theory this is captured by a difference in the status of an aphonic element occupying the subject position of the embedded infinitival clause:⁶³

- (10') $[_S John^1\ promised\ Ann\ [_S e_1\ to\ sing]]]$
 (19') $[_S John\ asked\ Ann^2\ [_S e_2\ to\ sing]]]$.

In (18'), the interpretation of *e* is required (by the syntax and the meaning of the verb ‘promise’) to proceed by way of the interpretation of the subject expression ‘John’ (‘promise’ is a *subject-control* verb). In (19'), by contrast, the interpretation of *e* is required (by the syntax and semantics of the verb ‘ask’) to proceed via the interpretation of the object expression ‘Ann’ (‘ask’ is an *object-control* verb (in this construction)).

It is worth pointing out a syntactic difference between the aphonic DPs in (17) and (18), on the one hand, and those in (6)-(9), on the other: in the terminology of older syntactic theory, those in (6)-(9) are *movement-generated* (assuming a ‘raising’ operation), whereas those in (17) and (18) are *base-generated* (given the lexical entries for ‘promise’ and ‘ask’). For present purposes the important syntactic differences between types of aphonics can be put aside.

A word of caution. As in *Descriptions*, (i) my use in the present chapter of formulae of a semi-formal system of restricted quantification to capture the truth conditions of what a speaker says by uttering a sentence *X* in a way that *respects X*’s semantically relevant structure, and (ii) my appeal to a Chomskyan picture of syntax according to which *X* may be factored into a surface form (its PF in today’s lingo) and its LF, where the latter *reveals X*’s semantically relevant structure and may contain elements no counterparts of which appear in *X*’s surface form, does not mean the formulae of the semi-formal

⁶² If, as seems plausible, possessives are definite descriptions, and definite descriptions are quantifier expressions, then we already have our answer in one case.

⁶³ For simplicity I have not raised the names here.

language *are* LFs. They most certainly are *not* LFs, containing as they do all sorts of constituents no counterparts of which appear in LFs, most noticeably where incomplete descriptions are concerned. Like the notation of structured propositions, they are methodologically useful precisely because they enable us to characterize the failure of isomorphism between form and content and to do so using a notation that nonetheless respects the view that what is said bears some systematic relation to syntax and word meaning.

7. ‘Implicit’ and ‘Explicit’

In *Descriptions* I alluded to two common responses philosophers and logicians have made over the past fifty or sixty years to the problem of incomplete descriptions (‘two of the more popular approaches’ as Reimer (1998a: 96) quite rightly puts it).⁶⁴ And I suggested, perhaps rashly that attempts to spell them out in satisfactory ways might render them notational variants (when ‘all is said and done’, as I put it). I also said that *both* responses might be necessary in any final understanding of incompleteness. Assuming a Russellian account of descriptions, the problem of incompleteness is, in its most general terms, the problem of explaining how it is that an occurrence of a matrix $\phi(x)$ occurring in a quantified expression [*the* x : $\phi(x)$] is understood, in context, as true of fewer objects than its superficial form seems to require. Putting the problem this way allows us to see it, as I claimed we should see it, as an instance of a more general problem that affects the use of quantified noun phrases.

When this problem arises, we are faced with a case in which we seem to have slippage between language and the world. And there are only two things we can do about that: tinker with language, or tinker with the world. When we tinker with *language*: we do something about the *matrix* $\phi(x)$. When we tinker with the *world*, we do something about the *objects* that (potentially) satisfy the matrix. If we tinker with $\phi(x)$ let us say that we are adopting an ‘explicit’ approach to the problem; if we tinker with the objects let us call it an ‘implicit’ approach. (This choice of terminology will become clear.)

The distinction between the world-tinkering, implicit approach and the language-tinkering, explicit approach corresponds to a difference in focus and in attitude with respect to different parts of a DP (this is perhaps easier to see under the DP hypothesis):

- (1) $[_{DP}[_D \text{the}][_{{}_{NP}}[_N \text{table}]]]$
- (2) $[_D[_D \text{the}][_{{}_{NP}}[_A \text{red}][_{{}_{N}} \text{table}]]]$.

A personal pronoun can be viewed as a D that effectively serves as a full DP by virtue of merging (combining) with an aphonic NP, one that has no phonological matrix and so

⁶⁴ Reimer’s (1998a) paper is an illuminating discussion of the two approaches.

does not appear, so to speak, at PF.⁶⁵ For example, ‘he’ would have the following structure, where *e* is aphonetic:

- (3) $[_{DP}[_{D}he] [_{NP}e]]$.⁶⁶

We find Ds occurring with aphonetic NPs elsewhere:

- (4) $[_{DP}[_{D}one] [_{NP}e]]$ is broken
 (5) $[_{DP}[_{D}this] [_{NP}e]]$ is hot
 (6) $[_{DP}[_{D}his] [_{NP}e]]$ is broken⁶⁷

⁶⁵ For discussion, see Abney (1987), Cardinaletti (1994), Cardinaletti and Starke (1999), Chomsky (1995), Szabolcsi (1994). The idea that pronouns might be determiners appears to originate with Postal (1969). In the elementary exposition of DPs that follows, I simplify dramatically as on-going debates in linguistics about the details are not crucial to the philosophical issues of concern here.

⁶⁶ The postulation of expressions that are aphonetic despite having syntactic rôles and semantic properties is surely no more or less problematic than the postulation of expressions that are semantically empty despite having syntactic rôles and phonological properties (‘it’ in ‘it’s raining’, for example). The idea of an expression that is phonetically *and* semantically empty is harder to get one’s mind around, and on the interpretation of Chomsky’s present framework I endorse— syntax is whatever it is that relates PF and LF— the possibility of such an expression is straightforwardly excluded. The discovery or postulation of *any* expression constitutes a contribution to syntax in the first instance, but its existence is justified only if it is doing something at LF or PF. Consequently, the discovery or postulation of an aphonetic expression must be justified by its rôle at LF, and to this extent, it will contribute in one way to the project of producing a theory of utterance interpretation. (The general point should not be exaggerated, however. The discovery or postulation of an aphonetic, *indexical* expression, one every bit as flexible in its interpretation as the overt expressions ‘this’ or ‘that’ or ‘he’ (when used to make independent reference) does not fix interpretation: the interpretation of an utterance of a sentence containing an occurrence of an aphonetic indexical is always going to be a full-fledged *pragmatic*, i.e. inferential matter, the semantics of the aphonetic expression itself placing only non-deterministic constraints on interpretation.) So how does $[_{NP}e]$ contribute to LF? What is its semantic rôle? The answer I explore in Neale (forthcoming a) is that it is interpreted as a formula $x_k = x_j$ ($k \neq j$). the idea is that the DP $[_{DP}[_{D}he] [_{NP}e]]$ is interpreted as $[he\ x_k: x_k = x_j]$, assuming an axiom for *he* that is a trivial modification of the Russellian axiom for *the*, as Postal’s hypothesis would anyway suggest. For a brief sketch, see the end of Section 24 of the present chapter.

⁶⁷ In English the distinction between genitive and possessive determiners is unmarked, at least phonologically. This is not universal. In Icelandic, for example, third person genitive and possessive pronouns are quite distinct, lexically and also in respect of inflectional morphology. Thus (i) is translated as (ii) or (iii) according as ‘his’ is bound by ‘John’ [‘every man’]: (i) John [every man] loves his wife; (ii) *Jón* [sérhver maður] elskar konuna sína; (iii) *Jón* [sérhver maður] elskar konuna hans. In (ii), *sína* is a reflexive possessive (or possessive reflexive) that has to be understood as bound by (and hence within the scope of, i.e. c-commanded by) the subject expression. (There are curious exceptions in subjunctive sentences involving logophoricity. For discussion, see Reuland (2001), Neale (forthcoming a) and references therein.) It occurs in the feminine, accusative, and singular to agree with *konuna*, the noun it qualifies. In (iii) *hans* is the simple genitive, which (unlike *sína*) occurs in the masculine and enters into no agreement relations whatsoever with *konuna*. The Icelandic definite article typically takes the form of a suffix added to the noun (very likely this originated in a free-standing definite article which is rarely encountered in ordinary talk today). Without the suffix the noun is typically understood as indefinite. So, for example, in the nominative (used primarily for the subject of a verb) the feminine noun *kona* (‘woman’ or ‘wife’) becomes *konan* (‘the woman’ or ‘the wife’) when definite. In the accusative (used primarily for the direct object of a verb but also with some prepositions) *konu* becomes *konuna* when made definite, as in the example above. Apart from certain standardized exceptions, the counterparts of English possessive descriptions are formed using the definite rather than the indefinite form of the noun. There are two main types of exception. The first is where the noun to which the possessive is attached expresses a close family relation, with the notable exception of *kona* (which can translate ‘woman’ as well as ‘wife’). The second is where the genitive is used with proper names to indicate possession. The accusative form of ‘John’s wife’ is *konu Jóns*: so unlike cases involving

- (7) $[_{DP}[_{D}\text{many}][_{NP}e]]$ applied but $[_{DP}[_{D}\text{few}][_{NP}e]]$ were chosen.

There are differences between (4)-(7) that need not detain us, but in each case the postulated aphonic $[_{NP}e]$ may be replaced by a phonic NP like ‘plate.’ (With pronouns this can be done, but it feels epenthetic: ‘I, Stephen’, ‘you, dear reader’, ‘he, Chomsky.’ I suppose any decent theory positing (10) as the structure of ‘he’ should explain this.)

By contrast with pronouns, proper names can be viewed as Ns effectively serving as full DPs by virtue of merging (in English) with aphonic Ds:⁶⁸

- (8) $[_{D}[_{D}e]][_{NP}\text{John}]$.

Perhaps we find NPs with aphonic Ds elsewhere:

- (9) $[_{DP}[_{D}e]][_{NP}\text{whales}]$ are $[_{DP}[_{D}e]][_{NP}\text{mammals}]$.

This picture of DPs helps frame the two traditional ways of thinking about incomplete descriptions. If we tinker with *language*, we end up doing something about the *matrix* $\phi(x)$. Let us call this the ‘explicit’ response to the incompleteness problem. If we tinker with the *world*, we end up doing something about the *objects* that satisfy the matrix. Let us call this the ‘implicit’ response. Consider

- (10)
- | | | |
|--|-----|-------|
| | DP | |
| | D | NP |
| | the | table |

The D ‘the’ is the head of the DP ‘the table’. The implicit response purports to explain how we get away with using incomplete DPs by focusing on how the head node, D, or its projection, DP, is to be interpreted, both of these nodes corresponding to a quantificational expression. On the assumption that the quantificational structure of a sentence ‘the table is ψ ’ is represented reasonably using

- (11) $[the\ x: table(x)]\ \psi(x)$

there are two quantifiers to look at, the unrestricted quantifier *the x*, corresponding to the D node, and the restricted quantifier $[the\ x: table\ x]$, corresponding to the DP node. The implicit, world-tinkering, approach explains incomplete usage by limiting the number of objects satisfying *table(x)*, and the fact that there are two quantifiers to consider means there are (at least) *two* ways of effecting the required delimitation. Let *r* be some subset of the objects in the domain of quantification (‘*r*’ for ‘restricted’). In our formal language, *RQ*, we can represent the two ways of delimiting the domain as follows, subscripting ‘*r*’ onto the quantifier whose domain is to be delimited:

possessive pronouns (*konuna sína*) and cases involving genitive pronouns (*konuna hans*), in cases involving genitive names (*konu Jóns*), the suffix for the definite article is *not* added to the noun.

⁶⁸ In many languages (e.g. ancient and modern Greek), names appear regularly with the definite article. In certain contexts names may occur with the definite article and other determiners in English: ‘That John Smith is not the John Smith I was talking about.’

- (12) $[the\ x_r: table\ x]\psi(x)$ ‘the x (in r) such that x is a table’
 (13) $[the\ x: table\ x]_r\psi(x)$ ‘the x such that x is a table, (in r).’

I prefer a simple ‘ r ’ rather than the variable-containing ‘ $x \in r$ ’ to avoid the suggestion that in (12) and (13) the descriptive content itself is modified, as it is in, say, (14):

- (14) $[the\ x: table(x) \bullet x \in r]$ ‘the x such that x is a table and $x \in r$ ’.

For as I said in *Descriptions*, one of the central tenets of the implicit approach is that it ‘leaves the descriptive content untouched’ (1990: 95). In (12), the unrestricted quantifier *the x* ranges over the things in r ; in (13) the restricted quantifier $[the\ x: table\ (x)]$ ranges over the things in r . If r is a proper subset of the original domain, then the felicitous use of an incomplete description can, in principle, be explained. If A utters the sentence ‘the table is ψ ’, if r contains only objects within two metres of where A is sitting, and if there is exactly one thing in r that is a table, and exactly one thing that is a table in r , (12) and (13) will each present exactly one thing, and a table at that, to examine as a potential satisfier of $\psi(x)$. A roaring success, apparently, for both accounts of the domain limitation. It is not immediately obvious that either of (12) or (13) has an advantage over the other; but we cannot rule out the possibility that we might come across phenomena that bring out important differences, so let us keep both on the table, at least for now.

No *syntactic* thesis is implied by the implicit approach. It is compatible with the thesis that the English DP ‘the table’ contains an aphonic expression corresponding roughly to the subscript r in our formal language, but it is also compatible with the syntactic thesis that *what you hear is what you get*. That is, the postulation of an aphonic in ‘the table’ is no part of the implicit approach itself, it is, rather, a particular *syntactic* proposal for *implementing* it, one that might be motivated or rejected on (presumably) syntactic grounds.

An interesting version of the implicit response was produced to Barwise and Perry (1983), in connection with incomplete descriptions used non-referentially. Drawing upon Barwise and Cooper’s work on generalized quantifiers, Barwise and Perry introduce the notion of *persistence* as a property certain statements have. Suppose statements are evaluated not with respect to the whole world, they suggest, but with respect to specified *parts* of the world, ‘situations’ as they call them. To simplify matters in a way that does not bear on present concerns in any threatening way, let us prescind from time: if we consider the world at a time we can think of situations as just spatial parts of the world (the maximal situation). So, for example, Britain is *part of* the world. Now consider the following sentence:

- (15) the richest person lives in London.

The idea Barwise and Perry want us to entertain is that what is expressed by an utterance of (15) can be *true* with respect to Britain (because the richest person *in Britain* lives in London), and at the same time *false* with respect the world at large (because there is at least one person *outside* Britain who is richer than the richest person *inside*). Now any

particular utterance of (15) is uttered *in* a situation *S* and intended to be evaluated *at* some particular situation *S'*, which need not be identical to *S* and which may be smaller than the entire world (the maximal situation). And it is this fact, or so it is suggested, that explains how utterances of (15) work (for now, let us put aside so-called referential uses of descriptions).

The phenomenon can be inverted by considering utterances of (16):

(16) some princes are ministers.

What is expressed by an utterance of (16) can be *false* with respect to Britain, but at the same time *true* with respect to the world at large (because, say, some Belgian princes are ministers).

On Barwise and Perry's account, the idea is *not* that someone uttering (15) expresses different propositions according as the utterance is to be evaluated with respect to Britain or the world at large—the proposition that the richest man *in Britain* lives in London and the proposition that the richest man *in the world* lives in London, for example. Rather, the idea is that the proposition expressed is the *same*, but is true with respect to Britain and false with respect to the larger world. *Mutatis mutandis* for utterances of (16). In Barwise and Perry's terminology, what is expressed by an utterance of (16) is *persistent* in (roughly) the following sense: if it is true with respect to a situation *S* (for example, Belgium), then it is true with respect to every more encompassing situation *S'* (for example, Europe, or the northern hemisphere, or the whole world). By contrast, what is expressed by an utterance of (15) is *non-persistent*. It can be true at a situation *S* (for example, Britain) whilst being false at some more encompassing situation *S'* (for example, Europe, or the northern hemisphere, or the whole world).

The persistence of (16) and the non-persistence of (15) are traceable to the determiners each contains: 'some' is persistent (in a sense to be defined rigorously) whilst 'the' is not. Various problems with this and other versions of the implicit approach—indeed for the general idea that our utterances are evaluated with respect to less than the entire world—were raised almost immediately Westerståhl (1985) and Soames (1986) (think about 'the dog bit another dog') Before looking at them. Let us turn to the explicit response.

The explicit response leaves the world alone. It purports to explain the felicitous use of 'the table' by looking at the *other* node in (1), the one not occupied by a quantifier but by the nominal, the NP 'table'. The basic idea is explicitly *modal*: the nominal is often shorthand for, elliptical for, an abbreviation of at least one richer nominal the speaker *could have* used and *could* produce if asked to be more explicit. (Hence the name.) Consider the following dialogue:

A: The table is scratched

B: Which table?

A: The table I bought this morning (*Or*: The one I bought this morning).

According to the explicit approach, this type of dialogue is suggestive of what is going on when we make felicitous uses of incomplete descriptions. *B* is intended to interpret *A*'s

utterance of ‘the table’ as if it were an utterance of ‘the table I bought this morning.’ There need not be a unique description that *A* can supply, but there had better be at least one—and one that *B* could reasonably have been expected to construct at that—if the speech act is to be felicitous.⁶⁹ This is vague, of course. But that’s fine: it’s just a general description of an *approach* to completion that many philosophers think should be explored, together with its *target*, one that cannot be made more precise without looking at the *mechanisms*, a general theory of utterance interpretation that explains how hearers integrate linguistic and non-linguistic information to arrive at interpretations, a theory of the pragmatic, inferential processes involved in utterance interpretation.

So on the explicit response, someone producing an utterance of an incomplete DP (e.g. ‘every citizen’ or ‘the president’) is understood as expressing what he *would have* expressed more explicitly had he uttered a richer (‘complete’) DP (e.g. ‘every U.S. citizen’, ‘the U.S. president’), a DP he *could have* used in place of the incomplete one. The explicit strategy is sometimes called the *ellipsis* strategy in the literature, presumably in deference to the suggestions made by Quine (1940) and Sellars (1954), who talk, respectively, of *elliptical uses* and *elliptical utterances* of descriptions. In consequence, I shall use ‘explicit response’ and ‘ellipsis response’ interchangeably. A word of warning: I shall sometimes talk about ‘completing an incomplete description’, but this is just a bit of shorthand for ‘coming up with a richer description that does a good job in capturing the speaker’s intention’.

Just as there is no *syntactic* thesis implied the implicit response, so none is implied by the explicit response. There is certainly no implication, for example, that expressions are transformationally deleted between levels of grammatical representation in a Chomskyan grammar—indeed, on standard assumptions there *could not be* such a syntactic thesis because such deletions would violate the principle of *recoverability*, which requires deleted elements to be recoverable from linguistic context (see below).⁷⁰ Like the implicit approach, the explicit approach is meant only to describe how speakers intend their utterances of incomplete descriptions to be *interpreted* on particular occasions and to describe the interpretations hearers do seem to get. Obviously it involves no cognitive claim about the mechanisms whereby hearers manage to *come up with* particular interpretations on particular occasions: that is something that a theory of the pragmatic, inferential processes involved in utterance interpretation to explain.

⁶⁹ W. Blackburn (1988) stresses an important point about the explicit approach that I did not appreciate until I had virtually finished writing *Descriptions* (it was Blackburn’s article that made me appreciate it). In the context of its utterance, an incomplete description is typically understood with the force of any of a *batch* of richer descriptions the speaker *could have* used to make his point, plenty of which the speaker *could* produce if asked to be ‘more explicit’.

⁷⁰ Oddly, Stanley and Szabó (2000a) read just such a syntactic thesis into my summary of the explicit response in *Descriptions*, and by implication into the explicit responses of Quine (1940), Sellars (1954), Davies (1981), and Evans (1982) amongst others. For detailed discussion, see below.

Using the notation given above, the basic idea—which can be developed in a number of ways—is that sometimes the matrix $\phi(x)$ of a quantified DP is understood, in context, *as if* it were a richer matrix $\phi(x, a)$ containing an additional argument or a conjunction $\phi(x) \bullet \zeta(x)$ which the speaker could have readily produced. The meanings of some nouns make it clear an additional argument is called for: ‘the murderer (of x)’, ‘the king (of x)’. As pragmatists sometimes put it, interpretation requires ‘saturation’ of an ‘implicit argument’. The meanings of others clearly don’t invite saturation: ‘the table’, ‘the man’. Here pragmatists sometimes talk of ‘enrichment’, which is constrained only by the exigencies of the interpretation process.⁷¹ Talk of saturation and enrichment is not itself meant to constitute a *theory* in any interesting sense: the only *theory* involved is a general *pragmatic* theory, a theory of the cognitive processes involved in utterance interpretation, a theory that explains how hearers integrate linguistic and non-linguistic information in interpreting one another, a theory that explains not only how we interpret utterances of incomplete descriptions but also how we assign reference to names and pronouns, establish binding (where syntax falls short), and resolve potential ambiguities, and how we identify and interpret utterances replete with irony, metaphor, elision, anacoluthon, aposiopesis, and on top of all of this how we identify what a speaker is *implying* as well as saying. And these cognitive processes must be appealed to by *any* account of incompleteness, whether it involves specifying some richer description, specifying some suitable background domain restriction, or specifying some suitable value for some hitherto unappreciated, aphonic domain variable cohabiting a syntactic node with a nominal in the manner of Stanley and Szabó (2000a) and Stanley (2000, 2002a,b). However you look at it, it’s magic, and it betrays a misunderstanding of the issues to complain, that on the explicit approach the hearer performs an act of magic (the recovery of a magical ellipsis) no counterpart of which the hearer performs on an approach that requires the hearer to supply properties or sets or whatever as values for aphonic domain variables.⁷² Similarly, it will not do to claim, with Stanley (2002a: 158, n 12) that the

⁷¹ For discussion of the literature on saturation and enrichment see Carston (2002).

⁷² A hearer who had to assign the sorts of values to aphonic domain variables required by the theory Stanley (2000, 2002a,b) and Stanley and Szabó (2000a) propose would seem to be in a worse position than one who had to come up with richer descriptions. Linguistic pragmatism is quite happy to posit aphonic variables where syntactic theory requires them but does not see them as a universal panacea to problems of interpretation. A theory that posits the existence of aphonic domain variables in syntax is primarily a *syntactic* proposal concerning the LF of a sentence that may be uttered on different occasions to say different things, and it should not be confused with a theory that explains *how hearers assign values to these variables*. Interpretation of any postulated context-sensitive expression on a given occasion of utterance is itself a *pragmatic*, *richly inferential* matter, the product of integrating linguistic and non-linguistic information, something that is done by a pragmatic theory. As far as interpretation of incomplete matrices is concerned, the only substantive difference between the pragmatist and someone who postulates an aphonic element co-occurring with a nominal is that the latter insists that the search for and integration of contextual information in the interpretation process is triggered syntactically. I know of no good argument that an item in syntax is *necessary* for such a search and/or for such integration to take place—such an argument would have to come from psychology. Merely pointing to the well-known phenomenon of ‘implicit binding’ certainly does not *demonstrate* the existence of aphonic variables, as has been recognized since at least Evans’s (1977) account

explicit approach ‘simply amounts to a re-description of the phenomenon to be explained, rather than an account of it’ *if* the implication is that this is less of an account than positing a domain variable that takes on whatever value is required to make things work out correctly. Neither the explicit nor the implicit approach, nor Stanley and Szabó’s contextual variable approach constitutes a *theory* in any sense relevant to interpretation. Whichever way we go here, all of the work is done by pragmatic inference.

It is difficult to imagine anyone sympathetic to the explicit response ever viewing it as subject to the following strange and quite *ad hoc* constraint: any descriptions occur in a single sentence they must be completed in precisely the same way. At the end of boxing match between a Russian and a Swede I might say to you, upon hearing that the panel of eleven international judges has declared the Swede the winner by ten votes to one, ‘I know why it wasn’t unanimous.’ ‘Why?’ you ask? I reply with (15):

(15) the Russian voted for the Russian’.⁷³

Obviously I would be saying that the Russian judge (in this contest) voted for the Russian boxer (in this contest). The point is almost too obvious to be worth stating, but probably I should have mentioned it in *Descriptions* and then laid out the consequences as I saw them for the implicit approach and its relation to the explicit approach.

Why did I suggest, in *Descriptions*, (i) that attempts to spell out the explicit and implicit responses in satisfactory ways may render them notational variants, and (ii) that a full understanding of incompleteness may require both responses? Isn’t there a tension between these suggestions? The two suggestions are, in fact, intimately connected. Here was my thinking.

First, as Soames (1986) showed in his discussion of Barwise and Perry’s version of the implicit approach, examples like the boxing match present an awful problem, for there is no situation that can be carved out of the world (the maximal situation), and no restricted domain of individuals in the world that can serve as a background against which my utterance of ‘the Russian voted for the Russian’ in the envisioned context comes out true. You need a situation or a domain with two distinct Russians in it (the judge and the boxer), but once you have that you are doomed to declaring my utterance false. In short, the implicit theory fails here because no situation or domain contains exactly one Russian and two distinct Russians.

Second, at the time I was writing *Descriptions* I could see only three things the implicit theorist could do at this point. (a) Accept that the implicit response needs help from the explicit response, and explore the idea that there is some sort of principled way to integrate the two responses as well as some sort of principled way to identify where one rather than the other is doing the lion’s share of the work. (A daunting task.) Hence the suggestion that both responses might be needed to fully understand incompleteness. (b)

of the interpretation of E-type pronouns, where implicit binding does not involve positing a phonetic variable in underlying syntax.

⁷³ This form of example is used by Soames (1986) and Westerståhl (1985) in a related context. See below.

Attempt to replicate the explicit theorist's success with the boxing match example by bringing in two distinct situations or restricted domains, one containing judges, the other boxers. The problem here, as I saw it, was explaining what it meant to say that my utterance could be true against two background situations or domains, one that came into view only in connection with the first occurrence of 'the Russian' in my utterance, the second only in connection with the second occurrence. And it was always possible to bring *prime ministers* into the picture. 'The prime ministers of Britain, France, and Russia have been trying to influence the judges,' I might say. And I might continue by uttering (16):

(16) the Russian told the Russian to vote for the Russian.

(The explicit approach has no problem here, of course. My utterance of (16) is understood as elliptical for an utterance of 'the Russian prime minister told the Russian judge to vote for the Russian boxer.') And then it was always possible to throw in a few non-extensional doodads and bound variables:

(17) the Dane didn't know the Swede suspected that the Russian had been ordered by his president to vote for the Russian.

I had no idea how to bring together *two* let alone three situations or restricted domains in the required way. I couldn't even see what the required way was! (c) Give up one of the two fundamental tenets of the implicit approach. Not domain restrictions themselves, but the idea that domain restrictions do not get into the proposition expressed. Hence the suggestion that when all is said and done the implicit and explicit responses might end up notational variants. Not because the *explicit* approach needed modifying—it always seemed fine to me—but because the only way of rescuing the *implicit* approach from obscurity might involve putting the (originally) background situations or restricted domains into the proposition expressed, as in (18), where p , j and b represent the restricted domains or situations in question:

(18) $[the\ x: Russian(x) \bullet x \in p] [the\ y: Russian(y) \bullet y \in j]$
 $[the\ z: Russian(z) \bullet z \in b]$ x told y to vote for z .

Quite what this would amount to when incomplete descriptions occurred within the scopes of various non-extensional operators, I could never get clear about. (The explicit approach is at least *clear* here.) But as far as I could make out, only by working appropriate restrictions into the proposition expressed could distinct descriptions occurring in the same sentence take advantage of distinct domains of quantification. And of course it involves abandoning one of the two characteristic features of the implicit approach: the idea of an *implicit* restriction that does not make itself into the proposition expressed. And the resulting proposal has the air of a pointlessly formal restatement of the explicit response. Replacing the idea that an occurrence of 'the Russian' is interpreted as if it were an occurrence of 'the Russian judge' by the idea that it is interpreted as 'the

Russian in the set of the things that are judges'—or, for that matter, 'the Russian having the property of being a judge' does not seem like great progress.

8. Pronouns and Incomplete Descriptions

Officially, I favoured neither the explicit response nor the implicit response over the other in *Descriptions*, but my preference for the explicit response is evident from its appearance throughout chapters 5 and 6 in the discussion of descriptive pronouns—those most incomplete of descriptions. After uttering (1),

- (1) A man walked over to our table

I might continue with any of (2)–(6), perhaps in descending order of likelihood:

- (2) He said nothing
(3) The man said nothing
(4) The man (in question) said nothing
(5) The man who had walked over said nothing
(6) The man who had walked over to our table said nothing.⁷⁴

Sometimes, it is essential to use a description to avoid ambiguity. Compare the following:

- (7) Smith bought an apple. It cost 50¢
(8) ? Smith bought an apple and a banana. It cost 50¢.

In (8) the pronoun 'it' must be replaced by 'the apple' (or 'the banana', 'the former', or 'the latter') to make one's meaning clear.

This suggests that an utterance of a pronoun anaphoric on, but outside the scope of (and hence not bound by) a quantified expression is understood as if it were an utterance of a description constructible from linguistic and conversational context. Indeed, this is the basic idea behind the postulation of a natural class of pronouns, which might be seen, on such a view, as conventional devices of ellipsis, viz. D-type pronouns.⁷⁵ The occurrences of 'he' in (2), 'the man' in (3), 'the man in question' in (4), and 'the man who had walked over' in (5) might all be seen, in the context in question, as elliptical for 'the man who had walked over to our table.'⁷⁶

⁷⁴ Similarly where we have plurals: 'Two men walked over to our table.' 'They said nothing.' 'The men said nothing,' etc. The idea that certain 'minimum' definite descriptions—e.g., 'the man', 'the woman', and 'the thing'—might function like the pronouns 'he', 'she' and 'it' is entertained by Quine (1960: 102-3; 112-3). I imagine he was not the first to point this out.

⁷⁵ See e.g. Cooper (1979), Davies (1981), Evans (1985), Ludlow and Neale (1991), and Neale (1990).

⁷⁶ The label 'D-type' seemed apt to Ludlow and me for those pronouns Evans claimed were E-types. First, Evans's E-type pronouns were rigid terms whose references were fixed by description; Ludlow and I denied the existence of E-type pronouns. (By virtue of having their references fixed rigidly by description, on Evans's account E-type pronouns are meant to be semantically equivalent to names introduced by description ('let us call whoever invented the zip fastener 'Julius') Second, Sommers (1982) had already suggested 'D-

Unsurprisingly, nothing precludes an utterance of D-type pronoun being understood as if it were an utterance of a description containing a name or a pronoun. Indeed, this fact seems to provide a plausible explanation of what is going on in cases of so-called donkey anaphora. For example, an utterance of (9) might be understood as if it were an utterance of (9'):

- (9) John bought only one donkey and Paul vaccinated it.
 (9') John bought only one donkey and Paul vaccinated the donkey John bought.⁷⁷

Consider the following progression:

- (10) every man who bought only one donkey
 had to pay for it with cash
 (10') every man who bought only one donkey
 had to pay for the donkey with cash
 (10'') [every man who bought only one donkey]¹
 had to pay for the donkey he₁ bought with cash.

Utterances of (10) and (10') may well be understood as elliptical for utterances of (10'').⁷⁸ The interesting point here is that the description for which the pronoun 'it' in (10) and the

type' for descriptive pronouns. Third, the idea of D-type pronouns (but not under this label) was logically prior to as well as sportier and more manoeuvrable than the idea of E-type pronouns. Pragmatic considerations often require D-type pronouns to be understood with large scope; but if we find occurrences that may be understood with small scope, this would seem to provide convincing evidence against the E-type analysis as well as a challenge to other theories of anaphora. In fact, scope ambiguities can be found in attitude, temporal and modal contexts. ('A man murdered Mrs Smith. The FBI think the police think he is from out of town.') See e.g. Davies (1981), Evans (1985), Ludlow and Neale (1991), and Neale (1990).

⁷⁷ Giving 'only one donkey' large scope quantifier and construing 'it' as a variable it binds yields the wrong result. Someone uttering (9) would not be saying that only one donkey satisfies *John bought x and Paul vaccinated x*, for that is compatible with John having bought *two* donkeys which the utterance of the original conjunction is not.

⁷⁸ Again, as Evans (1977) shows, giving 'only one donkey' large scope and construing 'it' as a variable it binds yields the wrong result (in this an countless other cases). Someone uttering (10) would not be saying that only one donkey satisfies *every man who bought x had to pay for x with cash*. The pronoun 'it' is not functioning as a bound variable in (10) because it does not lie within the scope of its purported antecedent 'only one donkey'; it is a D-type pronoun, and so is effectively an 'unbound anaphor'. At least that is the account I gave in descriptions, and it still seems to me the best account going. (See also Davies (1981) and Ludlow and Neale (1991).) In some languages the distinction between bound and unbound pronouns shows up lexically. (i) and (ii) are ambiguous in English between bound and donkey readings according as /his/ is bound by the subject DP or 'merely anaphoric' on the embedded DP:

- (i) [every man who has [a son]²]¹ loves his_{1/2} wife
 (ii) [only [John]²]¹ loves his_{1/2} wife.

In Icelandic, by contrast, /his/ translates as the reflexive possessive *sína* where binding by the subject DP is intended, but as the genitive *hans* where 'mere anaphora' with the embedded DP is intended:

- (i') a. [sérhver maður sem á [son]²]¹ elskar konuna sína₁
 ('[every man who has [a son]²]¹ loves his₁ wife')
 b. [sérhver maður sem á [son]²]¹ elskar konuna hans₂
 ('[every man who has [a son]²]¹ loves his₂ wife')
 (ii') a. [aðeins [Jón]²]¹ elskar konuna sína₁
 ('[only [John]²]¹ loves his₁ wife')

incomplete description in (10') are taken to be elliptical may itself contain a pronoun, moreover a pronoun understood as a bound variable. That is, the truth conditions of utterances of (10), (10') and (10'') might be captured using the following formula of our semi-formal language:

- (11) $[every_x: man\ x \bullet [just\ one_y: donkey\ y]\ x\ bought\ y]$
 $[the_y: donkey\ y \bullet x\ bought\ y]\ x\ had\ to\ pay\ for\ y\ with\ cash.$

(This is not an LF, of course, but it may correspond quite closely to the LF of (10''). See below.) The description is Russellian but *relativized* in the sense that uniqueness is relative to choice of man who bought only one donkey. This sort of relativization should occasion no surprise, and I embraced it in *Descriptions*. The processes at work in interpreting an incomplete description or a descriptive pronoun are pragmatic and richly inferential, and it is clear hearers have no trouble coming up with (10'') when quizzed about (10) or (10').⁷⁹

The relativization in (10) and countless other donkey sentences—indeed in countless *non*-donkey sentences (in ‘every man loves the woman he married’, for example)—is sometimes called ‘covariation’, or ‘implicit binding’. The three labels are all trying to get at the same basic point: the interpretation of the donkey pronoun ‘it’ is *relativized* to or *covaries* with objects satisfying ‘man who bought only one donkey’. The label ‘implicit binding’ is, perhaps, slightly more loaded than the other two as it may lull (or help those with an agenda lure) the unsuspecting into thinking the concept of concern is a *syntactic* one, that it concerns the explicit binding of an aphonic variable at LF. And that is certainly not the case. In their long discussions of this type of implicit binding, Evans (1977), Davies (1981), Ludlow and Neale (1991), and Neale (1990, 1993a, 1994) discuss the phenomenon and do not posit special variables in underlying syntax for the quantifiers to bind. The last four of these works are highly sympathetic to LF, but there is no commitment to aphonic variables at LF being so-bound, and no claim that (10), (10') and (10'') share an LF.⁸⁰

- b. $[aðeins\ [Jón]^2]^1\ elskar\ konuna\ hans_2$
 (‘[only [John]²]¹ loves his₂ wife’).

For discussion, see Neale (forthcoming b).

⁷⁹ Neale (1990, 1993a, 1994) toyed with the idea of specifying the descriptive content of D-type pronouns using a simple algorithm and then pointed to problems suggesting retrieval was a looser pragmatic matter, perhaps guided or shaped by formal factors and 78, strong interpretive heuristics.

⁸⁰ It is odd that Stanley and Szabó (2000a,b) and Stanley (2000, 2002a,b) assume without argument that relativization (implicit binding) requires an actual variable at LF, given the existence of accounts of relativization in the literature that do not postulate aphonic variables at LF that the subject quantifier binds. What Stanley and Szabó *do* argue is that *there are* all sorts of examples that involve relativization. But *that* is something people who discuss relativization already knew. For a while I thought perhaps an implicit argument was lurking behind Stanley and Szabó’s assumption: it is impossible to entertain quantified thoughts without entertaining quantified natural language sentences containing variables that the quantifiers bind—the evidence would have to come from psychology, and I am aware of none. But even if this turned out to be true, would it be relevant? Surely I can utter a sentence *X* and thereby conversationally implicate something we might *describe* using a quantified sentence *Y*, but that does not mean *X* itself must contain a

So what do the LFs of (10), (10') and (10'') look like. This is something I took up in Neale (1993). Given the role LFs are supposed to play in syntactic theory, I suggest roughly the following:⁸¹

- (12) [every¹ man₁ who³ [[[just one]² donkey₂]² [_s e₃ bought e₂]]_i]¹
[it]² [_s e₁ paid for e₂ with cash]
- (12') [every¹ man₁ who³ [[[just one]² donkey₂]² [_s e₃ bought e₂]]_i]¹
[the² donkey₂]² [_s e₁ paid for e₂ with cash]
- (12'') [every¹ man₁ who³ [[[just one]² donkey₂]² [_s e₃ bought e₂]]_i]¹
[the² donkey₂ e₁ owns e₂]² [_s e₁ paid for e₂ with cash].

In (12) 'it' has been raised, which goes hand-in-hand with its interpretation as a quantifier, indeed as a description, a full descriptive content for which the speaker expects the hearer to come up with pragmatically, just as with the incomplete description in (12').

9. Incompleteness and Persistence

In order to bring into the picture certain mathematical concepts of importance, let us idealize in certain ways temporarily, abstracting from the context-sensitive and elliptical ways of ordinary speech. Following Barwise and Cooper (1981), let us say that a determiner *D* is *persistent* just in case the truth of

- (1) *D* A(s) is (are) B(s)

(e.g. 'some ants are black') guarantees the truth of

- (2) *D* A'(s) is (are) B(s)

(e.g. 'some animals are black') in any case where the set of things that are A' (e.g. animals) properly includes the set of things that are A (e.g. ants). So whereas 'some' is persistent (witness the fact that if 'some ant is black' is true, so is 'some animal is black'), 'every' is not (witness the fact that the truth of 'every ant is black does' does not guarantee the truth of 'every animal is black').

In *Descriptions*, I noted in passing that simple cases of matrix incompleteness arise naturally in connection with uses of non-persistent determiners such as 'every', 'no', and 'the', whilst 'derived' or 'inverted' cases arise naturally in connection with uses of

variable for the quantifier in *Y* to bind. It is hard to believe that relativized interpretations of utterances are going to present problems for general pragmatic processes given that these processes must also be capable of revealing conversational implicatures, irony, metaphor, and so on.

⁸¹ I here assume with Evans (1977) and Quine (1960) that a relative pronoun is a device of predicate abstraction, the subscript on its relative clause indicating the predicate's argument. Thus the subject DPs are construed as [every x_1 : man(x_1) • [λx_3 [just one x_a : donkey(x_2)] x_3 bought x_2)](x_1)]

persistent determiners such as ‘a’ and some ‘some’.⁸² That note exhausted what I felt needed to be said at the time because I did not then see incompleteness as particularly interesting or threatening to Russell’s Theory of Descriptions. Subsequent literature, has convinced me the phenomenon is more interesting than I once thought, and while writing the present works and reviewing the footnotes that were culled from *Descriptions* before it was first published, I came across two that went into rather more detail; realizing their importance to the current discussion, I have distilled them for use here.

Simple cases exemplifying the problem of incomplete matrices arise naturally in connection with the use of quantifiers that are non-persistent, such as ‘every’, ‘no’, ‘the’; inverted cases arise naturally in connection with those that are persistent, such as ‘some’. The general problem of incompleteness, recall, is to explain how it is that a matrix $A(x)$ is understood, in context, as satisfied by fewer objects than its superficial form seems to require. The contrast between the implicit and explicit approaches is best summarized thus: On the implicit approach, $A(x)$ is satisfied by fewer objects because it is understood as *ranging over* fewer objects; on the explicit approach, it is satisfied by fewer objects because it is understood as abbreviating $A(x) \bullet C(x)$.

The problem has simple and inverted forms in the following sense. In the *simple* case, an utterance of ‘ D A is B ’ is understood, in context, as expressing a *truth*, despite the evident *falsity* of the formula $[Dx: A(x)]B(x)$ (taken at face-value). Examples, in context, might be the following:

- (3) the table is covered with books
[false—there are millions of tables]
- (4) every student swam the Hellespont
[false—millions of students have never been to Turkey]
- (5) no student lives in China
[false—millions of Chinese students live in China].

In *inverted* cases, an utterance of ‘ D A is B ’ is understood, in context, as expressing a *falsehood* despite the evident *truth* of $[D_x: A(x)]B(x)$ (taken at face-value).⁸³ An obvious way of producing an inverted case is to introduce negation with large scope. Examples, in context, might be the following:

⁸² Neale (1990: 114, n. 45). Actually, derived cases arise naturally in connection with determiners that are *not anti-persistent*, an interestingly wider group. To foreshadow: (i) since ‘the’ is both non-persistent and non-anti-persistent, simple *and* inverted cases are readily produced; (ii) perhaps the (overlapping) categories *non-persistent* and *non-anti-persistent* (rather than the (non-exhaustive) categories *anti-persistent* and *persistent*) is the one at the core of the generalized quantifier analysis of *natural language* determiners.

⁸³ My use of ‘inverted’ in connection with forms of the Argument from Incompleteness is intended to mirror my use of the same word in connection with the Argument from Misdescription. In the *simple* case of the latter argument, used by Donnellan (1966), an utterance of ‘the A is B ’ seems to be *true* despite the apparent *falsity* of its Russellian analysis. In the *inverted* case of the argument, used by Hornsby (1978), an utterance of ‘the A is B ’ seems to be *false* despite the apparent *truth* of its Russellian analysis.

- (3') it is not the case that the table is covered with books
[true—there are millions of tables]
(4') not every student swam the Hellespont
[true—millions of students have never been to Turkey]
(5') it is untrue that no student lives in china
[true—millions of Chinese students live in China].

Alternatively, given familiar relations between quantifiers, inverted cases can be generated without negation using 'some' or 'a', which are persistent:

- (6) some students live in China
(7) there's a bottle in the fridge.

Suppose you are in my kitchen and ask me for a beer; I reply by uttering (7); you open the refrigerator and find a bottle of ketchup but no bottle (or can) of beer. Now it is certainly true that there is *a bottle* in my fridge, but not true that there is *a bottle of beer* in there. And so there is some inclination to say that I made a false statement, in this context, when I said 'there's a bottle in the fridge.'

Finishing off our square, we should find simple cases again if we introduce negation with large scope over persistent quantifiers; and in fact we do:

- (6') It is not the case that some student lives in China
(7') There is not a bottle in the fridge.

I want to appeal to certain mathematical properties of quantificational determiners, in particular persistence, in so far as they bear on the matter of the interpretation of the determiners 'the', 'this', 'that', and 'a'.⁸⁴ Thoughts about incompleteness should be cast aside for a while as I want to talk in the abstract manner of the mathematical logician *prior to* raising issues about, say, domain restrictions when discussing utterances of natural language sentences. For immediate purposes, then, two harmless fictions or idealizations will be maintained to facilitate a clear exposition of the main points and avoid distracting side-issues.⁸⁵ First, the matrices used in the examples should be taken at face-value, as requiring no completion. Second, each unambiguous English sentence should be taken to map onto exactly one truth-evaluable formula of the language of restricted quantification.

A formula $[D_x: A(x)]B(x)$ contains two open formulae, $A(x)$ and $B(x)$. The determiner Dx combines with $A(x)$ to form a restricted quantifier, and that quantifier combines with $B(x)$ to form a formula.⁸⁶ Let us say that $A(x)$ occupies 'first' position with respect to Dx

⁸⁴ The general interest of these properties to semantics and syntax is discussed in detail by Barwise and Cooper (1981). See also Westerståhl (1985, 1989).

⁸⁵ When I say the fictions or idealizations are harmless, I mean only that they are harmless for the particular concerns I have—they can be very harmful in other contexts.

⁸⁶ In an alternative system of 'binary' rather than restricted quantification, Q_x might combine directly with $A(x)$ and $B(x)$ in one move to form a formula $[Q_x](A(x), B(x))$, but as a way of having our notation closer to the structure of English we have opted to treat $A(x)$ and $B(x)$ differently.

while $B(x)$ occupies ‘second’ position in the formula $[D_x: A(x)]B(x)$. We can now examine the impact particular determiners have on these positions in some detail.

Talk of persistent quantifiers can be extracted straightforwardly from Barwise and Perry’s talk of persistent statements. The basic idea is this: a quantifier $[D_x: A(x)]$ is persistent if and only if a simple sentence containing it remains true if $[D_x: A(x)]$ is replaced by $[D_x: A'(x)]$, where the set of things satisfying $A'(x)$ properly includes the set of things that satisfy $A(x)$, as for example the set of animals properly includes the set of ants. Thus ‘some ants’ is persistent because if, say, ‘some ants are black’ is true, so is ‘some animals are black’. It will be useful to specify this more precisely.

1. On Being ($\uparrow 1$). A determiner D is $\uparrow 1$ (‘one-up’, ‘upward entailing in first position’, ‘persistent’), if the following inference is valid (for arbitrary A , A' , and B):

$$\begin{array}{l} (\uparrow 1) \quad [D_x: A(x)]B(x) \\ \quad \quad [every_x: A(x)]A'(x) \\ \quad \quad \hline [D_x: A'(x)]B(x). \end{array}$$

Thus:

$$\begin{array}{l} (8) \quad [some_x: ant(x)] black(x) \\ \quad \quad [every_x: ant(x)] animal(x) \\ \quad \quad \hline [some_x: animal(x)] black(x) \end{array}$$

Examples: ‘some’, ‘an’, ‘at least n ’, ‘a few’ (unlike ‘few’). Quick test in natural language: if DET is $\uparrow 1$ then (9') is entailed by (9),

$$\begin{array}{ll} (9') & D \text{ animal(s) is (are) black} \\ (9) & D \text{ ant(s) is (are) black} \quad (\uparrow 1) \end{array}$$

(The determiner ‘some’ may take a singular or a plural complement (‘some ant’, ‘some ants’); either way, it is $\uparrow 1$.)

The simplest form of the problem of incompleteness arises in connection with determiners that are *non*-persistent, $\nmid 1$ (‘not one-up’). It is evident from consulting (9) and (9') that ‘every’ and ‘no’ are $\nmid 1$. (Like ‘some’, the determiner ‘no’ may take a singular or a plural complement (‘no ant’, ‘no ants’); either way, it is $\nmid 1$).⁸⁷ The numerical determiners ‘exactly n ’ and the proportional determiner ‘most’ are also $\nmid 1$ (unlike ‘at least n ').

More importantly for our concerns, on the Russellian account we are assuming, ‘the’ is also $\nmid 1$. Instantiating (9') and (9), we get (10') and (10):

$$(10') \quad \text{the animal is black}$$

⁸⁷ So is ‘all’, which could, perhaps, be seen as a syntactically plural form of ‘every’, though nothing I shall say here turns on such an assumption.

(10) the ant is black ($\uparrow 1$)

Unpacking these as (11') and (11), in Russell's manner, makes it clear that (10') is not entailed by (10):

(11') there is exactly one animal and every animal is black.

(11) there is exactly one ant and every ant is black. ($\uparrow 1$)

Like 'some' and 'no', the definite article may take a singular or a plural complement ('the ant', 'the ants'); and, naturally enough, 'the' remains $\uparrow 1$ when it takes a plural complement. This is worth exploring.

A useful way of thinking about Russell's Theory of Descriptions is as follows: Whereas the determiners 'every' and 'all' introduce *universal* quantifications and the determiners 'a', 'an', and 'some' introduce *existential* quantifications, the determiner 'the' simultaneously introduces *both*. For some purposes, it is useful to think of definite descriptions as complex *existential* phrases; for others it is more useful to think of them as complex *universal* phrases. For example, if one focuses on the existential character of Russell's proposal it is easy to explain it to the novice by building upon Russell's proposal for *indefinite* descriptions because a sentence of the form 'an *A* is *B*' is analysed by Russell as (12), whilst one of the form 'the *A* is *B*' can be analysed as (13), which is logically equivalent to Russell's preferred formula:

(12) $\exists x(Ax \bullet Bx)$.

(13) $\exists x((Ax \bullet Bx) \bullet \forall y(Ay \supset y=x))$.

But if one focuses on the universal character of the proposal, the relationship between singular and plural definite descriptions comes more clearly into view. There is a just a cardinality difference between the truth condition of an utterance of 'the *A* is *B*' and that of an utterance of 'the *As* are *Bs*': the former is true if and only if every *A* is *B* and there is *exactly one A*; the latter is true if and only if every *A* is *B* and there is *more than one A*.⁸⁸

We can now confirm that 'the' remains $\uparrow 1$ when it takes a plural complement. Instantiating (9') and (9) again, we get (14') and (14):

(14') the animals are black

(14) the ants are black ($\uparrow 1$)

Unpacking these as (15') and (15) makes it clear that (14') is not entailed by (14):

(15') there are two or more animals and every animal is black.

(15) there are two or more ants and every ant is black ($\uparrow 1$)

⁸⁸ The determiner 'both' is surely the *dual* form of 'the': 'both *As* are *Bs*' is true if and only if every *A* is *B* and there are exactly two *As*. Singular, dual, and plural truth conditions may be put into a set-theoretic notation that renders the relations between them transparent

'the *A* is *B*' is true if and only if $|A-B| = 0$ and $|A| = 1$

'the *As* are *Bs*' is true if and only if $|A-B| = 0$ and $|A| > 1$

'both *As* are *Bs*' is true if and only if $|A-B| = 0$ and $|A| = 2$.

So the simple form of the problem of incompleteness should arise naturally in connection with uses of both ‘the *A*’ and ‘the *As*’, which in fact it does. By contrast, the inverted problem should arise naturally in connection with uses of both ‘some *A*’ and ‘some *As*’, and for uses of indefinite descriptions of the forms ‘an *A*’ and ‘a few *As*’.

It is frequently suggested that demonstrative descriptions (complex demonstratives) of the forms ‘this *A*’ and ‘that *A*’ are special forms of definite descriptions. But the persistence test suggests they are, at least when used demonstratively, special forms of *indefinite* descriptions. Instantiating (9′) and (9) again, we get (16′) and (16):

- (16′) that animal is black
 (16) that ant is black (↑1)

In order for the intuitive test to make sense, two assumptions must be made. First, ‘that *A* is *B*’ is not true unless the object ‘that *A*’ denotes is an *A*; second, the occurrences of ‘that animal’ and ‘that ant’ in (16′) and (16) are being used to denote the same object. (The assumptions seem harmless if we are interested merely in the codification of inference.) Since (16′) is entailed by (16) on these assumptions, ‘that’ is ↑1 and so patterns with ‘a’ and not with ‘the’ (which is †1).⁸⁹ The same is true of the plural form ‘those’, and of ‘this’ and its plural form ‘these’. So, for example, (17′) is entailed by (17):

- (17′) these animals are black
 (17) these ants are black (↑1)

In summary, while ‘the’ is †1, ‘a’, ‘that’, ‘this’, ‘these’, and ‘those’ are ↑1.

2. *On Being* (↓1) One way for a determiner to be †1 is for it to be ↓1 (‘one-down’, ‘downward entailing in first position’, or ‘anti-persistent’). A determiner *D* is ↓1 if and only if the following is valid:

$$\begin{array}{l} (\downarrow 1) \quad [D_x: A'(x)]B(x) \\ \quad \quad [every_x: A(x)]A'(x) \\ \hline \quad \quad [D_x: A(x)]B(x) \end{array}$$

Examples: ‘every’, ‘no’, ‘at most *n*’. Quick test: If *D* is ↓1 then (9′) entails (9) (rather than being entailed by (9), as in the ↑1 case):

- (9′) *D* animal(s) is (are) black (↓1)
 (9) *D* ant(s) is (are) black

⁸⁹ Actually matters are slightly more complicated than indicated here. One issue is highlighted by the fact that demonstrative descriptions are sometimes used in ways that require them to be understood as Russellian definite descriptions, and on such uses the determiner ‘that’ is †1 rather than ↑1. This would seem effectively to preclude the simplest unitary accounts of the determiner ‘that’.

What of the definite article? We have already seen that it is $\uparrow 1$. Is it $\uparrow 1$ because it is $\downarrow 1$? Or is it $\downarrow 1$ as well as $\uparrow 1$? Along with ‘exactly n ’ it is $\downarrow 1$, as demonstrated by the fact that (10') does not entail (10), and (14') does not entail (14):

- (10') the animal is black ($\downarrow 1$)
 (10) the ant is black
 (14') the animals are black ($\downarrow 1$)
 (14) the ants are black

So ‘the’ and ‘exactly n ’ are both $\uparrow\downarrow 1$ (‘one-flat’). This should come as no surprise: ‘exactly n ’ is $\uparrow\downarrow 1$ because ‘at least n ’ is $\uparrow 1$ whilst ‘at most n ’ is $\downarrow 1$; and ‘the’ is $\uparrow\downarrow 1$ because ‘some’ is $\uparrow 1$ while ‘every’ is $\downarrow 1$ (think about it set-theoretically). Certain derived partitive quantifiers are also $\uparrow\downarrow 1$: ‘exactly n of the’, ‘half of the’ and ‘most of the’, as well as the simple proportional determiner ‘most’, as the reader can easily verify.

The important point to bear in mind when we return to incompleteness is that simple cases arise naturally in connection with determiners that are $\uparrow 1$, such as ‘the’, ‘every’ and ‘no’, and inverted cases occur naturally in connection with those that are $\downarrow 1$, such as ‘the’ (again), ‘some’ and ‘a’—the referential commitments of ‘this’ and ‘that’ rule out significant talk of incompleteness.

Let us now turn to the impact of determiners on second position.

3. *On Being* ($\uparrow 2$). D is $\uparrow 2$ (‘two-up’, ‘upward entailing in second position’, ‘monotone increasing’) if and only if the following is valid:

$$\begin{array}{l} (\uparrow 2) \quad [D_x: A(x)]B(x) \\ \quad [every_x: B(x)]B'(x) \\ \hline [D_x: A(x)]B'(x) \end{array}$$

Examples: *every* and *some*. Quick test: If DET is $\uparrow 2$ then (18') is entailed by (18):

- (18') D ant(s) is (are) black or brown
 (18) D ant(s) is (are) black. ($\uparrow 2$)

4. *On Being* ($\downarrow 2$). Dx is $\downarrow 2$ (‘two down’, ‘downward entailing in second position’, or ‘monotone decreasing’) if and only if the following is valid:

$$\begin{array}{l} (\downarrow 2) \quad [D_x: A(x)]B'(x) \\ \quad [every_x: B(x)]B'(x) \\ \hline [D_x: A(x)]B(x) \end{array}$$

Examples, *no* and *few*. Quick test: If DET is $\downarrow 2$ then (18') entails (18).

It is clear from the fact that ‘some’ and ‘every’ are both $\uparrow 2$ that on a Russellian analysis ‘the’ will also be $\uparrow 2$ (think about it set-theoretically), and the relevant test confirms this. Instantiating (18') and (18) we get (19') and (19):

(19') the ant is black or brown.

(19) the ant is black ($\uparrow 2$)

And the latter entails the former, as the Russellian unpacking makes clear:

(20') there is exactly one ant and every ant is black or brown.

(20) there is exactly one ant and every ant is black ($\uparrow 2$)

In this respect, 'the' differs from 'exactly one', 'just one', and 'one and only one', which are $\uparrow\downarrow 2$.

The fact that 'a' and 'some' are $\uparrow 2$ suggests that 'this' and 'that' (and their plural forms) will be $\uparrow 2$. Again, the relevant test provides confirmation (making the same assumptions as before). For example, (21) entails (21'):

(21') that ant is black or brown

(21) that ant is black ($\uparrow 2$).

We have identified ten properties. Six of them can be characterized succinctly as follows:

DET is $\downarrow 1$ if we can move 'down' from truth to truth,

$\uparrow 1$ if we can move 'up',

$\uparrow\downarrow 1$ if neither:

($\downarrow 1$) DET animal(s) is (are) black

($\uparrow 1$) DET ant(s) is (are) black.

DET is $\downarrow 2$ if we can move 'down' from truth to truth,

$\uparrow 2$ if we can move 'up',

$\uparrow\downarrow 2$ if neither:

($\downarrow 2$) DET ant(s) is (are) black or brown

($\uparrow 2$) DET ant(s) is (are) black.

Like $\uparrow\downarrow 1$ and $\uparrow\downarrow 2$, the four further negative properties $\uparrow 1$, $\downarrow 1$, $\uparrow 2$, and $\downarrow 2$ can be characterized in terms of negation, conjunction, and $\downarrow 1$, $\uparrow 1$, $\downarrow 2$ and $\uparrow 2$, which the table takes as 'primitive'. But it is not obvious this is the most illuminating way of organizing things, at least if *meaning* (as opposed to *inference*) is our main concern. Perhaps we are better off viewing $\uparrow 1$, $\downarrow 1$, $\uparrow 2$, and $\downarrow 2$ as 'primitive' (as *not* allowing true up and down moves) and seeing the rest as Boolean products.

As far as definite, indefinite, and demonstrative descriptions are concerned, matters may be summarized thus:

$\uparrow 1$ a, this, that

$\uparrow\downarrow 1$ the

$\uparrow 2$ the, a, this, that

Ladusaw (1981) proposes that negative polarity items (e.g. 'ever' and 'only') occur felicitously only in $\downarrow$ environments. If this is correct, such items are precluded from both

positions of ‘the’, which is $\uparrow\downarrow 1$ and $\uparrow 2$. May (1985) accepts Ladusaw’s hypothesis and a Russellian semantics for singular ‘the’, but suggests *plural* descriptions may contain negative polarity items in *first* position. This would require ‘the’ to be merely $\uparrow 1$ with a singular complement but fully $\downarrow 1$ with a plural. Since other determiners (e.g. ‘no’, or some’) do not alter in respect of persistence in this way, something seems to be wrong here. Perhaps Ladusaw’s proposal should invoke $\uparrow$ rather than $\downarrow$. Is this supported by the fact that ‘most’ is $\uparrow\downarrow 1$? (Richard Larson has pointed out to me, issues about partitives intrude here.) The issues with demonstratives are also unclear: ‘that’ is $\uparrow 1$ and should not license negative polarity items in first position on either the original or revised proposal. Ludlow’s (2002) fresh insights on the vexed matter of the bearing of $\uparrow\downarrow$ properties on the nature of syntactic theory are important here.

10. Sentence Ellipsis

(1) and (1’) may be said to share the same LF and differ only at PF:

(1) John can waltz, but Paul can’t.

(1’) John can waltz, but Paul can’t waltz.

In (1), the PF results from the deletion of a VP (verb phrase) under stringent conditions of ‘identity’ that guarantee its ‘recoverability.’

In order to understand what is at stake here and later in this chapter, some history is helpful. The main ideas can be sketched without too much detail, although the nature of the battles of the 1960s and early 1970s means that much is still disputed.

Early generative grammar posited two levels of syntactic representation, usually called deep structure and surface structure. Deep structures were essentially the products of a lexicon and phrase structure rules such as the following:

$$\begin{aligned} S &\rightarrow NP + VP \\ VP &\rightarrow V + NP. \end{aligned}$$

Phrase structure rules were supplemented by *transformational* rules, which were meant to bridge the gap between deep structure and surface structure, and explain the relationships between certain sentences (for example, active-passive pairs) by deleting, adding, and reordering elements in phrase markers (phrase structure trees) under clearly specified conditions.⁹⁰ Katz and Postal (1964) put forward the hypothesis (often called ‘Katz-Postal’) that transformational rules did not alter meaning, from which it virtually

⁹⁰ Deep structures were meant to be generated by sets of *context-free* phrase structure rules, i.e. rules of the form, (i) $\alpha \rightarrow \beta_1 \dots \beta_n$. These may be contrasted with context-sensitive rules of the form (ii) $X\alpha Y \rightarrow X\beta_1 \dots \beta_n Y$, which specify that if α occurs in the context $X\alpha Y$ it can expand to $\beta_1 \dots \beta_n$. The mapping between deep structure and surface structure was effected by context-sensitive transformational rules. In the philosophy of language a related distinction can be found, between compositions that are what Davidson calls *semantically innocent* (i.e. context invariant) and compositions which are not.

followed that deletion transformations could take place only under conditions that made ‘recovery’ possible, something Chomsky stressed the same year:

A transformation can delete an element only if this element is the designated representative of a category, or if the structural condition that defines this transformation states that the deleted element is structurally identical to another element of the transformed string. A deleted element is, therefore, always recoverable. (1964: 41)

We are sorely in need of examples. In (2), ‘John’ is the subject of the main clause and ‘Paul’ is the subject of the embedded clause (here infinitival):

(2) John wants Paul to leave.

Compare (2) with (3), where there appears to be no subject of the embedded clause:

(3) John wants to leave.

One early and widely discussed transformational deletion was *Equi NP deletion*, where the embedded clause of a verb like ‘want’ was deleted on account of being ‘identical’ to the subject of the main verb. The rough idea was that (3) was derived from the underlying form (3’):

(3’) John wants John to leave.

(it could not be not be derived from, say, (2).) In short, we have something we might call ‘recoverability of deletion under identity’, although we need to specify exactly what is meant by ‘identity’ here. Recoverability was a fundamental condition governing deletion. To abandon it—as some generative semanticists did—was to abandon serious work on the syntax of natural language.⁹¹

Obviously there are all sorts of issues to take up, even in connection with this single example—is the deletion obligatory or optional? is *identity* to be construed as a *formal* or *interpretive* notion? And so on. At first blush, the second question seems to have a clear answer. If I have two friends called John, even if I am being *unclear*, surely I am not speaking *ungrammatically* when I utter (3’), intending to refer to one of the friends using the subject of the main verb, and to the other using the subject of the embedded clause. Only when I intend the main and embedded subjects to refer to the *same person* does *Equi* take place.⁹² So *interpretive* identity is required (as one would expect, assuming

⁹¹ Eight years after Chomsky’s remark appeared in print, the premier linguistics journal *Linguistic Inquiry*, published a piece by Fiengo and Lasnik called ‘On Nonrecoverable Deletion in Syntax.’ (1972: 528) There was no text, just the authors’ names and their institutional affiliation (they were both then at MIT) and a blank space.

⁹² In the right circumstances could I not use (3’) intentionally to talk about a single individual? ‘I want Tom to leave, you want Tom to leave, Mary wants Tom to leave, Fred wants Tom to leave, *Tom* wants Tom to leave, but he can’t until the curtain comes down.’ Perhaps the fact that heavy stress is needed on the embedded occurrence indicates that a rule of sentence grammar is being stretched, however.

Katz-Postal).⁹³ Something like *formal* identity also seems to be required, however, for otherwise (4) could be the result of *Equi* applying to (4') when 'Cicero' and 'Tully' share a subscript:

- (4) Cicero wants to leave
- (4') Cicero wants Tully to leave.

On the assumption this is not something we want from a grammar, it would seem that the operation of *Equi* is constrained by formal and interpretive identity, so let us use 'identity' in this way in this context.

A type of deletion of more theoretical interest is *VP deletion*. VP deletion raises questions about the 'scope' of transformational rules and about variable-binding and the interpretation of pronouns. Since the second of these will assume some importance later, it is good to get clear about the main issue right away.⁹⁴ VP deletion appears to take place in the derivation of (1) from (1'):⁹⁵

- (1) John can waltz, but Paul can't.
- (1') John can waltz, but Paul can't waltz.

The deletion is recoverable because it takes place only under identity: (1) cannot be obtained by *VP deletion* from (and hence cannot be understood in terms of), say, (5) or (6):

- (5) John can waltz, but Paul can't tango.
- (6) John can waltz, but Paul can't speak German.

Recoverability and Katz-Postal respected again. It is no more a 'merely pragmatic' fact that (1) must be read as equivalent to (1') than is the fact that 'himself' must be read as bound by 'John' in

- (7) John told Paul's wife a lot about himself.⁹⁶

VP deletion raises two immediate questions. First, transformational rules were meant to operate on sentences, but VP deletion appears not to respect sentence boundaries, utterance boundaries, or even speaker boundaries, witness the following natural dialogue between *A* and *B*.

⁹³ It was not appreciated at the time that even under interpretive identity (3) and (3') are not equivalent. In current theory, the subject or 'want' is an aponic PRP. 'John wants PRO to leave' must be read *de se*, meaning it cannot be interpreted as a copy of 'John'

⁹⁴ I discussed VP deletion in several footnotes in *Descriptions*, but what I said was mostly rubbish.

⁹⁵ For high-level introductions to VP ellipsis, see ch 9 of see Heim and Kratzer (1998) and May (2002). It is common to distinguish VP ellipsis from VP anaphora. Example involving the latter would be (i) and (ii), which contain the anaphoric elements 'it' or 'so': (i) Bill can waltz, but Tom can't do it; (ii) Bill can waltz, and so can Tom. On the difference between VP anaphora and VP ellipsis, see May (2002).

⁹⁶ There are good reasons for thinking that reflexives are not referential expressions but devices that combine with predicates to form larger predicates, making 'John loves himself' equivalent to 'John is a self-lover'. On the assumption that an intransitive verb 'snores' is understood as $(\lambda x(\text{snores } x))$ and an intransitive verb as $(\lambda y(\lambda x(x \text{ loves } y)))$, this will drop out on a suitable account of reflexives. See Neale (forthcoming *a*).

- (8) A: John can waltz.
 B: I know. It's a shame Mary can't.

This would appear to show that the simplest transformational account of VP deletion is in trouble, at least on a standard account of what a transformational rules are, that the phenomenon is a *discourse* phenomenon, albeit one that is *grammatically constrained*. As May (2002: 1095) succinctly puts it, to the extent that VP deletion can be thought of as a rule, it would appear to be more of a rule of *discourse* grammar than *sentence* grammar. There is no need to pursue this matter here. Let us imply note that syntactic theory has evolved considerably since the heyday of transformations. In Chomsky's (1995, 2002) current set-up we might say that (1) and (1') have the same LF, (1_λ) , but different PFs, the PF of the former given by something like (1_π) , after VP ellipsis:⁹⁷

- (1_λ) $[_S \text{ Bill } [_{\text{can}} [_{\text{VP}} \text{ waltz}]] \text{ but } [_S \text{ Tom } [_{\text{can't}} [_{\text{VP}} \text{ waltz}]]]]$
 (1_π) $[_S \text{ Bill } [_{\text{can}} [_{\text{VP}} \text{ waltz}]] \text{ but } [_S \text{ Tom } [_{\text{can't}} [_{\text{VP}} e]]]]$.

Katz-Postal can be transposed accordingly: interpretation remains constant under linguistic ellipsis (henceforth, this is what I shall mean by 'Katz-Postal').

The second thing VP ellipsis brings out is an important ambiguity, independently motivated and, in effect, predicted by Katz-Postal. There is only one reading of (9), but there are two readings of (10):⁹⁸

- (9) John loves his wife, but Mary doesn't
 (10) John loves his wife, but Paul doesn't.

On one reading of (10) Paul is being said not to love John's wife, on the other not to love his own wife. (The former is usually called the *strict* reading, the latter the *sloppy* reading.) The existence of the two readings of 'Paul doesn't' as it occurs in (19) follows from the fact that on standard accounts of the interpretation of pronouns there are two distinct, but truth-conditionally equivalent readings of 'John loves his wife':

- (11) $John(\lambda x(x \text{ loves John's wife}))$
 (12) $John \lambda x(x \text{ loves } x\text{'s wife})$.

The fact that there are two ways of construing (10) does not jump out at us until we look at inference patterns or consider more complex examples containing 'John loves his wife' as a proper part, (10) for example, or an attitude report like (13):⁹⁹

⁹⁷ Strictly speaking, (1_π) cannot be a PF in Chomsky's current theory because it is still annotated for phrasal structure. I think Chomsky has been misunderstood on this matter. On the sort of pure theory Chomsky is articulating, PF is just a set of instructions to the sound system and unless empirical evidence to the contrary emerges the null hypothesis is that these instructions need no access to information about phrase structure of the sort present at LF. To say this is not to say that the instructions constituting PF have been derived without information about phrase structure.

⁹⁸ Such examples are discussed extensively by Sag (1976) Williams (1977). For recent, user-friendly discussion, see Heim and Kratzer (1998) and May (2002).

⁹⁹ See Soames (1990). Similarly with examples like 'Only John loves his wife'.

- (13) Mary believes that John loves his wife.

Assuming Katz-Postal, we can now see what is going on in (10). On one reading, the occurrence of ‘his’ in the first clause is being used to refer independently to John; on the other reading, ‘his’ is bound by ‘John’. If the former, the recovered predicate is the one in (11); if the latter, the one in (12). Either way, we have recoverability. And we can also elegantly explain what is going on in an example like (14), which combines inference and VP deletion:

- (14) Every Englishman loves his wife, but John doesn’t, so John can’t be English.¹⁰⁰

Three final, quick examples of syntactic ellipsis will give us all we need. The first is *stripping*, where the surface form (15’) is optionally reduced to (15):

- (15) Smith took a job in Uruguay after the war, or Paraguay.
 (15’) Smith took a job in Uruguay after the war, or Smith took a job in Paraguay after the war

In (15), everything but ‘Paraguay’ has been deleted (‘stripped’) from the second disjunct under identity (of the material surrounding ‘Uruguay’ and ‘Paraguay’). Again, the deleted material is recoverable. From (15) we cannot recover (16), for example:

- (16) Smith took a job in Uruguay after the war, or Smith collected stamps from Paraguay after the war.

Relatedly, but via a process of syntactic ellipsis called *gapping*, (17) can be obtained from (17’) but not from (18):

- (17) Smith went to Uruguay, and Jones to Paraguay.
 (17’) Smith went to Uruguay, and Jones went to Paraguay.
 (18) Smith went to Uruguay, and Jones sent his wife to Paraguay.

¹⁰⁰ Some very interesting examples, discussed by Elbourne (2001), appear to demonstrate the intersection of sentence ellipsis and utterance ellipsis involving descriptions. (To Elbourne they demonstrate the need to see D-type pronouns as real descriptions at LF.) The facts are truly maddening. First off, the VP ellipsis in (i) seems fine, but the simple recovery (ii) is not,

- (i) every Englishman who bought a donkey vaccinated it (/the donkey), but John didn’t, so he can’t be English
 (ii) * every Englishman who bought a donkey vaccinated it, but John didn’t vaccinate it (/the donkey), so he can’t be English.

However, a recovery that contains a full description containing a bound pronoun is fine:

- (iii) every Englishman who bought a donkey vaccinated it, but John didn’t vaccinate the donkey he bought, so he can’t be English.

I really don’t know what to make of this. The idea that at LF these donkey pronouns *really are full descriptions containing bound pronouns* might seem to provide *half* the story: recovery is of what is present at LF (which, after all, is where the conditions for ellipsis must be satisfied); but why can’t the description in the *recovered* VP collapse into a D-type pronoun the way the occurrence in the *donating* VP does? I find this utterly baffling. Perhaps Elbourne’s packed paper contains the answer but I have not had time to study it.

Another example, involving two distinct deletions, might be the derivation of (19) from (19''), perhaps via (19'):

- (19) one coffee with milk and one without.
- (19') one coffee with milk and one coffee without
- (19'') one coffee with milk and one coffee without milk.¹⁰¹

Again the deletions are recoverable, which explains why (19) cannot be derived from (and so understood as), say, (20):

- (20) one coffee with milk and one tea without sugar.¹⁰²

¹⁰¹ I am not here concerned with the issue of whether the derivation proceeds via (19') or via: (i) one coffee with milk and one without milk. Either way, recoverability is not violated. The example takes on more interest in cases where deletions are based on previous utterances, as in the following:

Smith: Two coffees, please.
 Jones: Milk?
 Smith: One with, one without.

I have no firm opinion of how much syntactic theory itself is supposed to contribute to our understanding of such a dialogue, but it is clear that the only recoverable interpretation of Smith's second utterance is as a request for one coffee with milk and one coffee without milk, and this strongly suggests syntax is implicated. Such examples seem problematic for Stanley's (2000) bold claim that it is impossible to perform a linguistic speech act by uttering something less than a whole sentence. According to Stanley, every purported counterexample involves either (i) a case where a whole sentence is, in fact, being uttered, or (ii) a case in which no linguistic speech act is, in fact, being performed. An utterance of 'water' made in circumstances where it is indeterminate whether to construe the remark as tantamount to an *assertion* that the speaker wants water or a *request* for water does not (because of this indeterminacy) constitute a linguistic speech act, Stanley says. Presumably he would say the same about 'one with and one without': it is indeterminate whether the speaker is *asserting* that he wants one coffee with milk and one coffee without milk rather than requesting one coffee with milk and one coffee without milk. If the speaker has not performed a linguistic speech act, why is recoverability satisfied?

¹⁰² It is sometimes claimed that syntactic deletion is involved in the following cases, where (ib) is derived from (ia), and (iib) from (iia):

- (i) a. This donkey is faster than that donkey
 b. This donkey is faster than that one
- (ii) a. This donkey is faster than the other donkey
 b. This donkey is faster than the other.

But here recoverability is violated. While it is perfectly natural to interpret an utterance of (ib) as equivalent to an utterance of (ia), this is not forced. Suppose we have been comparing a few *horses* in respect of speed. After a while Smith says to me, while pointing, 'I think I'm going to buy this one; he's a bargain'. I reply, pointing first at the horse Smith is thinking of buying, and then at my donkey, 'Not *that* one, surely! This *donkey* is faster than *that* one.' The word 'donkey' needs to be stressed, I think, but the example still seems to show that recoverability is violated and hence that the occurrence of 'one' in (ib) is no more obtained by syntactic deletion than the occurrence of 'one' in (iii): (iii) I'll take this one, please. (Imagine a shopper who has been examining several items and is aware that a sales assistant is watching. The sales assistant holds back until he is reasonably sure the shopper has made a choice, at which point he goes over and says to the shopper, 'Can I help you, sir?' The shopper replies by uttering (iii).)

What example (i) seems to show—example (ii) is not relevantly different, but notice that 'the other' and 'the other donkey' are both incomplete descriptions—is that the interpretation of 'one' cannot be determined automatically from grammatical considerations alone. In many contexts there will be a natural tendency to interpret an utterance of (ib) as equivalent to an utterance of (ia), but that interpretation is not *forced* by grammar. While there may be heuristics of interpretation that draw on syntactic information to provide default interpretations of utterances of 'one' and also utterances of some pronouns, when we find that the hypothesized default can be over-ridden very naturally in a particular syntactic structure, we have found a

It should be abundantly clear why *sentence* ellipsis could be of no relevant interest to anyone who agrees with Quine and Sellars that *utterances* of ‘the table is large’ are elliptical and as such do not counterexemplify Russell’s Theory of Descriptions. For there is nothing in the old idea that such utterances are elliptical to suggest the sentence ‘the table is large’ is itself elliptical in the syntactician’s sense.¹⁰³ Consider (21)-(23):

- (21) [s[_{DP} the table] [v_P is dirty]].
- (22) [s[_{DP} every table] [v_P is dirty]]
- (23) [s[_{DP} no table] [v_P is dirty]]

The idea of recovering syntactic deletions here is absurd; indeed invoking syntactic deletion here would involve making a category mistake, and it is hardly surprising no philosopher who sees utterances of these sentences as involving utterance ellipsis maintains that, for example, the phrase marker (1) is derived from an underlying phrase marker such as (21'),

- (1') [s[_{DP} the table we are sitting at] [v_P is dirty]].

Such philosophers are not even talking about the *sentence* specified by (21) being elliptical but about *particular dated utterances* being elliptical.¹⁰⁴

11. The Argument from Anaphora

We now have everything we need to discuss a syntactically sophisticated version of the Argument from Anaphora.¹⁰⁵ Consider (1):

- (1) *The man in the gabardine suit* is a spy. *He* tried to bribe me

The original version of the argument, which traces back to Strawson (1952) in its basics, goes like this. (i) The occurrence of ‘he’ can be understood as anaphoric on the occurrence of ‘the man in the gabardine suit’. (ii) If an occurrence of a pronoun β is

structure that does not provide enough syntactic information to sustain recoverability in the required sense, and hence a structure that we cannot see as obtained by deletion. (Two interesting examples from 1960s popular music, the first due to Dylan (1965), the second due to Lennon and McCartney (1966): (iv) The cops don’t need you and, man, they expect the same; (v) She told me she worked in the morning and started to laugh / I told her I didn’t and crawled off to sleep in the bath. (iv) involves what is often called VP anaphora because of the presence of the seemingly anaphoric element ‘the same’. (v) is particularly interesting given the various ambiguities, expectations about when sleep takes place, and other lines in the song.)

¹⁰³ Of course, an account of how we manage to interpret elliptical utterances will be needed at some point, just as a theory of how we interpret proper names and pronouns and so on will be required, and a theory of how we grasp what a speaker is conversationally implicating by his utterance. But it is hardly an objection to the (obvious fact) that there are elliptical utterances that we do not yet have a pragmatic-inferential theory capable of making clear empirical predictions.

¹⁰⁴ For detailed discussion, see below.

¹⁰⁵ The reconstruction and the criticism of the argument are taken from Ludlow and Neale (1991) and Neale (1990).

anaphoric on an occurrence of another expression α , then β is either a variable bound by α or a device that inherits its reference from α . (iii) The occurrence of the pronoun 'he' in (1) is not a bound variable. (iv) Therefore, it inherits its reference from the occurrence of 'the man in the gabardine suit'. (iv) Therefore, 'the man in the gabardine suit' in (1) is a referring expression. *Mutatis mutandis* for the indefinite description 'a cat' in (2):

(2) A cat is on the lawn; it looks like a stray to me

The argument is easily defeated. If, as Kripke (1977) and Lewis (1979) suggest, a (putatively) anaphoric pronoun may refer to an individual *raised to salience* by the utterance of a description, then either (i) or (ii) can be rejected and an alternative account provided of the referential nature of the pronouns in (1) and (2). Or, if, as the D-type theory maintains, 'he' and 'it' are just incomplete descriptions, (1) and (2) might be seen as elliptical for (1') and (2') respectively:

(1') The man in the gabardine suit tried to bribe me

(2') the cat on the lawn looks like a stray to me.

In which case the pronouns are not referential, and again (ii) is rejected.

Before looking at the syntactically sophisticated version of this argument, we need to pull a few things together. It is a familiar point, noted earlier, there are two readings of /Paul doesn't/ as it occurs in (3):

(3) John loves his wife, but Paul doesn't.

On one reading of (3) Paul is being said not to love John's wife, on the other not to love his own wife. (The *strict* reading and the *sloppy* reading.) It was also noted that the existence of the two readings follows from the fact that on one plausible of the interpretation of pronouns there are two distinct, but truth-conditionally equivalent readings of (4)

(4) John loves his wife

(4') $John((\lambda x(x \text{ loves John's wife}))$

(4'') $John(\lambda x(x \text{ loves } x's \text{ wife}))$.

On the reading given by (4'), the occurrence of /his/ in (4) is being used to make *free* (*independent*) reference to John; on the reading given by (4'') /his/ is *bound* by /John/, giving us the reading upon which (6) is a singular instance of (5) on its bound variable reading:¹⁰⁶

(5) every man loves his wife.

¹⁰⁶ For discussion, see Heim and Kratzer (1998). Evans (1977, 1980) argues that the bound/free distinction in connection with (4) and (5) can be cashed out effectively in terms of a distinction between using a pronoun in a *referentially dependent* way and using it to make *independent reference*, at least assuming a Fregean account of quantification. For discussion, see Soames (1989) and Neale (forthcoming a).

On their bound readings, someone uttering the singular (4) and the quantified (5) is predicating the same thing of John and every man, that John satisfies (6), and that every man satisfies (6):

- (6) *x loves x's wife*

The crucial feature of bound readings is that they are possible only when the pronoun lies within the scope of (is c-commanded by) its antecedent.¹⁰⁷ Thus only *free* readings of the pronouns are possible in (7)-(9) (which is not to say they may not be co-referential with 'John' by virtue of being used *freely* to refer to John:

- (7) /his wife loves John/
 (8) /the woman he married loves John/
 (9) /the woman who married him loves John/

I say that only free readings are available here for two reasons. First, notice that only free readings of the pronouns are available in the quantified counterparts of (7)-(9):

- (7') /his wife loves every man/
 (8') /the woman he married loves every man/
 (9') /the woman who married him loves every man/.

Now if we are serious about giving a unified account of what is going on in quantified and singular pairs like (4) and (5), we should be just as serious about giving a unified account of what is going on in the pairs (7')/(7), (8')/(8), and (9')/(9). Second, there is an important syntactic distinction between, on the one hand, (4) and (5), and, on the other, (7')-(9') and (9)-(11). In (4) and (5), the pronoun is within the scope of the purported antecedent, making a bound reading possible. Not so in (7')-(9') and (7)-(9).

But surely, it might be suggested, for each of (7)-(9) there is a reading upon which the pronoun and 'John' are *co-referential*. True, but irrelevant. We are discussing the asymmetric notion of *binding* not the symmetric notion of *co-reference*, remember.¹⁰⁸ Free occurrences of /his/, /he/ and /him/ in (7)-(9) can still be used to make *free reference* to whoever /John/ is being used to refer to. But on such uses they are *not* bound by /John/: they are merely co-referential with it.

And the difference between, on the one hand, the singular forms (7)-(9), and on the other, their quantified counterparts (7')-(9'), is just that in utterances of the former, the pronouns can be used freely to refer to whoever 'John' is being used to refer to, John now being as salient as any other potential target of an independent occurrence of a free pronoun. With utterances of the quantified examples, by contrast, there is no

¹⁰⁷ Whilst this is close to universally accepted where the 'antecedent' is a quantified DP, some linguists and philosophers (including Evans) hold out where the antecedent is a referring expression. Along with Reinhart (1983) and others, I think this is a serious mistake. For discussion, see Neale (forthcoming *a*).

¹⁰⁸ The asymmetry point is made clearly by Evans (1977, 1980), in terms of referential dependence vs. coreference.

corresponding individual to target (i.e. no individual who is also the intended referent of the relevant DP, which is quantificational).

With all of this in mind, let us now look at an interesting argument for a semantically distinct referential reading of definite descriptions, due in its essentials to Larson and Segal (1995) although I shall state the argument differently. Compare, (7'')-(9'') below with (7)-(9) and (7')-(9') above:

- (7'') /his wife loves the man I talked to/
- (8'') /the woman he married loves the man I talked to/
- (9'') /the woman who married him loves the man I talked to/.

Seemingly contrary to expectations, examples (7'')-(9'') appear to pattern with the *singular* examples (7)-(9), not with the quantified examples (7')-(9'), in permitting anaphoric links between the direct object and the pronoun buried in the subject DP. Since binding is out of the question—witness (7')-(9')—how is this to be explained? Answer: the pronoun must be functioning as a referring expression that inherits its referent from the direct object, just as in (7)-(9) where the direct object is a referring expression. But in that case the descriptions in (7')-(9') must also be functioning as referring expressions.

This argument, is essentially a syntactically elaborate version of the original Argument from Anaphora and appears to fall short of its target for the same reasons. Using a description referentially—no one doubts referential *usage*, only its semantic significance—can make a particular individual salient (just as using a name can), rendering that individual a natural target for an occurrence of a free pronoun. The Russellian does not deny that the pronouns in (7'')-(9'') may be used to refer to the individual the speaker is referring to using the descriptions; at the same time, he agrees with the syntactician that these pronouns are *not bound by* those descriptions. There is no contradiction here. Indeed, once I get to the promised rapprochement between the Unitarian and the Ambiguity Theorist, the analysis just given will seem more compelling. If the description is used non-referentially, as it might be in a particular of, say, (10)

- (10) his prime minister walks behind the monarch,

plausibly it is D-type.

That we might be on the right track seems to be supported by the following contrast:

- (11) his wife loves a man I was just talking to
- (12) his wife loves a man.

It is easier to use the richer 'a man I was just talking to' referentially here than it is to use 'a man' so.

12. Recent Claims

The distinction between sentences (i.e. sentence types) and utterances (i.e. particular dated utterances of sentences) is so central to the philosophy of language that there can be only misery in store for those who miss it. One of the philosophers who did most to impress upon us the sentence-utterance distinction was Strawson (1950), in the same article on which he presented the alleged problem for Russell's theory raised by incomplete descriptions.

Sellars (1954) rightly took issue with Strawson on this matter (in the context of a broadside against Strawson's (1950, 1952) critique of Russell) and was himself very much alive to Strawson's sentence-utterance distinction, talking explicitly about elliptical *utterances* of 'the table is ψ ' and never about an elliptical *sentence* 'the table is ψ '. So it is striking that the sentence-utterance distinction is missed entirely by Strawson (1954) in his reply to Sellars, and again by Stanley (2002a) in a recent discussion of the Sellars-Strawson debate containing some rather wild claims about Sellars and Strawson, and about more recent contributors to the descriptions debates.

Amongst other things, Stanley (2002a) asserts:

- (1) that 'an influential paper of Wilfred [*sic.*] Sellars ... has led some philosophers of language to the view that there is a *legitimate notion of ellipsis*' [*my italics*] a '*process*' [*my italics*] distinct from syntactic ellipsis which helps us understand the gap between sentence meaning and what is said although 'no evidence has been given by Sellars for the existence of such a *process*' [*my italics*]' (2002a: 151).
- (2) that on 'the proper use of the term "ellipsis",' 'an *utterance*' [*my italics*] is elliptical just in case its surrounding linguistic context quite literally supplies the missing words', that 'the terms "elliptical" and "incomplete" only seem apt when the phenomenon involves words that are supplied to incomplete *sentences*' [*my italics*] by linguistic context' (2002a: 156).
- (3) that 'Sellars' discussion has engendered much confusion in subsequent literature on context-dependence' (2002a: 157).
- (4) that Strawson 'with characteristic incisiveness ... in his reply to Sellars ... in the very same journal issue ... gives a clear explanation of the oddity of both Sellars claim and his vocabulary' (2002a: 155).
- (5) that 'all of Strawson's points are correct,' (2002a: 156);
- (6) that some recent discussion 'takes off from an uncritical acceptance of Sellars' vocabulary and distinctions' (2002a: 156).
- (7) that 'Sellars vocabulary and distinctions suffer from ambiguities, unclarities, and errors, which ... continue to disrupt the literature on the topic today' (2002a: 156).

- (8) that ‘every paper on ‘incomplete’ definite descriptions that uncritically accepts Neale’s vocabulary of ‘explicit’ and ‘implicit’ ... must be examined to see whether it is vitiated by Neale’s reliance on Sellars’ discussion’ (2002a: 156).

These claims build on three more by Stanley and Szabó (2002a), who, in connection with the matter of incompleteness claim

- (9) to provide an ‘exhaustive’ and ‘comprehensive survey of the space of possible analyses’ (2000a: 219)
- (10) to bring to bear ‘considerations which militate against all but our own proposal’ (2000a: 219)
- (11) to ‘consider and reject ... the ‘explicit’ approach to quantifier domain restriction (*sic.*) discussed ... by Stephen Neale’ (2000a: 219).

All of this seems to be in the service of a claim made by Stanley (2000):

- (12) that ‘all effects of extralinguistic context on the truth conditions of an assertion are traceable to elements in the actual syntactic structure of the sentence uttered’ (2000: 391).

Strong stuff. Being of a pragmatist bent, I think the chances of (12) being true are at best, slim. But that is not my concern here.¹⁰⁹ I am concerned with (1)-(11). And here, Stanley’s substantive claims are easily dismissed once we examine what Sellars and Strawson actually *say*, and what those who ‘rely’ on Sellars and ‘uncritically accept’ talk of implicit and explicit approaches to incompleteness actually *say*. The ‘confusion’, the ‘errors’, and the ‘unclearities’ are all Stanley’s, and it is *he* who ‘disrupts’ the literature. Once we slice through the rhetoric and confusions, it becomes clear that those who have ‘uncritically’ accepted informal talk of ‘explicit’ and ‘implicit’ approaches to incompleteness found in the literature have nothing to fear.

13. Incompleteness, Ellipsis, Modality

The modality involved in describing utterance ellipsis is not difficult to understand, and it is surprising that Stanley and Szabó profess to have ‘no clear conception of [its] intended interpretation’ (2000b: 295) and ‘as a result, no clear conception of the “explicit approach”’ to incomplete descriptions (2000b: 295). The modal characterization of the explicit approach goes back at least to Sellars: ‘a given utterance of it is elliptical and states what would be non-elliptically stated, for example, by “The table *over here* is large” ...’ (1954: 200). Given that I *quote* this modal characterization (2000: 286), given that I take Sellars’s (1954) discussion as the *locus classicus* of the explicit approach in

¹⁰⁹ For discussion, see Neale (forthcoming *b*).

Descriptions, it is surprising that Stanley and Szabó (2000b), when explicitly discussing Sellars (1954) and Neale (1990, 2000), claim the modal characterization is original with *me* in (2000).¹¹⁰

Rather more worrying is Stanley and Szabó's (2000b) claim not to *understand* the modal characterization. What is so hard? Particular utterances of (1)

(1) the table is dirty

are understood as equivalent to utterances of longer sentences the speaker might have uttered in lieu of uttering (1): one utterance might be understood as equivalent to an utterance of (1'), another an utterance of (1''), yet another and utterance of (1'''), and so on:

(1') the table over here is dirty

(1'') the table we are about to sit at is dirty

(1''') the table you promised to clean is dirty.

There is nothing mysterious here.¹¹¹ This way of talking reflects the way ordinary speakers—Stanley and Szabo included, as we shall see—describe the situation when asked. What *is* mysterious, however, is why Stanley and Szabó profess incomprehension. Is the modality harder to grasp than the analogous modality in Stanley's (1997) discussion of rigidity?

If one utters a sentence containing a rigid term ... then one has *asserted something different* from what one *would have asserted* [*my italics*] had one uttered, on that very same occasion, any sentence differing from the original one only in the substitution of a non-rigid term for the rigid one' (1997: 135)

Consider two terms (word-types) *t* and *t'*. We shall say that *t* has the same content, with respect to a context *c*, as *t'*, just in case, for any sentence *S* which contains *t* as a constituent, an utterance *u* of *S* in *c* *would have* [*my italics*] the same assertoric

¹¹⁰ The modal character is also explicit in the discussion by Davies (1981) that I explicitly drew upon in *Descriptions*. It is also explicit in the discussions of Bach (1987, 1994) and Blackburn (1988), which I did not discuss (but which Stanley and Szabó *do* discuss) and in my claim in that Evans (1982) suggests, in connection with one of his examples, that 'he (the speaker) *would* be able to complete the description in a uniquely appropriate way' (1990: 100). (Even Quine (1940) seems on the verge of a modal characterization—but of course the statement that there is a enriched description the speaker could have used involves quantifying into a modal context!) Critics of the explicit approach such as Reimer (1992, 1998), Schiffer (1995), Ostertag (1998, 1999), Devitt and Sterelny (1999), and Devitt (1997a) also allude to its modal character. Devitt and Sterelny, who endorse the explicit approach in connection with incomplete descriptions used attributively but not in connection with those used referentially, are particularly explicit: 'The way to save the Russellian is to see [incomplete descriptions] as *elliptical*: speakers have in mind a longer description which, if asked, they would produce to complete the brief description uttered' (1999: 107). Stanley and Szabó's claim that the modal characterization is novel with me in 2000 is also straightforwardly contradicted by Stanley's (2002a: 157, n. 11) quotation of Sellars.

¹¹¹ And certainly the people mentioned in the previous footnote do not find it mysterious. Nor does Devitt (this volume) who finds the explicit approach 'generally hopeless' in connection with incomplete referential uses. Whatever it is that Devitt finds 'generally hopeless', it is not because he finds its modal character mysterious. He has it down: 'an incomplete definite is elliptical for a longer description that the speaker could supply' (p. 000).

content as an utterance of a sentence which results from S by replacing t' for t' (1997: 135).

Or harder to understand than the analogous modality in Stanley's (2002c) statement of his 'Expression-Communication Principle'?

For all S, S', c, c', such that c and c' agree on all contextual features relevant for determining what is said by S and S', S relative to c, S' relative to c', express the same proposition if and only if an utterance of S *would communicate* [*my italics*] the same thing as an utterance of S' in every context c'' meeting the following three conditions . . . (2002c: 329).¹¹²

Or harder to understand than the analogous modality in remarks by *Stanley and Szabó themselves*. When setting up the incompleteness problem in connection with a particular utterance of (1) made by Lisa to Max,

(2) every bottle is empty

Stanley and Szabó say, 'we can plausibly assume that by uttering [(2)] Lisa conveyed to Max the proposition that every bottle she just bought is empty' (2000a: 231). This is just the sort of thing the advocate of the explicit approach says, of course, and it is instructive that Stanley and Szabo avail themselves of an explicitly richer description in just *stating* the general problem. The modals are not far behind: '*Had* [*my italics*] Lisa been more explicit, she *could* [*my italics*] have conveyed the same proposition by uttering [(3)] instead' (ibid):

(3) every bottle I just bought is empty.

And there are more to come. The difference between the various approaches they are going to discuss (their own included) they say, 'lies in the way they spell out the relationship between Lisa's actual [*my italics*] utterance of [(2)] and her *hypothetical* [*my italics*] utterance of [(3)]' (2000a: 232) And according to 'one very plausible' version of one of these approaches (the 'syntactic ellipsis approach'), 'although the sentence articulated by Lisa is different from the sentence she *would have* [*my italics*] articulated in the *hypothetical* [*my italics*] case, there is no difference between what was and what *would* [*my italics*] have been uttered' (2000a: 232).¹¹³

I shall have more to say later about the explicit but innocuous modality later. Right now I want only to point out that the idea of syntactic theory recovering the underlying syntactic structure (or LF) of (1) from any of (1'), (1'') or (1''') involves a horrible category mistake, which is surely why no advocate of the explicit approach has ever advocated it. has ever held it. (And, of course, it would involve positing an infinite number of sentences (not just alphabetical variants) with the same surface form, which would fly in the face of both syntactic theory and common sense.) By contrast, there is no

¹¹² I do not mean to be suggesting that there is any merit to the Expression-Communication Principle, only that there is no obvious mystery about its modal character.

¹¹³ There is no inconsistency in Stanley and Szabó's remark as technical uses of 'articulated' and 'uttered' are in play here.

category mistake (or clash with common sense) in the following idea: upon hearing an utterance of (1), *B* might interpret *A* as expressing what *A* might have expressed more explicitly by uttering (1'), (1'') or (1'''), or any of a number of similar sentences, *B* exercising his inferential, pragmatic abilities and exploiting facts about the context of utterance, beliefs about the speaker and so on.

Now of course the philosopher who recognizes utterance ellipsis *might* consider making the following (improbable) psychological claim: when *A* intends an utterance of (1) to be understood as expressing what he might have expressed more explicitly by, say, an utterance of (1'), this is because *A* first considered the longer sentence but decided to deliver a more economical one. But even if, as seems unlikely, a philosopher were to go along with this, his position would still have no bearing whatsoever on syntactic theory. He would be making no relevant claim about the *syntax* of the sentence actually uttered.

14. Strawson (1950)

One of the virtues of Strawson's 'On Referring' is its insistence on the position (also insisted upon by Grice) that *referring* and *saying* (or stating) are primarily things that *speakers* (rather than words) do.¹¹⁴ I am highly sympathetic to this stance. With Grice and Strawson, I see utterer's meaning (a notion to be explicated in terms of utterer's intentions) as conceptually prior to linguistic meaning, and I see speaker's reference (a notion also to be explicated in terms of utterer's intentions) as conceptually prior to the notion of linguistic reference (if such a notion is needed, which is far from clear).

Whilst 'On Referring' has many virtues, it has some serious vices—as Paul Grice once put it, 'a brilliant paper marred only by a discussion of Russell's Theory of Descriptions'.¹¹⁵ Strawson's failure to grasp a fundamental assumption underlying Russell's proposals is the matter I want to focus on here, for it is this failure that leads Strawson (1954) to a quite ridiculous reading of Sellars (1954). The matter concerns the distinction between sentences and utterances of sentences, and the bearers of truth and falsity.¹¹⁶ Russell, according to Strawson (1950: 325), fails at crucial points to distinguish the following:

- (A1) a sentence,
- (A2) a use of a sentence,

¹¹⁴ For detailed discussion, see the collections of papers by Grice (1989) and Strawson (1971), and Neale (forthcoming *b*). Every now and then a paper that explicitly discusses Grice's notion of what is said and claims to endorse it gets even this basic point wrong, talking relentlessly about what is said by a sentence relative to a context.

¹¹⁵ The major vices in Strawson's positive proposals are discussed in ch 2 of *Descriptions*, where plenty of references of earlier discussions can also be found.

¹¹⁶ This particular vice in Strawson's portrayal of Russell is discussed in ch 2 of *Descriptions*. See also Evans (1982). It appears to have been overlooked by some recent commentators.

(A3) an utterance of a sentence.

Similarly (B1)-(B3), in which 'sentence' is replaced by 'description'. And this failure, he says, is one of the things that leads Russell into error.

In order to forestall confusion and a silly defence of Strawson's reply to Sellars, let us be quite clear what Strawson intends by (A1) and (A3). *Sentences*, in the sense of (A1), are sentence-*types*; this Strawson makes clear when he says that 'one and the same sentence [may be] uttered on . . . various occasions' (1950: 325). Strawson is also clear about *utterances*, in the sense of (A3): 'two men who simultaneously uttered the sentence ['the king of France is wise'] in the reign of Louis XIV made two different utterances of the same sentence' (1950: 326).¹¹⁷

In summary, the same sentence (e.g., 'the present king of France is bald') may be used on one occasion to say something true, on another occasion to say something false. Thus sentences are not the sorts of things that are true (or false). If it is insisted that we need *objects* to be bearers of truth (falsity)—rather than just saying that *what the speaker said* was true (false), we can take particular dated utterances, sentences-relative-to-occasions of utterance, the propositions expressed by particular dated utterances, or what have you. There are many interesting discussions of this matter in the literature; in order to provide some continuity with the rest of the present article and avoid entanglements with Strawson's use of 'use', let us say that particular utterances are the sorts of things that are true or false.

Strawson is quite right, of course, to distinguish *sentences* from *utterances*, and to distinguish between the *linguistic meaning* of a particular sentence and *what someone says* by uttering it on a given occasion, if only because of indexical expressions like 'I', 'this', 'present', 'here', and so on. But as Sellars (1954), Russell (1959), and Grice (1970) all point out in their commentaries on Strawson, the question of how to treat indexicals is, in the first instance at least, quite distinct from the question of how to treat descriptions.

¹¹⁷ Notoriously, there is some latitude in the interpretation of (A2), but I suspect it is not harmful to Strawson's overall position. The remark just quoted is the first half a longer remark: 'two men who simultaneously uttered the sentence ["the king of France is wise"] in the reign of Louis XIV made two different utterances of the same sentence, though they made the same use of the sentence.' (1950: 326). He goes on to say of 'two men who uttered the sentence, one in the reign of Louis XV and one in the reign of Louis XIV' that 'each made a different use of the same sentence; whereas . . . two men who uttered the sentence simultaneously in the reign of Louis XIV, made the same use of the same sentence' (1950: 325-6). Although there is enough here to illustrate to the charitable reader what Strawson has in mind by 'use' in (A2), taken in an uncharitably strict way it does leave an open parameter: Is it just the fact that in both cases it is Louis XIV the men are talking about that guarantees sameness of use here? Or is the simultaneity of the utterances also playing a rôle? It should be clear that simultaneity *per se* is not really the issue, it is just something Strawson invokes so as not to get sidetracked into talking about 'is'. That is, as far as sameness of use of the expression 'the king of France' is concerned, it is not simultaneity that is relevant but who it is being used to talk about, so two men who utter 'the king of France is wise' five minutes apart in the middle of the reign of Louis XIV, make the same use of 'the king of France' for Strawson, even if his view of 'is' (and tense more generally) requires him to say that they did not strictly make the same use of the same *sentence* because they referred to different moments (I have no idea whether Strawson's view of 'is' (and tense more generally) does require him to say this; but in any case the matter is not remotely relevant to anything I say in this paper).

Of course definite descriptions may *contain* indexical components ('the *present* king of France', 'the man who gave me *this*', etc.); but all this means is that there are descriptions to which the Theory of Descriptions *and* a theory of indexicality apply. No substantive issue turns on Russell's failure to separate sentences from utterances explicitly when talking about descriptions. As Evans (1982) stresses, Russell is so very obviously concerned with the proposition expressed by a particular utterance—rather than the more abstract notion of the linguistic meaning of sentence-types—that it is very difficult to lend any sort of sympathetic ear to Strawson on this point. If Russell were pushed into being more precise, he would say that only in the following sense is the *sentence* 'The king of France is bald' equivalent to the *sentence* 'there is one and only one king of France and everything that is king of France is bald': the *proposition expressed* by a particular dated utterance of the former is equivalent to the proposition that would be expressed, in the same context, by a particular dated utterance of the latter.¹¹⁸

15. Sellars (1954)

The distinction between a sentence and a particular dated utterance of that sentence, so rightly stressed by Strawson, is respected and taken very seriously by Sellars (1954) in his criticisms of Strawson. A passage quoted earlier contains much that is important and sets the scene nicely for an examination of more of the text. Concerning the *sentence* 'the table is large,' he says,

a given *utterance* [*my italics*] of it is elliptical and states what would be nonelliptically stated, for example, by 'The table *over here* is large'...the context functions to give the statement the force, for example, of 'The table *over here* is large' (1954: 200).

There are three matters we need to focus on in connection with this passage:

(i) Like Quine before him, Sellars is *not* using 'elliptical' as a predicate of sentences. This facts jumps straight out in Sellars's wording: he explicitly uses 'elliptical' as a predicate of *utterances*. If Smith and Jones know that there are several As, an utterance of the 'the A' made by Smith on Monday might be elliptical, whilst an utterance made by Jones on Friday might be non-elliptical because Smith and Jones both know (and know that the other knows) that every A but one was destroyed in a fire on Wednesday. The moral, drawn earlier, is this: wherever we find talk of 'incomplete descriptions' in the literature, if we are really interested in meaning and communication and not merely in playing with formal systems, we would do well to take as much care as Sellars by construing it as talk of incomplete *utterances* of descriptions. And, of course, when we

¹¹⁸ Strawson's positive proposals about descriptions in 'On Referring' and later works have taken quite a drubbing in the literature, particularly at the hands of Sellars (1954). For a summary of some the problems and relevant references, see *Descriptions*, pp. 24-28 and pp. 53-56, nn. 22-29. Strawson's positive proposals are not my concern here, so I restrict my attention to his misreadings of Russell and Sellars.

discuss Sellars's use of 'ellipsis' we must always bear in mind that he is explicitly talking of *utterances* being elliptical and *not* of sentences being elliptical. (Notice the examples of utterance ellipsis in Sellars's own prose above.)

(ii) Sellars's way of stating his main point is explicitly *modal*. Indeed, the modal character jumps straight out in the passage just quoted: the elliptical utterance states what *would be* nonelliptically stated in the same context by an utterance of the richer sentence

(1) the table over here is large.

(iii) Sellars suggests an elliptical utterance can be used to express what would be expressed, in the same circumstances, by a complete utterance of a sentence whose words include *indexical* ('ego-centric') words not present in the words of the incomplete utterance.¹¹⁹

Given (i)-(iii), it is abundantly clear that Sellars is not primarily interested in *sentence* ellipsis in his discussion of descriptions. And it is equally clear that he is not claiming that words or phrases present at some level of syntactic representation are *deleted* by some sort of transformational operation in the course of deriving surface syntax.¹²⁰

Now there is a feature of Sellars's discussion that might today mislead those unable to see issues in the philosophy of language (even those with a long history) except through the lens of generative linguistics. (One should often *try* to see the issues through the lens of generative linguistics, in my view; my complaint is with those who can see the issues *only* through that lens.) In preparing the ground, Sellars starts out with a very simple and obvious example of utterance ellipsis, one that is straightforwardly explicable in terms of *sentence* ellipsis. That is, he discusses an example that involves the utterance of something that might be thought to be less than a whole sentence:

Jones: Seven is divisible by three.

Smith: Seven is not divisible by three.

Jones: Seven *is*.

¹¹⁹ This idea can also be found in Reichenbach (1947). Talking of the description 'the train' as it occurs in 'the train will arrive at 7 P.M.' Reichenbach says, 'The necessary addition then is understood. It usually consists in a reference to a preceding utterance; for instance, it may be assumed in the form 'the train of which we spoke'.' (1947: 258). This meshes nicely with Reichenbach's general picture of token-reflexivity. Reichenbach goes on to make some very instructive remarks about the relationship between descriptions and demonstratives.

¹²⁰ The idea of levels of syntactic representation related by transformational rules did not hit mainstream linguistics until at least 1957 when Chomsky's book *Syntactic Structures* was published, and it was not until the mid 1960's that Chomsky's ideas about syntactic structure were felt strongly in philosophy departments. Chomsky's long manuscript *The Logical Structure of Linguistic Theory* had been floating around since 1955 but was not published until much later, by which time *Syntactic Structures* (1957) and *Aspects of the Theory of Syntax* (1965), as well as a number of Chomsky's research papers, were already quite well known to philosophers. The idea that Sellars had in mind syntactic deletion in 1954 in connection with the sentence 'the table is large' is preposterous and hopelessly anachronistic. Similarly, Quine (1940) and Reichenbach (1947).

And it is quite clear he might have made his point just as well using a simple case of VP ellipsis, in the sense discussed earlier, without any heavy stress on ‘is’:

- (2) seven is not divisible by three, but six is.

According to Sellars, what someone says by uttering ‘seven is’—or, as he tends to put it, what an utterance of ‘seven is’ says—is a function of context, but not in the same way that what someone says by uttering ‘this is red’ is:

Correctly made *utterances* [*my italics*] of the latter sentence are *complete* even though they say what they do by virtue of their context. A context of a certain kind (e.g., the occurrence of a pointing gesture) is part and parcel of the grammar of the referring word ‘this’. On the other hand, ... *utterances* [*my italics*] of ‘Seven is’ are as such *not* complete and are only made complete by the context in which they are uttered. That the context serves in this way to complete them is as much a matter of linguistic convention as is the role of ego-centric expressions. Let us call this type of ambiguity *ellipsis* and say that in ellipsis the context completes the *utterance* [*my italics*] and enables it to say something which it otherwise would not, different contexts enabling it to say different things. (1954: 200).

Now this may not be a particularly sophisticated way of making the point—isn’t it more natural to talk about *utterers* rather than utterances saying things? shouldn’t the rôles of linguistic and non-linguistic context be distinguished?—but Sellars’s basic point is clear enough.

Once Sellars has made the general point that some utterances are obviously elliptical, he moves on to utterances of sentences containing the description ‘the table’, consistently continuing to apply the word ‘elliptical’ to utterances: ‘A given *utterance* [*my italics*] of [“the table is large”] is elliptical and states what would be nonelliptically stated, for example, by “The table *over here* is large.”’ (1954: 176). Surely no-one could construe Sellars as believing or claiming that an utterance of

- (3) the table is large

constitutes an utterance of something that is not a whole sentence. (3) is just as much a well-formed and complete sentence as ‘the table over here is large’ or ‘the table beside me is large’. Again, we see why Sellars is so careful to distinguish sentences from utterances: he is concerned with elliptical *utterances*, in the sense he characterizes so clearly in modal terms. ‘Seven is’ may well be an elliptical *sentence*, but that has no bearing on the matter at hand: elliptical utterances of sentences containing definite descriptions like ‘the table.’ (The example of utterances of the elliptical *sentence* ‘seven is’ is used only to get the ball rolling; it requires a monumental act of silliness to see Sellars as claiming that ‘the table is large’ is an elliptical *sentence*. Utterances of both ‘seven is’ and ‘the table is large’ are elliptical, the former as a matter of linguistic convention (because it is an elliptical *sentence*), the latter because of the way the world is, or is at least perceived to be by A.)

Once the distinction between the speech act notion of an elliptical utterance and the syntactic notion of an elliptical sentence is clear (as it should have been all along), it is

easy enough to distinguish the rôles of speech act context and syntactic context in resolving utterance ellipsis (and in the case of the latter, also sentence ellipsis). In a case of an utterance of ‘seven is’, *B* is alerted to the utterance ellipsis by (at least) the sentence ellipsis. In the case of an utterance of ‘the table is large’, by contrast, there is no sentence ellipsis. *B* is tipped off by an obvious *non-linguistic* fact: the world contains more than one table. And it is *B*’s knowledge of this fact—rather than a grammatical fact—that leads him to exploit context in the search for a complete interpretation of the utterance.

16. Strawson (1954)

In a 1954 reply to Sellars’s paper, Strawson’s original misunderstanding of Russell is transposed into a misunderstanding of Sellars, leading to a single misguided criticism. Strawson broods on the obvious difference between the *sentences* (or at least *expressions*) ‘seven is’ and ‘the table is large’, then draws all the wrong conclusions about Sellars’s discussion because he misses the fact that Sellars is concerned with the elliptical and incomplete nature of certain *utterances*.

Concerning ‘seven is’ or ‘James is’ (uttered in response to the question, ‘Who is going to drive?’), Strawson says ‘it seems very reasonable to call [these] *sentences* [*my italics*] ... incomplete or elliptical’ (1954: 223). He then goes on to say,

If one had to justify these phrases, I think one would be inclined to say that the *sentences* [*my italics*] were *formally, linguistically* deficient, that they did not come up to a certain standard of how a nonconversational English *sentence* [*my italics*] should be composed; one would point out that in their conversational setting, the deficiency is remedied by the *linguistic* context, that the surrounding *remarks* supply the missing words (1954: 223).

This seems basically correct as far as *sentences* are concerned, although one can well imagine cases in which a gesture or some salient feature of *non-linguistic* context remedies a linguistic deficiency.¹²¹ But, to repeat, *Sellars is not talking about sentence ellipsis*, and Strawson’s failure to appreciate this leads him into a sequence of errors:

But Sellars next suggestion, I find utterly puzzling. For he says that such a *sentence* [*my italics* SN] as ‘The table is large’ is incomplete or elliptical in the same sense as [‘seven is’] (1954: 223).

¹²¹ See the discussion of, for example, Nausicaa’s farewell (*Odyssey* viii. 461-2) in Hahnemann and Neale (forthcoming). When confronted with textual anomalies in Greek epic poetry and classical drama, scholars have tended to effect ‘repairs’—wishful emendations, *ad hoc* interpolations, and strained dependencies abound—to make offending lines conform to the exigencies of Greek syntax or metre. In a range of cases the case for repair is weakened considerably once the performative aspect of both genres is appreciated. Gestures produced by rhapsodes and actors can supply vital information not found in the words they accompany, and these together with intonation contour, metrical embellishment, and tone of voice may convey information to an audience that is simply unavailable to the reader. On this account, the performance of epic poetry must have been closer to that of drama than is often assumed, the rhapsode supplying elements without which certain parts of the performance involving direct speech would be incomplete and at times unintelligible.

Now according to Sellars, if ... I were to replace the words 'This pond' by the words 'the pond' [in 'This pond is used by children for sailing boats on'] I should be replacing a complete and non-elliptical *sentence* [my italics] by an incomplete and elliptical *sentence* [my italics]! (1954: 223).

In general, Sellars thinks that a *sentence* [my italics SN] containing a singular 'the'-phrase can be rendered nonelliptical only by supplementing the 'the'-phrase by some phrase containing what he calls an ego-centric expression (1954: 223)

Strawson has missed the central point of Sellars's discussion in spectacular style because he unremittingly overlooks his own distinction between sentences and utterances, and so his 'objection' to Sellars—which amounts to nothing more than the (correct) observation that the word 'elliptical' does not seem to apply to the *sentence* 'the table is large'—is stillborn.

17. Repetition

In a recent discussion of Strawson's response to Sellars, Stanley (2002a) repeats Strawson's simple error. 'With characteristic incisiveness,' says Stanley, 'Peter Strawson in his reply to Sellars in the very same journal issue . . . gives a clear explanation of the oddity of both Sellars claim and his vocabulary.' (2002a: 155). Anyone familiar with the Sellars-Strawson exchange knows this can go nowhere. Stanley presents, as a single quotation, virtually two whole pages (pp. 222-24) of Strawson's reply, the very pages that are replete with the error already discussed. Once he has finished quoting, Stanley says, 'All of Strawson's point are correct. First, the terms 'elliptical' and 'incomplete' only seem apt when the phenomenon involves words that are supplied to incomplete *sentences* [my italics] by linguistic context' (2002a: 156). If Sellars were talking about *sentence* ellipsis who could possibly disagree!¹²² Stanley goes on

Second . . . once one distinguishes between the kind of processes operative in resolving the reference of occurrences of demonstratives, and the kind of processes operative in resolving elliptical *sentences* [my italics] (such as in Sellars' example 'seven is' and Strawson's example 'James is') it seems bizarre to claim, as Sellars does, that the processes at work in resolving the context-dependence of 'The table is large' are the latter rather than the former (2002a: 156).

Indeed it would be bizarre to claim it, and presumably this is why Sellars does *not* claim it.

¹²² Actually, Stanley makes a mess of even this truism when analyzing Strawson's comments. He says, [Strawson] states that Sellars is justified in using the term 'elliptical' for *utterances* [my italics] of sentences like 'seven is' and 'James is', where the surrounding linguistic context 'supplies the missing words' to the incomplete *sentences* [my italics], because in the ordinary sense of the word 'elliptical', an *utterance* [my italics] is elliptical just in case its surrounding linguistic context quite literally supplies the missing words. This point is about vocabulary, the *proper* [my italics] use of the term 'ellipsis'. (2002a: 156).

When Stanley examines Sellars's text himself, rather than relying on the long quotation exemplifying Strawson's 'characteristic incisiveness' and 'clear explanation', he fails to see that the passages he has carefully selected for quotation are *precisely* the ones in which Sellars makes it clear he is talking about *utterance* ellipsis:

As Sellars writes, '... each of these dialogues contains an *utterance* [my italics SN] of 'Seven is'; and it is clear that what is communicated by these *utterances* [my italics SN] is a function of the contexts in which they are uttered' (2002a: 154)

According to Sellars, 'Correctly made *utterances* [my italics SN] of ['This is red'] are *complete* ... On the other hand, the two *utterances* [my italics SN] of 'Seven is' are as such not complete' (2002a: 154).

Once the quotation marks are off, Stanley veers all over the road:

Sellars then proceeds to argue that *sentences* [my italics SN] containing 'incomplete' definite descriptions such as 'The table is large', suffer (*sic.*) from . . . what he calls 'ellipsis'. [*Footnote omitted.*] That is, he claims that a *sentence* [my italics SN] such as 'The table is large' is incomplete in *just the same way* as Jones' *utterances* [my italics SN] of 'Seven is.' [*Footnote omitted.*] According to Sellars, an *utterance* [my italics SN] is elliptical for a *sentence* [my italics SN] like 'The table over there is large', in just the same way as Jones' first *utterance* [my italics SN] of 'Seven is' is elliptical for [the sentence? SN] 'Seven is divisible by three'. Sellars claim here is startling, as is his use of 'ellipsis' to describe the phenomenon of incomplete descriptions' (2002a: 154-5).

The only thing that is 'startling' is Stanley's slipping and sliding between 'sentence' and 'utterance'.¹²³

Stanley berates Sellars's use of 'elliptical', failing—as he must, since he is blindly following Strawson—to realize that Sellars is talking about *utterances* not sentences. And even if Stanley were able to demonstrate conclusively that some *other* feature of Sellars's overall philosophy renders *Sellars's* use of 'elliptical' *incoherent* (rather than merely different from the one he and Strawson mistakenly read into the paper), this would have no bearing whatsoever on the very clear and quite reasonable uses of the same word by Quine (1940), Kripke (1977), Brinton (1977), Salmon (1982) and Bach (1981), by me in *Descriptions*, by any of the people Stanley claims 'uncritically accepts Neale's vocabulary of 'explicit' and 'implicit', or by anyone else who wishes to talk very naturally about elliptical speech.¹²⁴

The motivation for Stanley's (2002a) digression on Sellars and Strawson is, perhaps, not entirely transparent.¹²⁵ He overtly seeks (i) to berate my mention of Sellars's (1954)

¹²³ For further examples of this, see the passage from Stanley quoted in the previous footnote.

¹²⁴ Perhaps Stanley thinks Sellars had no right (in 1954) to use the word 'elliptical' in the way he did because *twenty years later* it would be used in a narrower way in technical work by generative linguists—presumably such a complaint would apply equally to Quine's use of 'elliptical' in 1940. By these dim lights, then by parity of reasoning Quine (1960) and Geach (1962) had no right to use the ordinary word 'pronoun' in the way they did because *a decade or so later* it would be used in a narrower way in technical work by generative linguists so as to exclude the reflexive pronouns 'himself' and 'herself'? ('anaphors' in the sense of Chomsky (1981)).

¹²⁵ I am grateful to Robyn Carston for pointing this out to me.

paper in *Descriptions* when describing one of the two traditional responses to the problem posed by incomplete descriptions (the ‘ellipsis’ or ‘explicit’ response), and (ii) to berate and silence those who have used or accepted my innocuous labels (‘explicit’ and ‘implicit’) for these two traditional responses. But a little digging reveals the digression is also meant (iii) to add weight to Stanley and Szabó’s (2000b) response to criticism of their (2000a) spectacular misrepresentation of the explicit approach to incompleteness. Here is what Stanley says:

Sellars discussion has engendered much confusion in subsequent literature on context-dependence. Sellars’ paper has been given new life thanks to the influential discussion of ‘incomplete’ definite descriptions in Neale (1990, pp. 95ff.), from which much work on the topic in the subsequent decade borrowed its vocabulary. Neale’s discussion takes off from an uncritical acceptance of Sellars’ vocabulary and distinctions. But Sellars vocabulary and distinctions suffer from ambiguities, unclarities, and errors, which, because of Neale’s discussion, continue to disrupt the literature on the topic today. [Footnote 10: In particular, every paper on ‘incomplete’ definite descriptions that uncritically accepts Neale’s vocabulary of ‘explicit’ and ‘implicit’ approaches to context-dependence must be examined to see whether it is vitiated by Neale’s reliance on Sellars’ discussion] (2002a: 157).

It should be clear from all I have said already that the confusion engendered by Sellars’s discussion is all Stanley’s, the disruption to the literature entirely his doing, and the errors all his. But let us dig a little deeper as doing so will cast everything in a clearer light and provide some perspective on the explicit approach.

18. Misrepresentation

Stanley and Szabo (2000a) make an interesting positive proposal in connection with incompleteness. It can be stated thus: the incompleteness associated with the utterance of a DP is to be explained on the assumption that it contains an aphonic, indexical, domain variable assigned a value ‘by context’. The overarching idea is a familiar one, discussed by Westerståhl (1985), for example, but Stanley and Szabó add a novel twist: although the variable is syntactically real, it is not attached to, dominated by, or associated with either of the quantificational nodes, D (‘the’) or DP (‘the table’), in ‘the table’ as one might have thought; rather, it is associated with the nominal expression N (‘table’); as they put it, the variable ‘cohabits’ the node with ‘table’. We might call this a syntactic proposal with semantic import or a semantic proposal implemented syntactically, it doesn’t matter. What is key, however, is that it has a very clear *syntactic* dimension.

DP	
D	N
the	<table, e>

Strictly speaking, this approach to incompleteness is a version of neither the pure explicit nor the pure implicit approach, although it draws upon features of both.¹²⁶ From the implicit approach it borrows the idea of invoking quantifier domain restrictions.¹²⁷ What makes it different from the implicit approach is that instead of implicit, background domain restrictions, Stanley and Szabó want *foreground* domain restrictions, i.e. they want them to be part of the proposition expressed. And on this score, their approach is in harmony with the *explicit* approach. Finally, it differs from both approaches in having a crucial (and controversial) *syntactic* dimension: an N node in a phrase marker is not ‘inhabited’ by just a noun: it is ‘cohabited’ by a noun *and an aphonic domain variable*. In my opinion, this proposal faces serious syntactic and semantic obstacles, but that is a topic for another occasion.¹²⁸ My purpose here is to scrutinize what Stanley and Szabó say about the explicit approach in the build up to presenting their own proposal.

The proposal appears very late in their paper, after a drawn out assault on what they see as competitors to the proposal they are going to present. Indeed most of the paper appears to be an attempt to blast the logical terrain so hard that when their theory finally emerges it will face no competition. Two related fantasies infect their discussion. First they identify the explicit response with an absurd *syntactic* thesis. Second, they claim (i) to provide an ‘exhaustive’ and ‘comprehensive survey of the space of possible analyses’ (2000a: 219), (ii) to bring to bear ‘considerations which militate against all but our own proposal’ (2000a: 219), and (iii) to ‘consider and reject ... the ‘explicit’ approach to quantifier domain restriction (*sic.*) discussed ... by Stephen Neale’ (2000a: 219). In short, they claim it is *impossible* for any theory except their own to succeed. The reason? Their theory, and theirs alone, can capture ‘implicit binding’. For many sentences X, they claim, the binding by an overt quantifier in X of an aphonic domain variable in X is both necessary and sufficient to capture the relevant facts, and their cohabitation theory is uniquely positioned to do the job.¹²⁹

It is a seriously incomplete discussion of what is possible that omits what is actual. And, as it turns out, (i)-(iii) are false because Stanley and Szabo never even ‘consider’ the explicit approach. What they actually ‘consider and reject’ is something quite different, a straw man they call the *syntactic ellipsis* approach.

¹²⁶ This is not a problem for the labels ‘implicit’ and ‘explicit’, which merely applied to two approaches to incompleteness I detected in the literature. The labels were intended to be neither exhaustive nor exclusive (I was not trying to bifurcate the ‘logical space’, to provide an exhaustive or comprehensive survey of the space of possible analyses, as some claim to do in this area).

¹²⁷ This is why I suggested in ‘On Being Explicit’ (2000) that it was a version of the implicit approach. That was actually a very bad way to put the point, however, because Stanley and Szabó’s approach differs from the implicit approach in an important way I am about to specify.

¹²⁸ See Neale (forthcoming, *b*)

¹²⁹ Pragmatists have accepted implicit binding for some time, particularly in discussions of donkey anaphora. See, for example, Evans (1977), Davies (1981), Neale (1990, 1993), and Ludlow and Neale (1991).

First, a minor (but telling) point. Strictly speaking, Stanley and Szabó are wrong to talk about ‘the “explicit” approach to *quantifier domain restriction* [my italics] discussed ... by Stephen Neale’. What I call the ‘implicit’ approach to incompleteness is the one that invokes quantifier domain restriction, however. What I called the ‘explicit’ approach was a response to the incompleteness problem that does *not* invoke quantifier domain restrictions. Evidently, Stanley and Szabó are using ‘quantifier domain restriction’ in a new way, as a label for the general phenomenon of incompleteness in connection with DPs. I won’t quibble with this.

Now to the main point. Before they launch their attack on what they claim is the explicit approach, Stanley and Szabó festoon it with a mad thesis I am confident no advocate has ever advanced: on the explicit approach, according to Stanley and Szabó, an utterance of a sentence containing an incomplete description (or other DP) is an utterance of a sentence that contains an ‘unarticulated portion’ (2000a: 232), a ‘covert expression that cannot be heard’ (2000a: 233), a ‘syntactic constituent that has no phonological manifestation’ (2000a: 233). This is truly what they say. And in a footnote appended to the last of these remarks (n. 17), they refer the reader to ‘the discussion of the “explicit” approach in Neale (1990), pp. 95ff., as well as the articles cited in [Neale’s] note 49, p. 115.’ My note attributes the explicit approach to Quine (1940), Grice (1981) and others, and the sentence to which the note is appended attributes it to Sellars (1954).

The alleged commitment of the explicit approach to syntactic constituents that ‘cannot be heard’ amounts to a thesis in generative grammar, Stanley and Szabo claim, a thesis to the effect that expressions present at some level of syntactic representation (LF presumably) are deleted by a syntactic operation in the mapping from to surface form (PF presumably). As they see things, the superficial syntactic structure of Sellars’s example (1) is meant to be derived syntactically from the underlying (1’):

- (1) [s_[DP the table] [v_P is large]]
 (1’) [s_[DP the table over there] [v_P is large]].¹³⁰

In their jargon, ‘over there’ is *uttered* but not *pronounced*.

This is extraordinary. To claim that philosophers who have explicitly followed Quine, Sellars and others in talking about elliptical utterances have been arguing for, or assuming, *syntactic deletion* is tantamount to claiming that these philosophers know just enough syntactic theory to know what syntactic deletions *are*, but not enough to know that deletions must be *recoverable*! Certainly no-one would ever charge Stanley and Szabó with such a thing, for they are far from ignorant about the relevant aspects of syntactic theory, including the recoverability condition. As they say of VP deletion, it is ‘a syntactic phenomenon, due to some sort of syntactic rule of reconstruction or copying, or PF deletion under a syntactic parallelism condition.’ (2000a: 226, n. 9).

¹³⁰ Actually the LF will have a little more structure as the description will have been raised. For the present point, we can abstract from this.

Yet they claim syntactic deletion is the heart of position *I* was discussing when I sketched the traditional explicit approach in *Descriptions*, the approach associated with Quine (1940), Sellars (1954) and company. This is pure fantasy. I never discussed such a position in *Descriptions*—the relevant passages are reproduced below—and for good reason: *Descriptions* was meant to be a serious (if boring) book in the philosophy of language, not a work seeking to get mileage out of demolishing straw men. To the best of my knowledge, *no-one* sympathetic to the explicit approach has ever claimed that when someone utters (1), he is uttering a sentence which contains an expression ('over there') that is *uttered* but not *pronounced*.

Quite why Stanley and Szabó confuse (or embroider) the explicit approach with an extraordinary thesis about *syntax* is one of life's mysteries: (i) Very clear warnings against precisely such a confusion (or embroidery) are voiced by Bach (1981, 1987, 1994)—whom Stanley and Szabó *cite* in various connections—and by Ostertag (1998, 1999); (ii) *not one* of the well known exponents of the explicit approach (see below) says anything to suggest it is a syntactic thesis or has a syntactic dimension.

The fact that Stanley and Szabó talk of 'the "explicit" approach . . . discussed, *for example* [*my italics*], by Stephen Neale' (2000a: 219), shows they have *one* thing right: I am not the *only* person to have discussed it. In their eyes, my discussion is, I suppose *representative*. Of course I was far from being the first person to discuss it, all I did was pin a *label* on it: Quine (1940), Reichenbach (1947), Sellars (1954), Vendler (1967), Donnellan (1968), Kripke (1977), Brinton (1977), Sainsbury (1979), Davies (1981), Grice (1981), Wettstein (1981), Salmon (1982), Evans (1982), Soames (1986), Blackburn (1988), and no doubt many others, discussed it years before I threw my hat into the ring.

Stanley and Szabó are making at least three mistakes when they attribute to the advocate of the explicit response a story about syntactic deletion: (1) they are overlooking the fact that the earliest advocates, Quine (1940), Reichenbach (1947), and Sellars (1954), proposed the explicit approach before ideas about levels of syntactic representation and transformational deletions had been formulated and absorbed; (2) they are ignoring the fact that an examination of the relevant literature reveals no advocate of the explicit approach—from Quine in 1940, to me in 1990—ever viewed it, or even *discussed* viewing it, as involving a process of syntactic deletion between levels of syntactic representation; and (3) they ignore the fact that anyone who knew about syntactic deletions or syntactic ellipses in generative grammar would ipso facto know that such deletions or ellipses must be *recoverable*.

Let us now turn to textual matters and the history of the debate after Strawson and Sellars, to works written after the rise of generative grammar.

(i) Donnellan (1968) says something that might be read as a tentative endorsement of the explicit approach in connection with at least some incomplete descriptions used *attributively*. Perhaps he knew about deletions in generative grammar at that time, but either way there is no indication that he is suggesting syntactic deletion when he says, in connection with these attributive cases, that 'The reply one is inclined to make on

Russell's behalf is that in the loose way of everyday speech the context is relied upon to supply further qualifications on the description to make it unique.' (1968: 204, n. 5).

Donnellan goes on to criticize the explicit response in connection with incomplete descriptions used *referentially*, but nowhere does he suggest he sees it as involving syntactic deletion.

(ii) Kripke's (1977) brief discussion of incomplete descriptions appears in a paper prepared from a transcript of a talk given at MIT in the early 1970s. Although sympathetic to a unitary Russellian account of descriptions, which he was defending primarily on methodological grounds from the sorts of arguments Donnellan had mounted, Kripke suggested that ultimately such an account might be compromised by referential uses of radically incomplete descriptions such as 'the table', which he saw the explicit approach as struggling to capture. His wording suggests very strongly that at the time of his MIT lecture, or at least at the time of editing the transcript, the type of utterance ellipsis that Quine and Sellars were discussing was an established part of the general background for discussions of Russell's Theory of Descriptions:

The considerations I have in mind have to do with the existence of 'improper' definite descriptions, such as 'the table,' where uniquely specifying conditions are not contained in the specification itself. Contrary to the Russellian picture, I doubt that such descriptions can always be regarded as elliptical with some uniquely specifying conditions added. (Kripke, 1977: 8).

There appears to be no intention on Kripke's part to view this sort of ellipsis as involving syntactic deletion of the sort discussed by MIT linguists.

(iii) Brinton (1977) says something that suggests utterance ellipsis was at the time widely regarded as part of the Russellian picture, if not something Russell himself had in mind:

Russell claims that many or most occurrences of definite descriptions must be regarded as *elliptical* for 'complete' descriptions or at least for descriptions which would be thought by their users to be 'complete'. The description to be analyzed, then, is often not the description actually used by the speaker, but some longer description. Analysis takes us beyond the speaker's actual utterance (1977: 402).¹³¹

Brinton's use of 'actually' and 'actual' give his summary a standard modal flavour: the description to be analyzed in an utterance is often 'not the description actually used by the speaker', he says, 'but some longer description'; we go 'beyond the speaker's actual utterance'; and the clear implication here is of a longer description that *could have been* uttered.¹³²

¹³¹ I have not myself come across any passage in which Russell explicitly claims that utterances of sentences containing definite descriptions may be elliptical for complete utterances. But Brinton is quite likely correct that Russell would have concurred (and perhaps even did concur) with Quine and Sellars.

¹³² As noted earlier, when people who are sensitive to the distinction between utterance ellipsis and syntactic ellipsis—as Brinton is—use expressions like 'elliptical description' or 'incomplete description', it is obvious they are using a bit of shorthand and are talking about elliptical (or incomplete) *utterances* of the syntactically non-elliptical (complete) descriptions like 'the table'.

(iv) Sainsbury (1979), using more or less the same language as Donnellan, says that where we have an incomplete description, 'we treat the context as supplying a further determination of the predicate involved in the description' (1979: 115). There appears to be no commitment to syntactic deletion in this remark (and Sainsbury assures me he did not have it in mind).

(v) Bach (1981) provides a very nice example of an incomplete description used attributively and invokes utterance ellipsis in his appraisal of the situation as well as a modal fact:

If the emcee of a quiz show announces, 'The winner gets a trip to Hawaii,' he is using 'the winner' attributively, since he presumably does not know (and certainly does not intend the audience to think he knows) who the winner will be. Obviously he takes 'the winner' as elliptical for 'the winner of the game about to be played here,' and had he used that description, he would have used it attributively. Having not used it, he intended the audience not to figure out which winner he had in mind but how the description he used was to be completed (1981: 224)

(Notice the unproblematic use of the modal 'had he used that description'.) Later Bach goes on to say that,

there are many sentences which are almost always used nonliterally as elliptical for other sentences. For example, 'Ed doesn't look tired, he is tired' would likely be used with a suppressed 'merely' before 'look' to be inferred by the hearer, since the speaker would not be stating that Ed does not look tired but is tired anyway. Similarly, if I say, 'I drink only Scotch,' I would be stating not that I drink nothing but Scotch but merely that the only liquor I drink is Scotch ... The phenomenon of elliptical speech is commonplace; indeed, it often seems stilted not to suppress words that can easily be inferred Using incomplete definite descriptions elliptically ... is just another case of this familiar phenomenon (1981: 238)

There is much in these passages that I find conducive; and it is clear that Bach's talk of the 'suppression' of words is not meant to be understood as a claim about syntactic deletion between levels of syntactic representation.¹³³

(vi) Davies (1981) makes explicit use of the modal 'as though', and there is no suggestion he is entertaining the idea of syntactic deletion:

descriptions which do not, as they stand, apply uniquely to any object, (even taking into account contextual limitation of the domain of quantification) are often used, and are interpreted *as though* [*my italics*] they contained further descriptive material (in the form of a restrictive relative clause) (1981: 162).

(vii) Grice (1981) discusses incomplete descriptions, in connection with negation and scope:

¹³³ Bach (1987) elaborates on these ideas a few years later and echoes Sellars's (1954) point that someone using 'the table' is likely using it as 'short for some complete description like 'the table nearby', 'the table previously mentioned', 'the table I am pointing to, [*footnote omitted*] or something of the form 'the table which is G'' (p. 105). Although Bach is here discussing a version of the explicit approach to incompleteness, I should stress that on his account completion does not take place at the level of what is said, but at the level of what is meant.

Consider utterances of such a sentence as *The book on the table is not open*. As there are, obviously, many books on tables in the world, if we are to treat such a sentence as being of the form *The F is not G* and as being, on that account, ripe for Russellian expansion, we might do well to treat it as exemplifying the more specific form *The F which is ϕ is not G*, where ' ϕ ' represents an epithet to be identified in a particular context of utterance (' ϕ ' being a sort of quasi-demonstrative). Standardly, to identify the reference of ' ϕ ' for a particular utterance of *The book on the table is (not) open*, a hearer would proceed via the identification of a particular book as being a good *candidate* for being the book meant, and would identify the reference of ' ϕ ' by finding in the candidate a feature, for example, that of being in this room, which could be used to yield a composite epithet ('book on the table in this room'), which would in turn fill the bill of being an epithet which the speaker had in mind as being uniquely satisfied by the book selected as candidate. If the hearer fails to find a suitable reference for ' ϕ ' in relation to the selected candidate, then he would, normally, seek another candidate. So determining the reference of ' ϕ ' would standardly involve determining what feature the speaker might have in mind as being uniquely instantiated by an actual object, and this in turn would standardly involve satisfying oneself that some particular feature actually is uniquely satisfied by a particular actual object (e.g. a particular book). (1981: 276-77)

This is a rich passage I shall draw upon later. For immediate purposes, however, it is enough to note that Grice nowhere talks of 'the *F*' being obtained from 'the *F* which is ϕ ' by syntactic deletion.

(vii) Evans (1982) seems to follow Donnellan in connection with incomplete descriptions used attributively, providing the following example to show how the explicit approach works:

travelling in a car through the United states, I might pass through a town whose roads are particularly bumpy, and in consequence say 'They ought to impeach the mayor.' I do not intend my audience to identify the object spoken about as one of which he has information; I intend merely that he take me to be saying that the mayor of this town, through which we are passing, ought to be impeached, and this statement is adequately represented quantificationally (1982: 324).

This is a nice, straightforward statement of the unadorned explicit approach. There is no mention here or anywhere else in Evans's discussions of descriptions or in his lengthy discussions of syntactic theory, of deriving incomplete descriptions from complete descriptions by syntactic deletion.

(viii) Salmon (1982) does not seem to have syntactic deletion in mind when he says that on the explicit approach an incomplete description is 'elliptical for some more fully descriptive phrase to be supplied by presumed shared background assumptions in the context of use, or something similar' (p. 39) 'elliptical for some more fully descriptive phrases floating in reach just overhead' (p. 44).

(ix) Nor does Soames (1986) appear to have syntactic deletion in mind either when he talks of the explicit (or 'traditional') approach involving 'contextual supplementation' and of the hearer needing to 'extract content from context' (1986: 278 and 301, n. 7).

(x) Bach (1987) explicitly warns against viewing the elliptical/explicit approach as involving syntactic deletion (pp. 73-74), and provides the following elegant summary of the situation, before presenting his own theory:

We much prefer not to make fully explicit what we mean when that can easily be inferred, and, indeed, speech is stilted when words that can easily be filled in are not suppressed. Using an incomplete definite description in place of a complete or a more elaborate one is just one case of this widespread phenomenon. So if a speaker were to utter ['the table is covered with dust'], surely he would not be asserting that one and only one table in the universe is covered with dust. He is likely to be using 'the table' as short for some complete description like 'the table nearby', 'the table previously mentioned', 'the table I am pointing to', [*footnote omitted*] or something of the form 'the table which is *G*'. Whenever a description 'the *F*' is used as short for a complete description 'the *F_c*', I will refer to latter as its COMPLETION, which the hearer is intended to identify as such. . . . if a speaker utters ['the table is covered with dust' he] is implicitly asserting what could be made explicit by ['the table I am pointing to is covered with dust'] (1987: 104-5)

It is clear from this passage and the pre-emptive warning that Bach does not think this description of the facts presupposes syntactic deletion or syntactic ellipsis.

(xi) Blackburn (1988) discusses incomplete descriptions very fruitfully and with great clarity.¹³⁴ He suggests that Russell's (1959) response to Strawson (1950) might be defended in the face of an example involving an incomplete description such as 'the table', by saying that the speaker uses the incomplete description 'in place of some other description (perhaps 'the table over here')' (1988: 268), and notes that 'such a position has been defended by Wilfrid Sellars, among others' (1988: 268). Blackburn then goes on to discuss Sellars's account and calls it the 'traditional' account:

the view that when a speaker uses a sentence containing an incomplete definite description, the sentence is (usually) used elliptically for some other sentence containing a complete description. [*Note omitted*] When this occurs, what the speaker says is determined by the non-elliptical sentence, the sentence containing the complete description. . . . the position seems to provide an accurate account of what occurs in at least some cases. Consider, for example, the following exchange:

John: Who is the prime minister of this country?

Mary: The prime minister is Brian Mulroney

The traditional account provides a plausible account of what is going on here; it is natural to suppose that Mary is using the incomplete description 'the prime minister' elliptically for 'the prime minister of this country'. . . . an incomplete description is used elliptically in place of some complete description (1988: 268).¹³⁵

Blackburn nowhere suggests the traditional/elliptical/explicit account involves syntactic deletion or syntactic ellipsis. Indeed he is as clear as can be that on this account the idea

¹³⁴ The material in Blackburn's paper is based on material from his dissertation, *Reference and Descriptions* (University of Toronto, 1986). Had I known about Blackburn's dissertation and paper before I had written most of *Descriptions*—his paper was brought to my attention by Nathan Salmon after I had finished all but the footnotes—I would certainly have drawn on it. My conclusions would not have been different, but I would have been able to save myself some work by citing Blackburn's arguments and observations.

¹³⁵ As noted earlier, when people who are sensitive to the distinction between utterance ellipsis and syntactic ellipsis—as Blackburn is—use expressions like 'elliptical description' or 'incomplete description', it is obvious they are using a bit of shorthand and are talking about elliptical (or incomplete) *utterances* of the syntactically non-elliptical (complete) descriptions like 'the table' or 'the prime minister'.

involves using an incomplete description *in place of* a complete one, which is of course the core of the explicit approach.¹³⁶

We have reached 1990 and my own discussion (and labelling) of the explicit response in *Descriptions*. Did I say, or even imply, that I was construing those who had made the explicit response as endorsing, or even *discussing*, a theory that invoked a process of syntactic deletion? And did I suggest anywhere that I subscribed to such a view? Did I ever *mention* such a view? Most certainly not, on all counts.

First, *ex cathedra*: ‘explicit’ is *my* label, and I hereby explicitly state that the explicit response does not and never did involve or invoke syntactic deletion or syntactic ellipsis. Second, the idea is ridiculous and in any event violates recoverability—no philosopher taking generative grammar as seriously as Bach, Davies, Evans, Grice, Sainsbury, Soames and I did was ever going to proceed down that road. Third, as should be abundantly clear by now, there is nothing in the work of those people I explicitly mentioned as putting forward the explicit response—namely, Quine, Sellars, Donnellan, Sainsbury, Davies, and Evans—to indicate commitment to, or even contemplation of, the idea that it involved syntactic deletion. Fourth, everything of relevance I said about the matter in *Descriptions* is contained in the following passages, which, it will come as no shock to learn, make no reference to syntactic deletion:

- (1) On the explicit approach (taken by, e.g. Sellars (1954)), a particular utterance of ‘the table’ might be elliptical for ‘the table *over there*’. [Footnote: For similar suggestions see Quine (1940), Vendler (1967), Lewis (1973), Cresswell (1973), and Soames (1986).] (1990: 95).
- (2) On the explicit approach, the quantifier ‘everybody’ (as it is used on this occasion) is elliptical for ‘everybody at the dinner party I had last night’, or some such ‘narrower’ quantifier (1990: 95).¹³⁷

¹³⁶ Blackburn sees certain types of examples as providing difficulties for the traditional ellipsis approach, difficulties that can be overcome by a modification: where there are many competing completions, none of which can be said to determinately capture what the speaker intended, no single determinate proposition is expressed:

Rather, there is some more or less vague *class* of propositions, each determined by some expansion of the description—‘The *F* that is *H*’, ‘The *F* that is *K*’, and so on. (Contrary to what Russell suggests, it is probably a mistake to think that these descriptions are any sense present in the mind of the speaker. We can say instead that the speaker would be prepared to fall back on these descriptions if he were asked to identify which *F* he meant.) According to the modified version of the traditional account, when the speaker says ‘The *F* is *G*’ he asserts the truth of the propositions in the resulting class (or, perhaps, of some weighted majority of these). (1988: 271)

Stanley and Szabó (2000a) profess not to understand Blackburn here. (See below)

¹³⁷ The ‘occasion’ in question in (2) was an utterance of ‘everybody was sick’ made in response to a question about a dinner party I had last night. When I talked about syntactic ellipsis in *Descriptions* I was careful to talk of syntactic *deletion* rather than *ellipsis* in order to keep syntactic ellipsis and utterance ellipsis separate, for example in my brief (and useless) discussion of VP deletion on pp. 218–19. As noted several times already, when people who are sensitive to the distinction between utterance ellipsis and syntactic ellipsis use expressions like ‘elliptical description’, ‘elliptical quantifier’, ‘incomplete description’, and ‘incomplete quantifier’, obviously they are using a bit of shorthand and are talking about elliptical (incomplete) *utterances* of syntactically non-elliptical (complete) expressions like ‘the table’ or ‘everybody’.

(3) Clearly I do not mean to be asserting that everyone in existence was sick, just that everyone *at the dinner party I had last night* was. In some fashion or other, this is discernible from the context of utterance. [Footnote: This point is made by Quine (1940), Sellars (1954), Sainsbury (1979), Davies (1981), and Blackburn (1984).] (1990: 95).

(4) According to the explicit approach, incomplete quantifiers are *elliptical* for proper quantifiers. As Sellars puts it, the descriptive content is ‘completed’ by context (1990: 95);

(5) Let’s suppose that an incomplete description is elliptical for a proper (i.e., uniquely-denoting) description recoverable from the context of utterance (1990: 96).

(6) Evans suggests that he (the speaker) *would* be able to complete the description in a uniquely appropriate way, and supplies a plausible completion using referential rather than descriptive material. (1990: 100).

(7) Even the descriptive content of an overt description is not always fixed by purely linguistic factors (1990: 201).¹³⁸

(8) there is simply no difference between saying that there is an ‘implicit reference’ to the victim and saying that the incomplete description ‘the murderer’ is elliptical for a uniquely denoting description, such as ‘the murderer of *that man*’ (where ‘that man’ refers to the victim), or ‘the murderer of *him*’ (where ‘him’ refers to the victim), or ‘*his* murderer’ (where ‘his’ refers to the victim), all of which contribute the same thing to the proposition expressed, viz., the descriptive condition *murderer-of-b* where *b* is the victim himself rather than some description of *b*. Wettstein is just mistaken in claiming that ‘the murderer’ is ‘not elliptical for some Russellian description’ (*ibid*); the proposition expressed by (3) [‘the murderer is insane’] is given by (4):

(4) [the *x*: *x* murdered *b*] (*x* is insane).

It is unclear why Wettstein thinks the description in (4) is non-Russellian (1990: 99)

As Zvolensky (2000) points out, the last of these passages makes it quite clear I do *not* have a syntactic deletion account in mind. I was responding to Wettstein’s objection to the explicit approach that there are non-equivalent richer descriptions for which an attributive utterance of ‘the murderer’ is elliptical, from which no principled selection can be made. If I had intended a syntactic ellipsis theory, I could not have said what I said in (8) by way of responding to Wettstein. My point was that it was not the *words* in the richer description that mattered but *what the words contributed to the proposition expressed*, and that on that score there could be nothing to choose between, for example, ‘the murderer of *him*’ and ‘the murderer of *that man*’ for a direct reference theorist like Wettstein, ‘him’ and ‘that man’ being used solely to refer to the dead individual *b*, on direct reference accounts. My use of [the *x*: *x* murdered *b*] in the formalism giving the truth conditions of the utterance indicates clearly that I privileged no particular linguistic completion in the scenario in question, and that the formula itself was not an LF (LF was not introduced in any case until later in the book). Was I committing myself to direct

¹³⁸ The word ‘even’ introduces the last of these remarks, and the word ‘overt’ was italicized because I was, in that passage, commenting first on the fact the descriptive content of descriptive (i.e. D-type) pronouns is not always determined by purely linguistic factors.

reference here? No. I was committing myself to the position that the direct reference theorist can get no mileage out of on attributive use of ‘the murderer’ in the scenario at hand. Am I prepared to commit myself to direct reference? No. I am happy to accept that in this case and many others no one richer description is better than many others the speaker and hearer could easily drum up. Why people worry about this baffles me.

Perhaps Stanley and Szabó’s real target is not me but *more recent* advocates of the explicit response, those drawn to it *since* the emergence of the label ‘explicit’ in 1990. If so, then they should say so, name names, and not attribute syntactic deletion to me and my explicit ancestors. Other critics of the explicit approach such as Bach (1987), Reimer (1992, 1998a), Schiffer (1995), Bezuidenhout (1997), Ostertag (1998, 1999), Devitt and Sterelny (1999), and Devitt (1997, this volume) have not claimed to see it as involving syntactic deletion or syntactic ellipsis. Indeed, some of these critics—e.g. Bach and Ostertag—have explicitly stated that syntactic deletion or syntactic ellipsis is *not* part of the story.

Here is Ostertag (1998, 1999):

the notion of ellipsis appealed to is not the one familiar from syntactic theory . . . whatever the relationship is that holds between [a sentence containing an incomplete description] and [a sentence containing a completed one] it is not syntactic. (1998: 20).

There seems no reason to suppose that the relation between, say, [‘the table is covered with books’] and its completion is anything like that between: ‘John left and Bill did too’ and ‘John left and Bill left’. (1999: 144, n. 3).¹³⁹

One can only speculate why Stanley and Szabó identify the friendless syntactic deletion approach with ‘the “explicit” approach ... discussed, for example, by Stephen Neale’ (2000a: 219). One might, perhaps uncharitably, see them as desperate for a target, for shortly after their misrepresentation of the explicit response they provide arguments meant to refute it as part of a build-up to their own theory. But their arguments against the ‘explicit approach’ engage at most the antecedently implausible syntactic deletion theory—which they call ‘plausible’ (2000a: 232) before setting out their arguments against it.¹⁴⁰

¹³⁹ Students are warned that it would be a mistake to think of syntactic deletion as playing a role in a theory of incompleteness by Larson and Segal (1995) in their textbook *Knowledge of Meaning*. Taking care not to attribute such a view to anyone, Larson and Segal point out that such an approach would involve appealing to the idea of linguistic material that is ‘semantically significant but phonetically unpronounced’ (1995: 331), a locution that appears to foreshadow Stanley and Szabo’s talk of a ‘covert expression that cannot be heard’ (2000a: 233), and a ‘syntactic constituent that has no phonological manifestation’ (2000a: 233).

¹⁴⁰ This is not the only place where Stanley and Szabó construct the flimsiest of straw men. In one note (2000a: 232, n. 16) they consider what they call the ‘ambiguity approach’ to incompleteness according to which a noun such as ‘bottle’ is ambiguous, an infinite number of distinct lexical items ‘corresponding to it’, including ‘bottle₂₀₆’ which means the same as ‘bottle₁ Lisa just bought’ (where ‘bottle₁’ is ‘what we usually find as the first entry in any dictionary, and means what we would ordinarily expect’ (2000a: 232, n. 16)). Stanley and Szabo do not cite anyone as holding this ridiculous view; nonetheless they feel compelled to mount an argument against it.

It is unclear whether Stanley and Szabó (2000b) accept they earlier misrepresented the explicit approach, as there is much that is murky in their later discussion. They start out by saying,

According to Neale, we misinterpret the intended meaning of this phrase ['the explicit approach'] in the two places in which we use it. Of course nothing in our paper rests on this point. ... We should emphasize that, even after studying Neale's illuminating reply, we are unclear what he means by 'the explicit approach.' According to Neale's original explanation, 'incomplete quantifiers are elliptical for proper quantifiers' (1990, p. 95). The current explanation differs from the ones provided in (1990) by the addition of modal vocabulary; 'the matrix of a quantified NP *can* be interpreted, in context, as elliptical for a richer (complete) matrix that *could* have been used in its place (Neale, 2000) The purpose of the new formulation is meant (*sic.*) to bring out the modality Neale sees as implicit (*sic.*) in Sellars's use of the term 'elliptical'. But we have no clear conception of the intended interpretation of these additional modal expressions, and, as a result, no clear conception of 'the explicit approach' (2000b: 295.)

This is all glib. First frequency of the phrase 'explicit approach' in their paper is not the issue for they first *identify* the explicit approach with what they call the 'syntactic ellipsis' theory, then use their own label many times! (The label 'explicit approach' and my name also appear prominently in their short abstract.) Second, quite a lot turns on their misrepresentation for they claim in their original paper to provide an 'exhaustive' and 'comprehensive survey of the space of possible analyses' (2000a: 219) and to bring to bear 'considerations which militate against all but our own proposal' (2000a: 219). Third, it is simply false that I see modality as *implicit* in Sellars's use of the term 'elliptical': I see it and always have seen it as *explicit* in Sellars's *explicit* claim, that an elliptical utterance states 'what *would be* nonelliptically stated' in the same context by a complete utterance [*my italics*]. Stanley and Szabó omit to mention that in *Descriptions* I explicitly took Sellars's discussion as the *locus classicus* of the explicit approach; and as we have seen already, Sellars *explicitly* uses a modal characterization (as do many others). Fourth, given their own uses of analogous modals discussed earlier, how can Stanley and Szabó claim to have 'no clear conception of the intended interpretation' of such 'additional' modal expressions? (The 'additional' modal expressions they are referring to are just 'can' and 'could' in my statement that 'the matrix of a quantified DP can be interpreted, in context, as elliptical for a richer (complete) matrix that could have been used in its place' (2000: 286).) Fifth, after all of the passages from *Descriptions* quoted above, after all of the quotations from Quine, Sellars, Davies, and Evans above (all of which appear in *Descriptions*, and some of which are *explicitly* modal), and after the points I made explicitly in the short piece to which Stanley and Szabó (2000b) are replying, it takes an act of will to be 'unclear' about what I mean by 'the explicit approach'.¹⁴¹ Critics such as Reimer (1992, 1997), Bezuidenhout (1997), Ostertag (1998,

¹⁴¹ Stanley and Szabó deploy this rhetorical gambit elsewhere. In a note discussing Blackburn's (1988) 'modified traditional account' of incomplete descriptions, they correctly present Blackburn as suggesting that an utterance of a sentence containing an incomplete quantifier matrix expresses 'the set of propositions

1999), Schiffer (1995), Devitt (1997*a,b*, this volume), Devitt and Sterelny (1999), and others seem to have had no trouble whatsoever understanding it and generating some genuine problems.¹⁴²

19. Gödelian Completions

Trivially, no logical problem of incompleteness arises for a Gödelian description $(\lambda x)(Fx \bullet x=a)$, the sort used in spelling out Gödel's slingshot argument.¹⁴³ This suggests a fool-proof way of *regularly* interpreting utterances of incomplete descriptions used referentially? Consider the following example due to Schiffer (1995). Ferdinand Pergola is known by you and me to uniquely satisfy the matrices of many distinct descriptions ('the author of *Smells and Tickles*', 'the man staggering up to the podium in front of us', etc.); seeing Pergola trip, I say to you 'the guy is drunk'; you understand my remark. According to Schiffer, no *single* uniquely satisfied matrix, 'however complex' is salient enough to underlie both my intention and your comprehension.

I suggest *simple* rather than *complex*: you understand me as saying something whose truth conditions are captured by (1), where *a* stands for Pergola:

- (1) $[the_x: guy(x) \bullet x=a] x \text{ is drunk}$

The underlying idea is simple and natural: completion of an incomplete description used referentially is effected *non-descriptively* (as non-descriptively as possible, at any rate). It is something like this I believe Grice had in mind in the passage quoted earlier.

Before saying more, I want at least to mention a matter I cannot do justice to here because of a host of complex psychological and semantic issues it raises. One of Schiffer's main aims is to present a dilemma for a certain theory of propositional attitude reports—the so-called hidden indexical theory—and he goes about this by first presenting a dilemma for anyone wishing to hold a unitary Russellian account of descriptions in conjunction with a direct reference theory of simple indexical and demonstrative pronouns.¹⁴⁴ Schiffer asks us to suppose that I had uttered 'he is drunk' (rather than 'the

expressible by sentences which result from [that sentence] by expanding the quantifier phrase . . . into an expression that the speaker "would be prepared to fall back on". (2000a: 237, n. 23) and then add, 'We are not sure what is meant by this latter phrase' (ibid.). It is pretty obvious what Blackburn means, as Schiffer (1995) and other commentators recognize. An apparent problem for versions of the explicit approach like Blackburn's is mentioned by Brinton (1977): in a particular situation a speaker will often assent to alternative completions, which lead to complete matrices that happen to be true of different individuals. There is no evidence that this worry is implicated in Stanley and Szabo's professed uncertainty about what Blackburn means.

¹⁴² The rhetorical gambit again: Stanley and Szabó say, 'We avoided using Neale's phrase "the implicit approach", because we had no real grasp of it' (2000b: 295). Again, other critics appear to have had no problem. For insightful discussions of both the implicit and explicit approaches, see Reimer (1992, 1998a).

¹⁴³ See Neale (2001).

¹⁴⁴ An expression is directly referential, in Kaplan's (1989) sense iff its contribution to the proposition expressed is just the object to which it refers. There is no commitment in *Descriptions* to direct reference. Indeed, I talked of *object-dependent* (rather than *singular*) propositions precisely to counter any thought I was

guy is drunk') in the example involving Pergola. And he says the Russellian has no good basis for preferring a direct reference theory of my use of 'he' to a theory that treats it as an incomplete description with more or less the same content as my use of 'the guy'. This is because in the two cases under consideration there would be no discernible difference in my communicative intentions, and these intentions are the only psychological states relevant to determining what I would have meant by my utterance—what I would *say* is part of what I would mean, so it would be backed by a communicative intention. The Russellian must, it would seem, either (a) deny that my demonstrative use of 'he' is directly referential (and offer an alternative treatment, presumably description-based), or (b) deny the claim of indiscernible communicative intentions, or (c) deny that psychological facts alone determine the issue of what I meant. (a) might have some plausibility. My utterance of the demonstrative pronoun 'he' might be interpreted as equivalent to an utterance of a description we interpret as the Gödelian [$the_x: male\ x \bullet x=a$], where a refers to Pergola. Thus, a formal representation of the truth conditions of my utterance of 'he is drunk' might be given by (2):

(2) [$the_x: male\ x \bullet x=a$] x is drunk.

I am inclined to think this would be a more plausible line of defence than (b)—(c) is out of the question for the pragmatist—but more would need to be said about the nature of pronouns and about relationship between formal representations of truth conditions, the LF's of English sentences, and the thoughts we seek to convey.

The upshot of all this is a welcome rapprochement between those who have argued for a semantically significant distinction between Russellian and referential uses of descriptions and those who have argued for a unitary Russellian theory. In a sense, everyone was right and everyone was wrong. When a description is used referentially it is being used in such a way that there is an obvious Gödelian completion (whether the description actually needs one or not). On this view, if I use a description 'the ϕ ' referentially, to refer to John, then what I say entails that John is ϕ , which seems right. If, in Schiffer's example, I had said not 'the guy is drunk' but 'the idiot is drunk', then surely I am committing myself to Pergola's being an idiot, and any purely referential alternative is going to have to find some way of accommodating this.

The difference between what I called referential_D and referential_N uses in ch 3 of *Descriptions* can now be thought through in terms of the type of expression most plausibly substituting for a , if (2) were to be rendered explicitly in English, a simple demonstratives or a name. Do we now say that it is part of the *meaning* or *semantics* of 'the' that on one use it invites a Gödelian completion? Or should we see such completions as simply an abstraction over regular interpretations? I doubt it really matters

assuming direct reference. There is no problem in maintaining that a formal language (or mentalese, for that matter) contains directly referential terms while English itself does not. On my account, of course, where natural language is concerned it is *speakers* rather than expressions that refer. On the use of 'directly referential' and 'rigid' in connection with what a speaker refers to and says, see Neale (forthcoming, *b*).

how we think about this until we have some understanding of the cognitive mechanisms involved in interpretation, which is presumably a long way off. Either way, the philosophical point remains intact: descriptions are Russellian, and the phenomenon of referential usage is a special case of the phenomenon of incompleteness. I said ‘a special case’ not ‘just a case’; the case is special in the sense that it is highly regular, perhaps even conventional. Which brings me to the next point.

20. The Argument from Convention

Devitt (1997*a,b*, this volume) and Reimer (1998*a*) have presented an intuitive and powerful argument for an ambiguity in definite descriptions. I shall call it the Argument from Convention: referential uses of descriptions are common, standard, regular, systematic, and cross-linguistic; indeed so much so that it would be a bit rich to deny that such uses are *conventional*, a direct function of linguistic meaning in a way that referential uses of other quantified DPs are not.¹⁴⁵ I think Devitt and Reimer are right about this; but rather than undermining the unitary Russellian analysis *per se*, the Argument from Convention undermines only the standard, wooden, Gricean explanation of referential usage (like the one I sketched in *Descriptions*), which amounts to no more than a generalized conversational implicature story.¹⁴⁶ But the proposal on the table here is unruffled by the Argument from Convention. Indeed, the proposal seems to *explain* the convention in question, and this is why it seems to me a synthesis of everything that is right and nothing that is wrong in the old unitary Russellian theory and in semantical ambiguity theories that associate object-dependent (rather than object-independent) propositions with what is said when a description is used referentially. Moreover, it seems to do a much better job than either in yielding blueprints that place the right truth-conditional constraints on what is said, particularly where modalities are concerned.¹⁴⁷ In

¹⁴⁵ I do not want to put too much weight on the word ‘convention’ here. Perhaps ‘Argument from Regularity’ is a better label, but I have been using ‘Argument from Convention’ for some years (against Kent Bach’s advice) so I stick with it here. For an illuminating discussion of the pitfalls involving ‘conventional’, ‘standard’, ‘regular’ and so on see Bach (1995).

¹⁴⁶ All conversational implicatures must be calculable in Grice’s sense, even those he takes to be generalized. It is crystal clear from the first sentence of ‘Logic and Conversation’ onwards that the conversational implicatures associated with uses of the words ‘and’, ‘or’, ‘if’, ‘every’, ‘a’, ‘the’ etc.—i.e. those words corresponding to formal devices in logical theory—are generalized implicatures for Grice, the ones of philosophical importance, the ones that really bothered him (unlike the particularized ones which have no philosophical significance). It is equally clear that ‘generalized’ has no *theoretical* import for Grice in the context of his account of the properties an implicature must have if it is to count as conversational (hence the calculability requirement) and that generalized conversational implicatures are quite different from conventional implicatures. Grice is so very careful here, and it is odd that he has been misunderstood.

¹⁴⁷ One advantage it has over the original unitary Russellian account is that it captures a reading of (i), pointed out to me some years ago by Irene Heim:

(i) It is always the case that the player on the left wins these games.

this last respect it also differs also from the earlier pragmatist proposals of Récanati and Bezuidenhout which seek broad truth-conditional agreement with the ambiguity theorist; furthermore, unlike its pragmatist predecessors, it does not require us to view the *type* of proposition (object-dependent or object-independent) expressed as semantically underdetermined, an idea that has always seemed to me to carry the whiff of ambiguity. (Given the traditional relation between proposition types and logical form, the attempt to use LFs in capturing logical forms, and minimal constraints on the compositional character of LFs, it would seem that the proposals of Récanati and Bezuidenhout will not be able to avoid positing two distinct LFs.)

21. Demonstratives as Gödelian Indefinites

Some who have posited a semantically distinct, referential reading of ‘the ϕ ’, have been inclined to posit one for ‘a ϕ ’ too—but others have pointedly not done this. Let us suppose the following truth-theoretic axiom helps yield a blueprint for interpreting an utterance of a sentence of the form ‘a ϕ is ψ ’:

- (1) $[a_k: \phi]\psi$ is true of a sequence s iff ψ is true of at least one sequence that ϕ is true of that differs from s at most in the k^{th} position.

As noted earlier, the simplest cases exemplifying the problem of incompleteness arise naturally in connection with uses of $\uparrow 1$ determiners such as ‘the’, ‘every’ and ‘no’, whilst the most straightforward ‘derived’ or ‘inverted’ cases arise naturally in connection with uses of $\uparrow 1$ determiners such as ‘a’ and ‘some’. The need for enrichment in connection with indefinite descriptions is clear when they occur within the scope of negation:

- (2) It is not the case that a student lives in Chicago.

If we are discussing students in my seminar, I can say something true by uttering (2), despite the obvious fact that many students live in Chicago.

It is sometimes suggested that demonstrative descriptions ‘that ϕ ’ and ‘this ϕ ’ are special forms of definite descriptions. But the persistence test, discussed earlier, suggests that when used demonstratively they are more like *indefinite* descriptions. But what does the difference between demonstrative and indefinite descriptions consist in? The answer may well be that whereas indefinites do not signal uniqueness, demonstratives do, But doesn’t that make them like definites? Yes and no. The difference is in the way uniqueness is meant to be secured. Demonstratives are Gödelian *by nature*. An act of reference is signalled *as a matter of linguistic convention* by the use of ‘that’ (rather than

Traditional considerations governing scope suggest that the description may not take scope over the temporal operator, yet there would still appear to be a reading of (i) upon which the speaker could be saying of Smith (who is the player standing on the left) that he always wins.

‘a’ or ‘the’). In short, they are Gödelian indefinites; the blueprint for an utterance of ‘that ϕ is ψ ’ is basically one specified by (3):

$$(3) \quad [an_x: \phi(x) \bullet x=that] \psi(x).$$

This would appear to deliver, by a route Lepore and Ludwig (2000) may find objectionable, an analysis equivalent to one they have suggested, capturing the *demonstrative* and *descriptive* features (or conventions governing the use of) demonstrative descriptions. The following truth-theoretic axiom would help yield the desired blueprint:

$$(6) \quad [that_k: \phi] \psi \text{ is true of } s \text{ iff } \psi \text{ is true of at least one sequence that 'that=x \bullet \phi' is true of that differs from } s \text{ at most in the } k^{\text{th}} \text{ position.}$$

There is no reason why ϕ itself might not be *another* identity statement, which might suggest a treatment of ‘that NAME’ as equivalent to $[an_x: that=x \bullet NAME=x] \psi(x)$, assuming an appropriate axiom for NAME, consistent with a speaker-based notion of reference, of course. There is much to discuss here, but it should be clear that, given this notion of reference, under the Gödelian analysis the usual problems for accounts of demonstrative descriptions created by the twin pulls of rigidity and descriptive content of what is said evaporate.¹⁴⁸

22. Relativization Again

It was noted earlier that descriptions may contain pronouns bound by exterior quantifiers:

$$(1) \quad [\text{every man here}]^1 \text{ loves the woman he}_1 \text{ married}$$

¹⁴⁸ Notice that (i) the third person personal pronouns have possessive forms (‘his’, ‘hers’ and ‘its’), (ii) demonstrative descriptions have possessive forms (‘this man’s’ and ‘that man’s’), but (iii) the demonstrative pronouns ‘this’ and ‘that’ do *not* have possessive forms (*‘this’s’ and *‘that’s’): ‘*this’s colour’ and *‘that’s shape’ are usually regarded as ill-formed and must give way to the ‘French’ forms, ‘the colour of this’ and ‘the shape of that’, respectively. Is the following the beginning of an explanation? (a) there is no overt nominal for ‘s’ to attach to in the DP ‘that’ because its syntactic structure is really $[_{DP} \text{that } [_{NP} e]]$, where the first element is actually the *determiner* ‘that’ and the second, $[_{NP} e]$, an aphonic nominal; (b) ‘s’ cannot attach (linearly speaking) to the aphonic $[_{NP} e]$; (c) it cannot attach to the determiner ‘that’ either because $[_{NP} e]$ is (although not present at PF) in the way; (d) since ‘that’ is not assigned case by a verb, preposition or other marker (such as ‘s’), Chomsky’s (1981) case filter is violated. (*Mutatis mutandis*, for ‘these’ and ‘those’.) Pursuing this line of thought might, as Richard Larson has pointed out to me, help to explain why we find a similar situation with respect to other determiners that freely license a ‘missing’ NP. Sentence (i) is fine, but attempts to add the possessive fail, as in (ii), unless an overt nominal is present, as in (iii):

- (i) $[_{DP} \text{many } [_{NP} e]]$ applied, but $[_{DP} \text{few } [_{NP} e]]$ were selected
- (ii) * $[_{DP} \text{few } [_{NP} e]]$ ’s credentials were good enough
- (iii) $[_{DP} \text{few } [_{NP} \text{applicants}]]$ ’ credentials were good enough.

Unfortunately, this does not explain why possessive forms of some other demonstrative and indexicals are usually regarded as ill-formed. For example, although ‘yesterday’s’, ‘today’s’ and ‘tomorrow’s’ are fine as possessive forms, *‘then’s’ and *‘now’s’ are not (although there was once a band called ‘Now’s Children’). Nor are *‘here’s’ or *‘there’s’.

- (2) [every man who bought only one donkey]¹ paid cash for the donkey he₁ bought.

The subject DPs bind-into the descriptions by virtue of *wholly*-binding the pronoun 'he'. The descriptions are Russellian but *relativized* in the sense that uniqueness is relative to choice of man satisfying the NP in the subject DP.¹⁴⁹

It was also noted that even utterances of incomplete descriptions and pronouns may be understood as elliptical for bound-into descriptions:

- (3) [every man who bought only one donkey] paid cash for the donkey.
 (4) [every man who bought only one donkey] paid cash for it.

But here, of course, there is no actual variable-binding because there is no actual variable in the sentences' LFs, at least not according to the pragmatist proposal I suggested in *Descriptions* and have reiterated here. And that's why I prefer to talk of 'relativization', as I did in *Descriptions*, when talking about the interpretive phenomenon, rather than 'implicit binding', which has been known to lead people astray. All of (1)-(4) exemplify relativization; but on the suggestion at hand only (1) and (2) make that relativization transparent through the actual binding of variables in syntactic structure. It is the task of a pragmatic theory to explain how hearers retrieve the content speakers intend (or at least something close enough to it).

One of the standard uses of the English expression 'in question' when it is attached to a noun phrase seems to be explicitly to invite the hearer to see the description as incomplete; indeed, it be used to indicate relativization:¹⁵⁰

- (5) [every man who bought only one donkey] paid cash for the donkey in question.
 (6) every man kissed the woman on his immediate left. At least one man married the woman in question.¹⁵¹

The demonstrative descriptions 'that donkey' and 'that woman' in (5) and (6) may also be used here, but oddity would accompany 'that donkey in question' and 'that woman in question.' Why is this? And why is it that 'a woman in question' is odd but 'one of the women in question' and 'every woman in question' are fine (the intrusion has to be paragogic rather than epenthetic)?

23. The Argument from Binding

A number of philosophers and linguists have argued that some occurrences of definite descriptions function as bound variables and hence as referential expressions, which if

¹⁴⁹ There is nothing conceptually problematic about an expression that is bound-into wholly-binding (or binding-into) another: (i) [every man here]¹ thinks [the woman he₁ married]² loves [her₂ mother].

¹⁵⁰ This is noted by Kripke in his 1973 John Locke lectures, unfortunately still unpublished.

¹⁵¹ Related devices are 'under consideration' and 'at hand'. Unlike 'in question' these seem less happy in environments where relativized interpretations are called for.

true would create a problem for a unitary Russellian analysis.¹⁵² The arguments are far from compelling. First, in all of the cases I have seen discussed in the literature, the Russellian analysis yields the *right* result once incompleteness is taken into account and the distinction between wholly-binding and binding-into is appreciated; second, in certain cases the bound variable analysis yields the *wrong* results while the Russellian analysis continues to deliver the right results; and third, the Russellian might actually have the germ an explanation of why definite and demonstrative descriptions seem able to function in this way whilst indefinite descriptions do not.

Let me begin with the following claim by Wilson:

definite descriptions have a number of related *pronominal* uses, and referential use [of descriptions] is the use of these pronominal descriptions in direct, singular reference. Pronominal uses...contrast systematically and *as a matter of semantics* with instances of attributive use.' (1991: 359)

Wilson's case against the unitary Russellian position (and hence the motivation for his own theory of 'pronominal descriptions') turns on the unargued assumption that certain occurrences of descriptions *must be* treated as bound pronouns, an assumption that is simply false. Consider the following examples, on which the case rests:

- (1) [every scientist who was fired from the observatory at Sofia]¹ was consoled by [someone who knew [*the fired scientist*]₁ as a youth].
- (2) [every scientist who was fired from the observatory at Sofia]¹ was consoled by someone who knew [*him*]₁ as a youth.

It is Wilson's opinion that

the italicised description [in (1)], like the pronoun ['him' in (2)] that could replace it, is a variable bound by the quantifier phrase that fronts the sentence. The repeated lexical material in the descriptor helps here to indicate the tie between the description and the quantifier phrase that binds it (1991: 360-1).

On this account, the truth conditions of what is said by utterances of (1) and (2) are given by (3), the underlined variable $\underline{x}$ inside the second quantifier doing the work of the italicised pronoun in (2) and the italicised description in (1):

- (3) [*every* _{x} : scientist x • x was fired from the observatory at Sofia]
[*some* _{y} : y knew $\underline{x}$ as a youth] (x was consoled by y).¹⁵³

¹⁵² See (e.g.) Kempson (1986), Wilson (1991), Larson and Segal (1995).

¹⁵³ Remember these representations are *not* LFs. Stan Dubinsky has suggested that the descriptions Wilson sees functioning like bound pronouns fail to satisfy a condition that bound pronouns should satisfy. Compare (i) and (ii):

- (i) [every scientist who was fired from the observatory at Sofia]¹ was consoled by the project leader he_1 thought had made the decision to terminate him_1
- (ii) ? [every scientist who was fired from the observatory at Sofia]¹ was consoled by the project leader he_1 thought had made the decision to terminate [*the fired scientist*]₁.

(i) is a perfectly natural sentence in which 'he₁' and 'him₁' are interpreted straightforwardly as bound by the subject DP. By contrast (ii), which differs from (i) only in the replacement of the pronoun 'him₁' by the

There are several things to note here.¹⁵⁴

(i) For some speakers, there is an important difference between (1) and (2): what someone says by uttering the latter, but not by uttering the former, can entail that every scientist fired from the observatory at Sofia was male.¹⁵⁵ In order to forestall immediate objections to Wilson's proposal, let us think of 'him' in (2) as shorthand for 'him or her' (or simply replace it by 'him or her'); I will, however, return briefly to this matter when I compare Wilson's proposal with the Russellian's.

(ii) It is well known that demonstrative descriptions (phrases of the form 'that ϕ ') can be used to signify an anaphoric link in much the same way as definite descriptions. In (1), for example, 'the scientist' could just as well have been 'that scientist'.

(iii) A wooden Russellian treatment of the description 'the fired scientist' in an utterance of (1) would yield an interpretation we can render as (4), which obviously fails to capture the intended interpretation of the utterance:

- (4) [every_x: scientist x • x was fired from the observatory at Sofia]
[the_z: fired scientist z] [some_y: y knew z as a youth] (x was consoled by y).

The reason (4) fails is obvious: it fails to relativize values of z to values of x in the way Wilson's bound variable treatment (in effect) does by treating 'the fired scientist' as an occurrence of x .

(iv) On a more subtle Russellian treatment, 'the fired scientist' as it occurs in an utterance of (1) is an incomplete description that is meant to be interpreted as if it were an utterance of richer description that is bound-into. Saul Kripke seems to have spotted all of this some time ago.¹⁵⁶ A natural enrichment manifests itself in ordinary talk where we

description 'the fired scientist₁', does seem a little unnatural. The situation is similar with (iii) and (iv), which differ from (i) and (ii) only in the replacement of the nominative pronoun 'he₁' by the description 'the fired scientist₁', and which thereby differ from one another only in the replacement of 'him₁' in (iii) by the description 'the fired scientist₁' in (iv):

- (iii) [every scientist who was fired from the observatory at Sofia]₁ was consoled by
the project leader *the fired scientist₁* thought had made the decision to terminate *him₁*
(iv) ? [every scientist who was fired from the observatory at Sofia]₁ was consoled by
[the project leader *the fired scientist₁* thought had made the decision to terminate *the fired scientist₁*]

All of this suggests to Dubinsky that the occurrence of 'the fired scientist' in (ii) and the second occurrence of it in (iv) do not function as variables bound by 'every scientist who was fired from the observatory at Sofia' since they do not behave in exactly the same way as the occurrences of the pronoun 'him' in (i) and (iii), which *do* seem to function as bound variables. If Wilson maintains the standard line that variables are directly referential, then Dubinsky's examples, if convincing, present a genuine problem. I am grateful to Anne Bezuidenhout for forwarding Dubinsky's observations and questions, and for her own thoughts on the matter.

¹⁵⁴ An alternative diagnosis is offered by Simons (1996).

¹⁵⁵ There is some fluidity here. Fifty years ago, I imagine fewer people 'felt' the entailment (cp. 'someone has left his pen behind' which would have been regarded as default-neutral as to gender); and, of course, fifty years ago there were fewer female scientists in the world, numerically and proportionately.

¹⁵⁶ Examples similar to Wilson's are discussed by Kripke in his 1973 John Locke lectures, and the schematic form of what I take to be the right Russellian response can be extracted from his discussion. Kripke and Scott Soames brought up the same idea in the discussion following Wilson's presentation of similar examples at a conference on anaphora at Princeton in October 1990.

might find the overtly relativized description ‘the fired scientist in question’ In a representation of the truth conditions of an utterance of (1), a Gödelian description containing variables on *both* sides of the identity sign gives us exactly what we want:

- (1') [every_x : scientist $x \bullet x$ was fired from the observatory at Sofia]
 [the_z : fired scientist $z \bullet z=x$] [some_y : y knew z as a youth]
 (x was consoled by y).

The matrix of [the_z : fired scientist $z \bullet z=x$] is understood as uniquely satisfied relative to values of x , exactly as in the donkey and other relativized examples discussed earlier. In short, the Russellian says that the incomplete description in (1) is not, *pace* Wilson, a bound variable, but just another incomplete description—one for which the speaker could provide a fuller description that is bound-into—a description *containing* a bound pronoun. It is an incomplete, relativized description whose natural completion contains an expression understood as a variable bound by the subject expression.

(v) The Russellian analysis is perfectly capable of making transparent the relationship between anaphora on singular and quantified expressions, something Wilson rightly sees as repudiating a ‘pronoun of laziness’ approach to pronouns anaphoric on singular terms. For example, an utterance of (5) will be analysed as (5'), and an utterance of (6) as (6') (imagine (6) uttered in a context in which it is presupposed that all of the fired scientists were gifted astronomers):¹⁵⁷

- (5) Hugo Wexler was consoled by someone who knew *the gifted astronomer* as a youth.
 (5') [Hugo Wexler_x] [the_z : gifted astronomer $z \bullet z=x$] [some_y : y knew z as a youth]
 (x was consoled by y)
 (6) every scientist who was fired from the observatory at Sofia was consoled by someone who knew *the gifted astronomer* as a youth.
 (6') [every_x : scientist $x \bullet x$ was fired from the observatory at Sofia]
 [the_z : gifted astronomer $z \bullet z=x$] [some_y : y knew z as a youth]
 (x was consoled by y).

These straightforward analyses neatly capture the fact—which Wilson wants to capture—that the interpretation of utterances of (5) and (6) concerns the satisfaction of the same condition. According to the Russellian, the condition in question is given by (7), which is open in x :

- (11) [the_z : gifted astronomer $z \bullet z=x$] [some_y : y knew z as a youth]
 (x was consoled by y).

Someone uttering (9) says that Hugo Wexler satisfies it; someone uttering (10) says that every scientist who was fired from the observatory at Sofia does.

¹⁵⁷ For convenience, I have raised the name as if it were a quantifier. There are otherways of obtaining the desired result. See Neale (forthcoming *a*).

(vi) The Russellian account immediately explains the semantic difference between (1), which contains ‘the fired scientist’, and (10), which contains ‘the gifted astronomer’: the respective analyses (1') and (10') are not equivalent:

- (1') [every_x: scientist $x \bullet x$ was fired from the observatory at Sofia]
 [the_z: fired scientist $z \bullet z=x$] [some_y: y knew z as a youth]
 (x was consoled by y).
 (10') [every_x: scientist $x \bullet x$ was fired from the observatory at Sofia]
 [the_z: gifted astronomer $z \bullet z=x$] [some_y: y knew z as a youth]
 (x was consoled by y).

According to Wilson, however, (10) is ‘little more than a stylistic alternative’ of (1) (1991: 361). Evidently much turns on what is meant by ‘little more’. If ‘the gifted astronomer’ in (10) is meant to be a bound variable just like ‘the fired scientist’ in (1), then Wilson’s analyses of (1) and (10) are *equivalent* and ‘no more than a stylistic alternative’ of (1). But if Wilson wants his analyses to be *non-equivalent*, he needs to tell us what ‘little more’ there is to (10).

To sum up, the Russellian has a perfectly good account of why sentences can contain descriptions that *appear* to be functioning as bound variables—they are bound-into. Far from presenting problems for a unitary Russellian theory of descriptions, the examples discussed by Wilson serve only to emphasize the elegance and extraordinary range of Russell’s Theory of Descriptions.

A question remains, however. Why is it that definite and demonstrative descriptions can behave in this way whilst indefinite descriptions cannot? Larson and Segal (1995) see the contrast between the sentences in (12) and those in (13) as supporting the contention that definite and demonstrative descriptions may function *as* bound variables:

- (12) [[every boy]¹’s mother] loves him₁/[that boy]₁/[the boy]₁
 (13) [[every boy]¹’s mother] loves *[a boy]₁/*[a unique boy]₁/*[exactly one boy]₁.

In fact, the strongest conclusion that can be drawn is that definite and demonstrative descriptions may be incomplete bound-into Gödelian descriptions but indefinites may not. If the analysis of demonstrative descriptions above is on the right track the explanation is straightforward: the demonstrative description just *is* the requisite indefinite.

24. Are Pronouns Ever Wholly Bound?

Flushed with the success of the Russellian analysis of Wilson’s and related examples, we might get carried away. Why not interpret the pronoun ‘him’ in (1)

- (1) Every scientist who was fired from the observatory at Sofia was consoled by someone who knew *him* as a youth

as equivalent in this context to a relativized description $[the_z: \text{fired scientist } z \bullet z=x]$, or $[the_z: \text{scientist } \bullet z=x]$, or with attitude contexts in mind, $[the_z: z=x]$.

Why not treat *all* pronouns traditionally treated *as* bound variables as devices that only *contain* bound variables, as descriptions that are bound-into? On such an account *all* anaphoric pronouns would be D-types—no natural language bound pronoun would actually *be* a bound variable, but all would be interpreted as bound-into. How might such an idea be motivated and implemented? These are questions I take up in Neale (forthcoming a), but I can sketch part of the story here.¹⁵⁸ Recall that I am taking seriously Postal's idea that third person pronouns are versions of the definite determiner 'the'. An occurrence of a DP 'he', on the syntactic theory I favour, has the structure $[_{DP}[_{D}he] [_{NP}e]]$. Since $[_{NP}e]$ is aphonetic, its existence in a minimalist framework must be justified by its role at LF. So how is it interpreted? I suggest it is interpreted as a formula $x_k=x_j$ ($k \neq j$), meaning $[_{DP}[_{D}he] [_{NP}e]]$ is interpreted as $[he \ x_k: x_k=x_j]$, assuming an axiom for *he* that is a trivial modification of the Russellian axiom for *the* (as Postal's hypothesis would anyway suggest):

$[he \ x_k: \phi] \psi$ is true of a sequence s iff ψ is true of every sequence ϕ is true of differing from s at most in the k -th position, and there is exactly one such sequence.

The issue is ϕ , of course. The morphosyntax of English will, in fact, insist upon the complement of 'he' being $[_{NP}e]$. If the semantics of $[_{NP}e]$ is given by $x_k=x_j$ ($k \neq j$) (at least when it functions as the complement of the determiner 'he', $[_{NP}e]$ in $[_{DP}[_{D}he] [_{NP}e]]$ has a definite role at LF. When we have a suitably placed co-indexed binder as in (2), we get a bound interpretation of 'he', revealed by (2'), interpreted as (2''):

- (2) [every man]¹ thinks he₁ is smart'
- (2') $[_{DP} \text{every man}]^1$ thinks $[_{S} [_{DP} he^2 \ [_{NP} e]_1]^2$ is smart]]
- (2'') [every $x_1: \text{man } x_1]$ (x_1 thinks ($[he \ x_2: x_2=x_1]$ (x_2 is smart))).

If there is no such binder, we have a *free* occurrence of 'he', used to make indexical reference to some individual. If gender is seen as important, as I suggested above in connection with Wilson's examples, then we restate the axiom for the determiner *he* in such a way that ϕ is a conjunction ($x_k=x_j \bullet \text{male } x$).

Talk of *bound* pronouns is still perfectly intelligible on this proposal. The D 'he' is not bound; indeed it is a binder. And the DP 'he' is not *wholly-bound* the way a bound variable is; rather it is just *bound-into*. The subscript on 'he' in (2) must now be understood as indicating that $[_{DP}[_{D}he] [_{NP}e]]$ is *bound-into* rather than wholly bound. Trivially, we now have a 'uniform theory' of pronouns, something many semanticists crave. Some have laboured to produce exotic accounts of variable-binding that might

¹⁵⁸ Coming from a different direction, a proposal similar to the one I explore has been articulated nicely by Elbourne (2001), who draws upon the idea of binding situation variables inside pronominal DPs. It seems to me that our proposals have a common spirit and differ primarily in execution and assumed ontology, although I am not sure Elbourne would agree.

draw in pronouns occurring both inside and outside the scopes (as standardly defined) of their antecedents. My suggestion is that uniformity might be found without modifying the standard notion of binding. The proposal also seems to solve a problem raised in connection with Chomsky's (1981) Binding Theory. Principle A of that theory states (roughly) that a reflexive pronoun must be bound by something in its own clause, correctly predicting that 'every man' cannot bind 'himself' in (3):

- (3) * [_S[every man]¹ thinks [_SI like himself₁]]

So how does 'every man' manage to bind 'himself' in the isomorphic structure (4)?

- (4) * [_S[every man]¹ thinks [_She₁ likes himself₁]].

Answer: it doesn't: 'every man' binds the non-reflexive 'he' (satisfying Principle B, which states (roughly) that a non-reflexive must be bound by something *outside* its own clause), and 'he' in turn binds 'himself' (in accordance with Principle A). That is, the structure of (4) is really something like (4'):

- (4') [_S [_{DP}every man]¹ thinks [_S [_{DP}he² [_{NPe}]₁]² likes himself₂]].

Fourth, it might also help with the matter of gender: occurrences of 'he', 'him', and 'himself' anaphoric on (and within the scope of) on a quantifier understood as say, '[every_x: man x]' might be understood as '[he x_k: x_k=x_j • male x_k]. Whether an occurrence of 'he' always makes this richer contribution to what is said is, however, debatable. There is more than one way to skin a cat.