

Saying Nothing:
In Defence of Syntactic and Semantic
Underdetermination

Mark Bowker



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of PhD
at the
University of St Andrews

25/09/2015

1. Candidate's declarations:

I, Mark Bowker, hereby certify that this thesis, which is approximately 60,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in September, 2011 and as a candidate for the degree of PhD Philosophy in September 2011; the higher study for which this is a record was carried out in the University of St Andrews between 2011 and 2015.

Date signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date signature of supervisor

3. Permission for publication: *(to be signed by both candidate and supervisor)*

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

a) No embargo on print copy

ELECTRONIC COPY

a) No embargo on electronic copy

Date signature of candidate

signature of supervisor

Please note initial embargos can be requested for a maximum of five years. An embargo on a thesis submitted to the Faculty of Science or Medicine is rarely granted for more than two years in the first instance, without good justification. The Library will not lift an embargo before confirming with the student and supervisor that they do not intend to request a continuation. In the absence of an agreed response from both student and supervisor, the Head of School will be consulted. Please note that the total period of an embargo, including any continuation, is not expected to exceed ten years. Where part of a thesis is to be embargoed, please specify the part and the reason.

Contents

Introduction	5
--------------------	---

Chapter 1: The Encoding Model

The basic model	10
Extending the model	12
Implicature	12
Indexicality	13
Ambiguity	15
Non-sentential articulation	16
Implicature	17
The Encoding Model and the Epistemic Content Constraint	19
The domain of the Epistemic Content Constraint	20
Exploiting ambiguity	22
Contextual salience	23
Mistaken belief	24
Interpretive pretence	26

Chapter 2: Underdetermination arguments

Schiffer against the ‘naïve’ hidden-indexical theory of descriptions	31
Schiffer’s indeterminate reference	35
Wettstein against Russell’s analysis of definite descriptions	38
Blackburn against Wettstein’s referential account	42
Stanley and Gendler Szabó against a syntactic ellipsis account of quantifier domain restriction	46
Clapp against traditional semantics	47
Underdetermination without traditional semantics	52
Completion and Expansion	52
Austinian semantics	56
The S-View	57

Discourse Representation Theory	60
The source of the problem	62
Chapter 3: The paradoxes of underdetermination	
Epistemic underdetermination	64
Of what is said	64
Of what is uttered	66
Beyond language	68
Benacerraf: numbers are not sets.....	68
The Problem of the Many	70
Non-epistemic underdetermination	71
Of what is said	74
Of what is uttered	76
Chapter 4: Responding to the paradoxes.....	
Embracing underdetermination.....	78
What is said and what is communicated: one to many	80
Local equivalence	84
Responding to the underdetermination arguments.....	90
Schiffer against the 'naive' hidden-indexical theory of descriptions.....	90
Wettstein against Russell's analysis of definite descriptions.....	92
Blackburn against Wettstein's referential account.....	94
Stanley and Gendler Szabó against syntactic ellipsis	95
Clapp against traditional semantics	96
Underdetermination without traditional semantics.....	97
Completion and expansion.....	97
Austinian Semantics	98
Discourse Representation Theory.....	99
How to build an underdetermination argument	99

Chapter 5: An account of the referential/attributive distinction.....	102
The referential/attributive distinction	102
The problem for Russell	106
A Russellian account of the distinction	108
Yablo on non-catastrophic presupposition failure.....	115
Loar on the semantics of singular terms.....	119
Developing the distinction	122
 Chapter 6: Objections and Replies	 128
Can we retain traditional semantics?.....	128
Revenge problem 1: The underdetermination of what is meant	129
Response	131
Revenge problem 2: The underdetermination of what is presupposed.....	136
Response	137
Pluralism about what is said	141
Speech act pluralism	142
The benefit of underdetermination	148
 Conclusion:	 150
 Bibliography	 151

Abstract:

According to the Encoding Model, speakers communicate by encoding the propositions they want to communicate into sentences, in accordance with the conventions of a language L. By uttering a sentence that encodes p, the speaker says that p. Communication is successful only if the audience identifies the proposition that the speaker intends to communicate, which is achieved by decoding the uttered sentence in accordance with the conventions of L.

A consequence of the Encoding Model has been the proliferation of underdetermination arguments, each of which concludes against some linguistic theory T, on the grounds that, were T true, audiences would be unable to know what was said by utterances of some particular linguistic form, and therefore unable to know what speakers intended to communicate by these utterance. The result, if we accept the conclusion of these arguments, is radical restriction of the domain of viable linguistic theory.

This Thesis defends an alternative model according to which there need be nothing encoded in an uttered sentence – nothing that is said by its utterance – for the audience to retrieve. Rather, there are indefinitely many ways to interpret uttered sentences – indefinitely many routes to the propositions that speaker intend to communicate – which proceed through different interpretations of what is said.

Introduction¹

According to the Encoding Model, communication proceeds as follows. The speaker has a proposition that they want to communicate to their audience. Coordination on this proposition is achieved through the conventions of a language. The speaker utters a sentence which, in a language shared by the speaker and their audience, is associated with the proposition that the speaker intends to communicate. By uttering such a sentence, the speaker *says* or *expresses* the proposition that they intend to communicate, which is the *semantic content* of the sentence uttered, or the proposition it *semantically expresses*. Audiences decode uttered sentences in accordance with the relevant conventions, thereby deriving whatever the speaker has said.

As a consequence of this model, several authors have endorsed the claim that, in a broad range of cases roughly characterised as cases of *literal speech*, an audience cannot know what a speaker intends to communicate if they cannot know what the speaker says. This is the *Epistemic Content Constraint*. If the speaker can know what the audience intends to communicate, and that the speaker says what they intend to communicate, then the audience can know what the speaker has said: it is whatever proposition they intend to communicate.

This assumption has led in turn to a proliferation of *underdetermination arguments*, each of which confronts some theory T with the *problem of underdetermination*. The problem is that, if T were true, then audiences would be unable to know what is said in a number of simple cases. By the Epistemic Content Constraint, therefore, these audiences would be unable to know what their speakers intend to communicate. As audiences clearly can know what speakers intend to communicate in these cases, T cannot be true.

Rather than taking these arguments to motivate the rejection of their respective theories T, this Thesis responds to the problem of underdetermination by offering an alternative model of communication, according to which it is no requirement on successful or normal communication that anything be said. Rather, utterances can be associated with indefinitely many candidates for what is said, each of which would communicate the same

¹ Presented with thanks to the generous financial support of the AHRC and the equally generous emotional support of my parents. Thanks to my supervisors Patrick Greenough, Andy Egan, Herman Cappelen, and Josh Dever, and to my examiners Derek Ball and Stefano Predelli. Thanks also to various members of the Arche Research Centre, past and present, for the many discussions that shaped this work.

information. Audiences needn't know what was said to know what the speaker intended to communicate. Knowing that the speaker intended to communicate p , the audience needn't be able to know whether the speaker said that p or that q , if saying either would communicate that p .

Chapter 1 presents the Encoding Model of communication. The basic formulation of the Encoding Model leaves it suited to explaining only a limited number of potential communicative situations. To broaden the scope of the model, it is elaborated to allow for cases in which the speaker means something other than, or in addition to, what they say, as well as various cases in which the audience must rely on the context in order to settle what is said by the articulation of a particular sentence, including cases of ambiguity, indexicality, non-sentential assertion, and implicature.

Despite these elaborations, the basic case envisioned by the Encoding Model remains much the same. It is the case in which the speaker utters a sentence which is associated with the proposition that the speaker intends to communicate, although decoding the sentence relies on far more than deductive inference, including the audience's knowledge about the world and their general inferential capacities. This has led to the underdetermination arguments that we encounter in Chapter 2. Chapter 2 presents only a sample of extant arguments, some of which disagree as to the source of the problem.

The first instance is Schiffer's (1995) argument against the 'naïve' hidden-indexical theory of descriptions. The second example of an underdetermination argument comes from Wettstein (1981) who argues against an defence of Russell's (1905) theory of definite descriptions against *prima facie* counterexamples involving 'incomplete' or 'indefinite' definite descriptions such as 'the table', according to which a sentence including an incomplete definite description is interpreted as a longer sentence that includes further domain-restricting descriptive material. The problem of underdetermination leads Wettstein to reject this defence of Russell's theory and propose instead that definite descriptions be completed, not with further descriptive linguistic material, but with demonstrative, or directly-referential, material. Blackburn (1988) shows that the problem of underdetermination can also be formulated in opposition to Wettstein's account.

Stanley and Gendler Szabó (2000) identify the problem posed for Russell's theory of definite descriptions by incomplete definite descriptions as instances of a more general

problem of quantifier domain restriction and extend Wettstein's argument to reject all syntactic ellipsis explanations of domain restriction. Stanley and Gendler Szabó defend instead the hidden-indexical account of quantifier domain restriction. Clapp (2002) explains that the problem of underdetermination is as serious a problem for a hidden-indexical account as it is for a syntactic ellipsis account.

Clapp suggests that the source of the problem of underdetermination is the assumption of Truth Conditional Compositionality, that is, the assumption that what is said by the utterance of a sentence depends only on the structure of the sentence and the meanings of the terms it contains. As we shall see towards the end of Chapter 2, however, underdetermination is as problematic for the authors that Clapp cites as rejecting Truth Conditional Compositionality as it was for the authors previously considered.

Chapter 3 distinguishes between epistemic and non-epistemic versions of the problem of underdetermination. The epistemic version of the problem is fuelled by the Epistemic Content Constraint, according to which knowledge of what the speaker intended to communicate requires knowledge of what they said. If what is said is underdetermined by the evidence available to the audience, then the audience cannot know what was said and, by the Epistemic Content Constraint, cannot know what was meant. The non-epistemic version of the problem takes matters a step further, arguing not only that the evidence available to the audience is insufficient for them to know what was said, but that all the evidence there is cannot determine what is said, that is, that nothing is said. Chapter 3 also explains how Unger's (1980) Problem of the Many feeds the problem of underdetermination by showing how it can be extended to singular terms.

Chapter 4 introduces the response to the problem of underdetermination, which is to show that underdetermination of what is said is not inherently problematic. The Epistemic Content Constraint is rejected. Audiences can know what a speaker intended to communicate, even though they cannot know what was said, when all of the candidates for what is said consistent with the audience's knowledge would communicate the same information. Central to this conclusion is the observation of *local equivalence*. Different propositions, defined by different truth-conditional profiles, can take the same truth value at any possibility under consideration. When two propositions are locally equivalent, the saying of either proposition will communicate all and only the same propositions.

We will see that, in each of the arguments from Chapter 2, the candidates for what is said are locally equivalent, with the result that the audience can know what the speaker intended to communicate without knowing which of the candidates was said. Indeed, the audience can know what the speaker intended even if, as per the non-epistemic argument from Chapter 3, nothing is said. By uttering the right sentence in the right context, a speaker can lead the audience to recognise what they intend to communicate, even though they leave open various paths to arrive at that conclusion; by, for example, interpreting the utterance as though one proposition was said or as though another proposition was said. The local equivalence of the candidates for what is said is essential if the arguments from Chapter 2 are to secure another of their key premises, that is, the premise that the audience can know what the speaker intended to communicate. Chapter 4, by providing instructions for the formulation of underdetermination arguments, will show how underdetermination arises, unproblematically, for any utterance.

Chapter 5 deploys the view motivated in Chapter 4 to offer a unified Russellian account of Donnellan's (1966) distinction between referential and attributive uses of definite descriptions. At the heart of the distinction is Donnellan's observation that some uses of definite descriptions communicate propositions other than those Russell takes them to say. As we shall see, Donnellan's observation is predicted by a Russellian account that recognises the local equivalence of propositions. Following this result, Chapter 5 presents an account of non-catastrophic presupposition failure. The account follows a similar line to that proposed by Yablo (2006) but, unlike Yablo's account, allows for a univocal account of definite descriptions, whether used referentially or attributively. Chapter 5 closes by considering Loar's (1976) account of the referential/attributional distinction and distinguishing it from the account offered in this Chapter.

Chapter 6 replies to various possible objections. First is the objection that, if nothing is said, then we cannot retain any of the views discussed in Chapter 2. On the contrary, such theories can be true, even if nothing is said. According to the alternative model of communication defended in this Thesis, *prima facie* inconsistent linguistic theories can be viewed as consistent. The view is not only consistent with any one of the linguistic views discussed in Chapter 2, but with their conjunction.

A response is then provided to a potential revenge problem suggested by Buchanan's (2010) argument that speakers cannot always mean propositions by their utterances, because there is often no proposition that satisfies the principle that if a speaker means a proposition *p* by her utterance *U* then her audience must entertain *p* if she is to understand *U*. In response, it is argued that Buchanan is assuming a mistaken view of what it is to *entertain* a proposition, according to which the entertainment of two propositions is more laborious than the entertainment of a single proposition.

A further revenge problem is that the response offered in Chapter 4 only delays the problem of underdetermination. Although the underdetermination of what is said is unproblematic if we assume a common background of presuppositions, the underdetermination of what the speaker presupposes remains problematic. If the audience cannot know what the speaker presupposes, then the audience cannot know what the speaker intends to communicate. This problem is significantly different from the original problem of underdetermination, which threatened failures of communication when speaker and their audience share a common stock of presuppositions. This new problem, however, threatens failures of communication when speaker and their audience fail to share a common stock of presuppositions.

If the audience cannot know what the speaker presupposes, then the audience cannot know what the speaker intends to communicate, but we have good reason to suppose that speakers and their audiences can coordinate their presuppositions. If they cannot, then communication will at best be imperfect, according to the account offered in this Thesis. On any account, however, we should expect a failure to coordinate our presuppositions to manifest in a failure (or, at least, imperfection) of communication.

The view of this paper is then distinguished from, and compared with, other views, including pluralism about what is said and Cappelen and Lepore's (2005) Speech Act Pluralism. We close by noting that the underdetermination of what is said is not a product of speaker error, but carries with it certain communicative benefits.

Chapter 1: The Encoding Model

The basic model

According to the *Encoding Model*, communication works as follows. The speaker has a particular thought that they want to communicate to their audience, which they encode into words in accordance with a language which they take their audience to share.² The articulation of those words constitutes the saying of the proposition the speaker means. The audience then recognises what the speaker intends to communicate by recognising what they say, according to the conventions of their shared language, and recognising that the speaker intends to communicate what they say. Communication succeeds only if the audience recognises what the speaker *means* or *intends to communicate*.³

Let's begin with a simple illustration. Suppose I want to communicate to you the thought that John runs. To do this, I articulate a sentence, which, in the conventions of our shared language is used to say that John runs; the sentence 'John runs' for example.⁴ Thoughts have proposition contents. At a minimum, therefore, knowing what thought I intend to communicate involves identifying some proposition that is the content of the thought I intend, which, again, minimally requires identifying some truth conditions. According to a plausible account of 'John runs', this sentence is associated, by the conventions of English, with the proposition that John runs, which is why I *say* that John runs in articulating this sentence. Equivalently, the *semantic content* of the sentence articulated is the proposition that John runs, which is why I *semantically express* the proposition that John runs, in articulating this sentence.

The meaning of the sentence is not normally taken to be a matter of brute convention, but as derivable from the meanings of the words in the sentence. I utter a sentence composed

² Aside from the communication of thoughts, speakers may intend to communicate feelings, attitudes, biases, and so on. Throughout this Thesis, we will be focusing on the communication of thoughts. We will also tend to focus on cases of vocal communication, but everything said should extend easily to other modes of linguistic communication.

³ We might think that this is a necessary and sufficient condition on successful communication, or we might take it to fall short of sufficiency, requiring, perhaps, that the audience identify the speaker's thought in the right way. This condition is equivalent, in my terminology, to the condition that the audience *knows*, or *identifies* what the speaker intends to communicate.

⁴ This fact needn't exhaust the conventions associated with the sentence, or even those conventions that determine the meaning of the sentence. Predelli (2013) refers to additional aspects of meaning as 'biases'. We will be focusing on un-biased meaning.

of two distinct terms, the noun ‘John’ and the verb ‘runs’. According to one account, the meanings of these terms are captured by their referents. In English, ‘John’ refers to John, while ‘runs’ refers to the property of running, represented as a function from an object x to a proposition (specifically, the proposition that x runs). This representation is intended to capture the interaction between the meanings of these terms, with the meaning of ‘John’ becoming argument to the meaning of the term ‘runs’ to determine the meaning of the sentence: the proposition that John runs, which is true if and only if John runs. This proposition can itself be represented mathematically as a function from worlds to truth values (specifically, the function from a world W to truth if and only if John runs in W and from W to falsity otherwise).⁵ If you understand English (or at least this fragment of English) then you can know what sentence I have uttered, the meaning of the sentence, and so the proposition I intend to communicate.

The Encoding Model is currently only suited to discussion of a very limited range of situations. For the model to find more general application, it can be developed in a number of ways. In the next section, the picture will be extended to accommodate implicature, impliciture, ambiguity, non-sentential articulation, and context-sensitivity. Note that these extensions to the basic model will take us beyond Sperber and Wilson’s (1986, 9) *code model*, as we allow that “there is a gap between the semantic representations of sentences and the thoughts actually communicated by utterances” which is filled “by inference.”

⁵ This is not the only formal representation of the proposition in play in contemporary philosophy of language. Many prefer representation of propositions in terms of the objects and properties that are conceived of as their constituents. So the proposition that John runs might be represented by the ordered pair $\langle \text{John, the property of running} \rangle$. Although such a *structured proposition* has a truth condition – it is true if and only if John runs – its properties are not exhausted by its truth condition. For the most part, we will stick with the representation of propositions as sets of worlds or functions from worlds to truth values. Possible worlds are possible states of the actual world – they are ways the actual world might have been – where the actual world is that maximal object in which we find ourselves. One way of incorporating the insights of Predelli (2005, Chapter 4), however, is by noting that one and the same world can be carved up in various ways, depending on our interests. There is a sense, therefore, in which any world (even the actual one) is constructed, rather than given. In deciding whether or not ‘John is running’ is true, it is not enough to know what is said by the sentence – that John runs – and to know what the world is like – that it includes John moving in such-and-such a way – I must also make some decision as to whether John’s movements count as running for the purposes at hand. Standards may vary significantly between a 100m sprint and a walking race and so the truth-value of ‘John is running’, even when said of the same individual at the same time. We’ll return to the topic briefly in the following Chapter.

Extending the model

Implicature

Contrary to the Encoding Model, speakers do not always intend to communicate what they say. Speakers may intend to communicate something in addition to what their words mean, or may not even intend to communicate what their words mean in the first place. As an example of the first sort, if you ask whether John, Abby, or Paula will deliver a message fastest, I may intend to communicate, by my response that 'John runs', that John will deliver the message fastest. I intend to communicate that John will deliver the message fastest but this is not what I say by my articulation. The meanings of the words and the sentence composed of them do not change between this and the simpler case above, so what I say is the same. Of course, another thing that I may intend to communicate in this case is precisely what I say: that John runs.

There are also cases in which speakers do not even intend to communicate what they say. Consider, for example, metaphor and figurative language. Figurative speakers can use the very same sentences, associated with the very same meanings, as literal speakers, but they don't intend to communicate the things they say by the sentences they articulate. The meaning of 'John is a pig' may be the same whenever it is uttered, but only the literal speaker intends to communicate that John is a pig. Consider also cases of hyperbole. The content of 'Smith is the dumbest person alive' is very plausibly false if there is a person dumber than Smith, yet what the speaker intended to communicate might be true in the very same situation because the speaker intended to communicate only the weaker proposition that Smith is very dumb.

Propositions communicated without being said are what Grice called 'implicatures'. Grice (1991, 26-27) posits a number of principles that govern cooperative enterprises, which include, among other things, communicative exchanges. Violation of any of these principles can be used to introduce an implicature. Central is the Cooperative Principle which commands: Make your conversational contribution such as is required at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. In general, obeying the Cooperative Principle requires that we obey certain derivative Maxims. The Maxims of Quantity demand that speakers offer as much information as is required for the purposes of the exchange, but no more. If I ask you where the nearest butcher's is, then in

general you should tell me where the nearest butcher's is but not also where the nearest grocer's is. The Maxims of Quality command that speakers should not say anything they believe to be false or for which they lack adequate evidence. The Maxim of Relation commands that speakers 'Be Relevant'. Grice doesn't expand much further on this Maxim, but the project is taken up extensively in Sperber and Wilson (1986). Finally, the Maxims of Manner command that speakers avoid obscurity, ambiguity, unnecessary prolixity, and disorder. Speakers can indicate to their audiences that they mean something other than, or in addition to, what they say by saying something that threatens to violate one of the Maxims. When I tell you that 'John is a pig', the recognition of what I have said and its patent falsity might be taken as an indication that I mean some proposition other than that encoded by my words.

Indexicality

According to the Encoding Model, the conventions of a language associate sentences with the propositions they are used to say. Natural languages contain an abundance of sentences that are *indexical* or *context-sensitive* in that they can be used to say different things in different contexts. What is said by an utterance of 'She is eating cabbage'? What thought or proposition is it associated with? In abstraction from some occasion of its use, the conventions of English don't provide an answer. One explanation is that 'She' has no reference in abstraction from a context of its use. Given a context in which 'She' refers to Paula, the sentence says that Paula is eating cabbage. Given a context in which 'She' refers to Margo, the sentence says that Margo is eating cabbage. To determine what is said by this sentence, we need to refer to a particular context, as well as the conventions of English.

As Frege (1979, 213) puts it "The sentences of our everyday language leave a good deal to guesswork. It is the surrounding circumstances that enable us to make the right guess. The sentence I utter does not always contain everything that is necessary; a great deal has to be supplied by the context". The conventions of English do not determine the referent of 'She' and so fail to determine what is said by an utterance of 'She is eating cabbage'. To know what is said and therefore what the speaker intends to communicate (on the assumption that they intend to communicate what they say), the audience has to know, not only the sentence articulated and the conventions of English, but also the context in which the sentence was

articulated. While perhaps not as a matter of infallible logical deduction, identifying the reference of ‘She’ in the context is not entirely a matter of complete guesswork. While the conventions of English fail to determine the referent of ‘She’, or what is said by the sentence uttered, they plausibly provide restrictions on viable referents and so some guidance as to what is said. Very plausibly, the conventions surrounding ‘she’ are such that this term can only refer to females, so ‘She is eating cabbage’ cannot be used to say anything that requires ‘She’ to be assigned a non-female referent.⁶ In identifying what is said, however, the audience must appeal to the context, as well as the conventions associated with ‘She’ in English.⁷

Nothing is said by an articulation of ‘She is eating cabbage’ in abstraction from a context but that is not to say that this sentence is completely devoid of meaning in abstraction from context. We can distinguish two levels of sentential meaning. The first, what is said, is, for indexical sentences, defined only relative to context, but while nothing is said by ‘She is eating cabbage’ in abstraction from context, this sentence still has a meaning in virtue of which we can identify restrictions on the propositions that it can be used to say. This is the level of meaning that Kaplan (1989) calls ‘character’ and which I may sometimes call ‘conventional’ or ‘context-invariant’ meaning.

Features of context relevant to identification of the referent of a indexical term are standardly taken to include, for example, the speaker (to determine the reference of ‘I’) their audience (to determine the reference of ‘you’), the time and location at which the articulation is made or interpreted (to determine the references of ‘now’ and ‘here’), and a number of salient objects determined by the history of the conversation and indications (such as gestures) from the speaker (to determine the references of terms like ‘this’ and ‘that’).⁸ Less standardly, Predelli (1998 a, b) takes the referents of indexical terms to be determined by the speakers intentions; Corazza, Fish, and Gorvett (2002, 11) take their referents to be determined by “the social or conventional setting in which the utterance takes place”; Romdehn-Romluc (2003,

⁶ I say “very plausibly” rather than “obviously” at least partly because uses of ‘she’ to refer to prized possessions such as cars and boats may be entrenched enough to be considered literal.

⁷ Stanley and Gendler Szabó (2000, 229) represent this interpretive role of context by the equation “what is uttered + linguistic meaning + context = what is said”.

⁸ Perhaps the clearest examples of context-sensitive expressions are the so-called *indexicals* like ‘I’, ‘you’, ‘now’, ‘here’, ‘this’, and ‘that’. Other candidates for context-sensitive terms include terms that require an orientation such as ‘to the left’, that require a comparison class such as ‘tall’, or that require some standard of precision such as ‘red’ and ‘knows’.

38-39) argues that “the reference of an indexical is determined by the context that [a competent and attentive version of the audience who it is reasonable to take the speaker as addressing] would identify on the basis of the cues that she would reasonably take [the speaker] to be exploiting”. We will return to some of these considerations later.

Ambiguity

According to the Encoding Model, the conventions of a language associate sentences with the propositions they can be used to say. We should be clear as to what we mean when we talk about ‘sentences’, however. In common parlance we can talk about ambiguous sentences. If a speaker utters a sentence that is ambiguous in language L, then the conventions of L fail to determine what is said by the utterance of that sentence. The sounds I make in reading ‘John lies’ are ambiguous because there are two different interpretations of ‘lies’, which, on one reading, contrasts with telling the truth and, on another reading, contrasts with sitting or standing. There is no proposition that is associated, by the conventions of English, with the ambiguous sentence ‘John lies’. Rather, the conventions of English associate each disambiguation with a different proposition. In deciding between these interpretations, the audience must appeal to features of the conversational situation or context other than the sounds the speaker has made and the language they are speaking, as when resolving context-sensitivity or implicature.⁹

In cases of lexical ambiguity there are a number of different candidates for the words that make up an utterance, as in the case of ‘John lies’. Even when the audience is clear on how they should interpret each of the individual words, however, it may be unclear what structure they should impose on them. The sounds you would make in reading ‘Every member of staff shares an office’ are often thought to be structurally ambiguous in that ‘every’ can take wide or narrow scope. Only on the narrow-scope reading will the sentence entail that there is one office that is shared by every member of staff; the wide-scope reading is consistent with situations in which each member of staff shares an office with someone, even though there is no single office that is shared by all of them.

⁹ Here we set aside the vexed question of whether proper names like ‘John’ are ambiguous, context-sensitive, or what.

We can distinguish between the sounds made or the marks on the page, which may be ambiguous both lexically and structurally, and the unambiguous object that is associated with a proposition by the conventions of the language (with whatever contextual contribution is necessary to resolve context-sensitivity). The first, I will call *what is articulated*; the second, *what is uttered*.¹⁰

Non-sentential articulation

According to our initial characterisation, speakers utter sentences, yet it is often appropriate to articulate non-sentential expressions in isolation, as when I respond to the question ‘What is John doing?’ by replying, simply, ‘Running’. This is consistent with the view that sentences are the objects of utterance, however. Plausibly, articulation of this single word is sufficient, in the context, for the speaker to make manifest the sentence that they intend to utter;¹¹ in this case ‘John is running’. Such cases, in which terms that can be inferred from the context are omitted from what is articulated, we call cases of *syntactic ellipsis*.¹² As the audience might have to appeal to context to discern the sentence uttered when the speaker articulates a grammatically complete but ambiguous sentence, so the audience must appeal to context to discern the sentence uttered when the speaker articulates a grammatically incomplete sentence.

The view that sentences are the objects of the uttering-relation, while of course not universally accepted,¹³ can be motivated from some widely held assumptions about the role of context in determining what is said on the basis of what is uttered. If the only role that context plays in the determination of what is said, given a particular sentence, is to assign referents to indexical expression¹⁴ then subsentential articulations are either elliptical for complete sentences, or language is radically indexical. Consider a subsentential articulation like ‘Running’. This term can be used in one context to say that John is running (as when I am responding to the question ‘What is John doing?’) and in another context to say that running is my favourite activity (as when I am responding to the question ‘What is your favourite activity?’). If what is articulated is not elliptical for some complete sentence uttered, then the

¹⁰ This terminology follows Stanley and Gendler Szabó (2000).

¹¹ Or, if you prefer, to make manifest the sentence that they intend the audience to interpret.

¹² The term is borrowed from Stanley and Gendler Szabó (2000).

¹³ See Stainton (1995) for an alternative.

¹⁴ See the following section for reasons to deny this claim.

single term ‘Running’ must refer to the proposition that John is running, in one context, and the proposition that running is my favourite activity, in another context. Plausibly, any subsentential expression can be used to say any proposition, given the right context, so the context-sensitivity required would be truly radical: any term can refer to any proposition.

Implicature

The Encoding Model assumed that the sentence uttered by a speaker will be associated with a propositional content by the conventions of the speaker’s language. Bach (1994) argues, however, that the meaning of complete a sentence often falls short of any complete proposition and interpreters have to engage in a process of *completion* to derive a truth conditional content from the meaning of the sentence. According to Bach (1994, 127) for example, the meaning of ‘Steel is not strong enough’ fails to determine a complete proposition. Rather, the sentence determines a propositional radical that requires completion in the form of “some contextually identifiable respect” in which steel is not strong enough (such as *for building a 500-storey building* or *to resist bending by Superman*) to determine a truth-evaluable proposition.

Even when the meaning of the sentence uttered does determine a complete proposition, speakers often intend to communicate, not that proposition, but a related one. In these cases, interpreters may have to engage in a process of *expansion* to derive the proposition that the speaker intended to communicate from what they have said. Someone who utters “John hasn’t eaten breakfast” may be naturally understood, in the right context, as meaning that John has not eaten breakfast that day. The meaning of the sentence itself, however, plausibly determines the stronger proposition that John has not eaten breakfast, tout court.

Bach calls propositions communicated by expansion or completion *conversational implicatures*,¹⁵ which are, according to Bach (1994, 126), distinct from implicatures because “In implicature one says and communicates one thing and thereby communicates something else in addition. Implicature, however, is a matter of saying something but communicating

¹⁵ Note that implicatures needn’t always be communicated. When I utter ‘I haven’t eaten breakfast’, the implicature may be that I haven’t eaten breakfast today, even though I implicate and communicate something else; perhaps that I am hungry.

something else instead, something closely related to what is said.”¹⁶ Bach (1994, 143-144) notes that his characterisation of implicature isn’t clearly Grice’s, given that Grice analyses metaphors, for example, as instances of implicature but someone who calls you ‘a pig’ doesn’t communicate that you are literally a pig. A further problem with this characterisation is that it fails to fit some classic examples of expansion. Someone who utters ‘Mary went to the edge of the cliff and jumped’ might naturally be understood as communicating that Mary jumped over the edge of the cliff. The meaning of the sentence itself, however, very plausibly determines the less informative proposition that Mary went to the edge of the cliff and, simply, jumped, which is consistent with her jumping over the edge, jumping backwards, jumping up and down on the spot, and so on. While this seems like a clear case of expansion, and so implicature, it meets Bach’s condition for implicature: the speaker means what they say – that Mary jumped *tout court* – but means something else in addition – that Mary’s jumping was over the edge. Very plausibly, one cannot mean that Mary jumped over the edge without also meaning that Mary jumped *tout court*, the latter being entailed by the former. Witness the apparent absurdity of saying ‘I mean that Mary jumped over the edge of the cliff. I don’t mean, however, that Mary jumped’. If you mean that Mary jumped over the edge, then part of what you mean is that Mary jumped, the other part being that her jumping was over the edge of the cliff.

An alternative way to distinguish between implicatures (in Grice’s sense, which includes metaphor) and implicatures is by noting that any expansion or completion must “be closely related to the conventional meaning of the words (the sentence) ... uttered.”¹⁷ Expansions and completions are alike in that they are derived by supplementing the meaning of the sentence uttered. Additional conceptual material can be added but the implicature must retain this kernel of meaning, which was derived from the structure and components of the sentence uttered. Implicature, however, needn’t retain any of the meaning of the sentence uttered. By uttering ‘Mary went to the edge of the cliff and jumped’, in response to the question ‘Has anyone been outside today’ I can, for example, implicate that someone has been outside today. What I say and implicate are, in this case, entirely distinct.

¹⁶ Here, Bach takes ‘what is said’ to identify the contextually-determined meaning of the sentence uttered, even when it falls short of a complete proposition. An alternative is to reserve this term for the proposition determined by the process of completion or expansion, but I take this to be nothing more than a matter of terminology.

¹⁷ This is Grice’s (1989, 25) necessary condition on what is said.

The Encoding Model and the Epistemic Content Constraint

We have seen how the Encoding Model can be developed to deal with implicature, ambiguity, indexicality, nonsentential assertion, and implicature. Despite these developments, however, the basic case remains that of literal speech, in which the speaker intends to say and communicate some particular proposition, which the audience can determine from their knowledge of the sentence uttered, the conventions of the language spoken, and the context.

In these cases, the Encoding Model suggests the Epistemic Content Constraint:

Epistemic Content Constraint: If the audience cannot know what is said, then the audience cannot know what the speaker intends to communicate.

In contrapositive form: if the audience can know what the speaker intends to communicate, the audience can know what was said. The speaker intends to communicate what they say. If the audience can know what the speaker intends to communicate, therefore, the audience can know what the speaker has said: it is whatever they intend to communicate.

Similarly, we might think, the speaker who knows what is said knows what is uttered, given that what is said is determined in part by the sentence uttered. This motivates the Epistemic Utterance Constraint:

The Epistemic Utterance Constraint: If the audience cannot know what is uttered, then the audience cannot know what the speaker intends to communicate.

In contrapositive form, if the audience can know what the speaker intended to communicate, the audience can know what the speaker uttered. If the audience can know what the speaker intended to communicate, then, by the Epistemic Content Constraint, they can know what the speaker said. If the audience can know what the speaker said, then they can know what the speaker uttered: it is whatever sentence the speaker could have uttered to say what they said. This reasoning is fallacious, given that the speaker could have uttered a number of different sentences to say what they did. Assuming that the speaker said that John runs, for example, the speaker might have uttered 'John runs' or the speaker might have uttered 'He runs' in a context such that 'He' refers to John. The above reasoning goes through, however, when the only candidates for the sentence uttered that are consistent with the evidence available to the

audience are non-synonymous. If the audience knows what is said and that the only candidates for the sentence uttered are S and S^* , each of which expresses a different proposition, then they audience can know which of S and S^* was uttered: it is whichever of these sentences expresses, in the context, what the speaker said.

We see both the Epistemic Utterance Constraint and the Epistemic Content Constraint endorsed by Stanley and Gendler Szabó (2000, 231) who describe as “immensely plausible” the assumption that “the hearer who grasps the proposition communicated will also know what sentence was uttered and what proposition was [said] by that sentence on the given occasion.”¹⁸

The domain of the Epistemic Content Constraint

According to the Encoding Model, in cases of literal speech, the speaker intends to communicate what they say. This needs some further caveats to accommodate a variety of further cases, however. The most significant of these are cases in which the speaker intends the audience to interpret a context-sensitive sentence as though it were uttered in a context other than the actual context. The speaker therefore intends the audience to interpret as though the speaker said something other than what was actually said, thereby intending to communicate, not what they actually said, but what they would have said in another context. The phenomenon is independently motivated and provides a response that undermines various arguments against traditional account of indexicals.

Stanley and Gendler Szabó (2000, 231) endorse the Epistemic Utterance Constraint and the Epistemic Content Constraints, more or less, in describing as “immensely plausible” the assumption that “in normal instances of *successful* communication, the hearer who grasps the proposition communicated will also know what sentence was uttered and what proposition was expressed by that sentence on the given occasion.” Stanley and Gendler Szabó explicitly restrict the claim to “normal instances” of successful communication but do not quite tell us what counts as a normal instance. The pair do tell us earlier that their entire discussion is

¹⁸ The original uses ‘expressed’ rather than ‘said’ but the two are equivalent in the language of that paper. Interestingly, the pair seem to take the Epistemic Utterance Constraint to hold even when the only candidates for what is uttered are two non-synonymous sentences. This interpretation of the Constraint will not be assumed in what follows.

restricted to *typical assertions*, where “(i) there is a single speaker and a single hearer, (ii) the speaker vocalizes a well-formed, meaningful sentence, and by doing so (iii) the speaker intends to convey a certain proposition.” (226) Even with these restrictions, however, it is easy to formulate *prima facie* counterexamples to the Epistemic Constraints. Even if you are unable to know what I have said, there are various ways to know what I meant. You might be able to come to know what I meant by testimony from a third party who does know what I said, or by induction from my past utterances, for example.

Clearly, however, coming to know what I meant by testimony or induction is only possible given a context in which there is a third-party who knows what I said, or one in which you have knowledge of my historical utterances. In the central cases that will occupy us throughout this Thesis, it is clear that the audience can know what the speaker means to communicate, given the context as stipulated. Yet the stipulated context is devoid of the evidence that the audience would need to know what the speaker means by testimony or induction. Although there are in fact various ways in which the audience can come to know what a speaker meant without knowing what they said, they all rely on features of the context that are not stipulated in the cases we will engage with in the following Chapter.

Slips of the tongue provide another kind of case that fall outside the domain of the Utterance and Expression Constraints. Suppose that I intend to say that John loves soup but, whether due to a slip of the tongue or a flawed understanding of English, I utter instead ‘John loathes soup’. If the audience recognises my mistake, then communication can succeed, the natural explanation being that the audience interprets based on what they take me to have intended to utter and say, rather than what I actually uttered and said. Extending the case so that nothing is uttered or said, a slip of the tongue might lead me to utter a meaningless string of sounds (that is, one that is not associated with any meaning in the relevant language). While something may be uttered, nothing is said, yet my audience may very well be able to know what I meant by interpreting as though I had uttered whatever sentence I intended. While it is clear that such cases are significantly different from those we will engage with in the following Chapter, we can bring them into the fold by replacing the Epistemic Content Constraint with

Epistemic Content Constraint*: If the audience cannot know what the speaker intended to say, then they cannot know what the speaker intended to communicate.

Although I might not intend to communicate what I actually say, I may still intend to communicate the proposition that I *intended* to say. Given that misspeaking will not feature in the cases we will consider, however, we will stick with the original formulation of the Constraints.

Exploiting ambiguity

Grice (1989, 35-36) notes an interesting category of cases, which he characterises as flouting his Maxim of Manner: be perspicuous! Part of being perspicuous is avoiding ineliminable ambiguity, yet Grice suggests there are some cases in which ineliminable ambiguity can be utilised to aid communication. Grice's example comes from William Blake: "Never seek to tell thy love, Love that never told can be." Rather than deal with the imperative, Grice considers "I sought to tell my love, love that never told can be", which Grice considers multiply ambiguous, as 'my love' can refer to the emotion of love, of the object of love, while 'love that never told can be' may be interpreted as saying that the feeling cannot be told, or that the feeling, once told, must cease to exist. Because of the skill of the poet and the fact that the ambiguity is kept up throughout the poem, Grice takes it that Blake uses ambiguities to convey both what is said on the first interpretation and what is said on the second.

The example is somewhat *recherché* and so there is plenty of room for dissent. We might deny, for example, that Blake's poetry is really communicating anything, or that, if it is, it is not bound by the usual Maxims of Conversation and so cannot be implicating anything. A second example is perhaps even more singular. Grice's tells the apocryphal tale of a British general who, having exceeded his mandate by capturing the province of Sind, reported his accomplishment by the message '*Peccavi*', supposedly Latin for 'I have sinned'. Although the Latin word is unambiguous, Grice suggests that an ambiguity arises after translation. The sounds you would make in uttering 'I have sinned' are at least phonetically ambiguous between that sentence and 'I have Sind'. Grice suggests that, in the right circumstances, the speaker may utilise the ambiguity to communicate what would be said on one disambiguation, as well as what would be said on the other.

While the examples are peculiar, the general point seems fair enough. If the speaker utters a sentence that is irreducibly ambiguous, the speaker says nothing at all and the audience cannot know what the speaker has said, yet the audience may be able to know what the

speaker intends to communicate. Indeed, the point extends beyond ambiguity. There are other ways of failing to say anything that can force the audience to consider potential implicatures. By using the term ‘She’ to talk about a male, a speaker may implicate that their referent exemplifies some property or properties that the speaker associates with females. On the assumption that ‘She’ can only be used to talk about females, the speaker has failed to say anything, and indeed, the failure to say anything may be central to the effect that the speaker wants to achieve.

Again, we can safely set these cases to one side. While there are cases in which the failure to say anything is central evidence of the speaker’s implicature, there is no similar and plausible analysis along these lines of the cases from the following section. For what is it worth, there is evidence that even Grice himself saw these cases as somewhat deviant. Grice (1991, 39) for example, says that the calculation of an implicature requires “knowledge of what has been said”. As nothing is said in these cases, they fail to meet a central requirement of Grice’s theory, showing that Grice did not allow these cases to drive his theory of implicature. In the cases of underdetermination from the next Chapter, it will be clear that the speaker is not relying on any such supplementary interpretive procedure. Indeed, it is implausible to suggest that the speaker is relying on an implicature of any kind.

Contextual salience

By ‘contextual salience’ I do not mean ‘salience *in* a context’ but rather ‘salience *of* a context’. It is commonly assumed that what is said by an utterance is determined by the actual context in which the utterance takes place. Even assuming this traditional view, it must be allowed that contexts other than the actual context can become more salient than the actual context, leading audiences to interpret the speaker as though they said, not what is said in the actual context, but what is said in the more salient context. We are creatures capable of both mistaken belief and imaginative pretence, and both of these capacities can lead to situations in which we take another context to be more salient than the actual context in which the utterance takes place.

Mistaken belief

Consider the following case.

It's Daniel's office hour, and he has gone to lunch. In order to save Daniel's students some time, Frank writes 'I am out to lunch' on a scrap of paper and attaches it to Daniel's office door. To the students who turn up for Daniel's office hour, the note communicates that Daniel is out to lunch and they leave without wasting their time by knocking and waiting for a response.¹⁹

According to a standard treatment, any occurrence of 'I' refers to its author; in this case the author of the note. The note can therefore only be used to say that the author of the note is out to lunch. Given that the note was in fact authored by Frank, it says that Frank is out to lunch, yet it communicates that Daniel is out to lunch. We have already seen that what is communicated can come apart from what is said when some additional interpretive procedure is involved, such as the resolution of implicature, but such an analysis of this case is a nonstarter. The interpreter does not first identify what is said as the proposition that Frank is out to lunch and then derive the proposition that Daniel is out to lunch on the basis of what was said. If the students believe that Daniel authored the note, there needn't be any way for them to recognise what was actually said.

¹⁹ These are similar to cases noted in Weatherson (2002). Weatherson argues against Romdehn-Romluc's (2002) account of indexicals by deploying cases in which it seems possible to trick a reasonable and attentive hearer into mistaking the content of indexicals. I may, for example, write a note that reads 'there is a faculty meeting today' and bury it at the bottom of a colleague's in-tray, among correspondences from two days earlier. According to Weatherson, the note says that there is a faculty meeting on the day of writing, but my colleague can be expected to mistakenly take it to say that there was a faculty meeting two days earlier. One way to explain this case is to accept that 'today' is associated with the linguistic rule that this term always refers to the day on which it is produced for interpretation, but that my colleague interprets with respect to a shifted context in which they read two days earlier. The note says that there is a faculty meeting today, but communicates to my colleague that there was a faculty meeting two days previous. The cases that interest us here are a superset of Weatherson's cases of trickery.

These cases highlight a limitation of Kripke's (1977:264) definition of speaker's referent, according to which "the speaker's referent of a designator" is "that object which the speaker wishes to talk about, on a given occasion, and believes fulfils the conditions for being the semantic referent of the designator". Kripke offers this definition only "tentatively", but offers nothing less tentative. This characterisation of speaker's referent clearly fails to fit this case, or others that will be presented later. In this case, Frank knows that Daniel fails to fit the conditions for being the semantic referent of 'I', but nonetheless, Daniel appears to be Frank's referent. Kripke (1977: 273, footnote 22) does briefly consider extending the analysis to include some "exceptional" cases mentioned in Donnellan (1966), but concludes against doing so, partly for simplicity, and partially because he considers them akin to cases of sarcasm. The cases presented here are not particularly exceptional, and not at all like sarcastic speech.

It seems clear how the sentence is interpreted: the referent of ‘I’ is taken to be Daniel rather than Frank. Corazza, Fish, and Gorvett (2002: 4-5) take this case to raise “an intriguing and important question of just how the reference of ‘I’ is determined” and conclude that the standard treatment is wrong: an occurrence of ‘I’ needn’t always refer to its author. In this case, they say, the occurrence of ‘I’ on the note refers to Daniel. But while they are correct that Daniel “is certainly what the audience of the utterance, the students, take [‘I’] to refer to”, there is a simple explanation of this fact that is consistent with the standard treatment of this term. Standardly, a note in the first-person attached to an office door would be written by the regular occupant of the office, so the students who turn up to Daniel’s office can be expected to have the mistaken, though entirely predictable, belief that Daniel authored the note himself. Even if the students believe that ‘I’ refers to the author of the note and interpret accordingly, they will take ‘I’ to refer to Daniel rather than Frank, due to their mistaken belief as to who authored the note. Frank constructs his note in the recognition that the students will interpret it relative to a context in which Daniel is the author, thereby taking it to say that Daniel is out to lunch, which is just what Frank means to communicate.

For the students who mistakenly take the note to have been written by Daniel, the context in which Daniel authored the note is more salient than the actual context in which Frank authored the note, which leads them to interpret as though the note said that Daniel is out to lunch, rather than what it actually said: that Frank is out to lunch. If what is said is sensitive to features of context (such as authorship) about which the audience can be mistaken, then we can construct cases in which the audience mistakes what is said. If a speaker can predict the way in which the audience will mistake the context, then they can use this to their advantage, intending to communicate, not what they actually say, but what they would have said in another context.

If the traditional account of ‘I’ is correct, then Frank does not intend to communicate what he says and the audience needn’t know what was said to know what Frank intended to communicate. The cases can be expanded into cases in which nothing is said. Take for example, the variation of the case in which the note is co-authored by Frank and Lisa. If a single occurrence of ‘I’ can only refer to a single individual, then the occurrence on Frank and Lisa’s note refers to no one. Frank and Lisa each have equal claim to be the referent, so ‘I’

refers to Frank if and only if it refers to Lisa. As ‘T’ can only refer to one of them, it therefore refers to neither.²⁰

Interpretive pretence

The phenomenon is not restricted to cases in which the audience mistakes the nature of the context. Even an audience who knows exactly what is said can choose to interpret relative to contextual conditions other than those which actually obtain.

Suppose that Frank makes a cake for a party he and Daniel are attending. While it cools on the table, he leaves a note which reads ‘Hello. I am a cake for the party. Please don’t eat me!’ The note does not communicate that Frank is a cake for the party. Rather, it communicates that the cake on the table is a cake for the party.

Again, we have a case in which the referent assigned to ‘T’ by the interpreter is not the referent of that term according to the standard treatment. For the case to be consistent with the standard treatment, however, we needn’t follow Corazza, Fish, and Gorvett’s reasoning and conclude that the deceptively simple term ‘T’ is associated with a surprisingly complex meaning that allows for reference, not only to the author of a note, but also to a cake prepared by the author of a note. Instead, we need only recognise that we are capable of interpreting sentences relative to contexts that we know fail to be actualised in order to facilitate communication.

A natural way to explain the case is to suppose that the note invites Daniel to imagine a (mildly) amusing context in which the cake itself authored the note, which goes some way towards explaining why such notes are (mildly) amusing. According to the standard treatment of ‘T’, the proposition determined by the sentence, relative to a context in which the cake authored the note, is the proposition that was in fact communicated: the proposition that the cake on the table is a cake for the party. In this case, the author doesn’t intend to communicate the proposition they say – the absurd proposition that the author of the note is a cake – but the proposition that would have been said in the salient (non-actual) context.

This case demonstrates clearly that interpreters often engage in a kind of pretence during interpretation, but the phenomenon is not restricted to (mildly) comical cases. We can

²⁰ This is an instance of the non-epistemic underdetermination that we will come to in Chapter 3.

easily construct variants of the Mistaken Belief case that exhibit a similar reliance on interpretive pretence:

The audience knows that Daniel often puts a note on his door when he leaves his office, but that sometimes Frank does it for him. A note on Daniel's door reads 'I am out to lunch' and communicates to the audience that Daniel is out to lunch.

In this case, the audience doesn't have any belief as to the identity of the actual author, so can't proceed by trying to identify what is said by the sentence in the actual context. Instead, the case can be explained in similar fashion to the case of the cake. Given the convention that first-personal notes of this sort are generally left by regular occupants of the offices to which they are attached, the note invites the interpreter to suppose that Daniel wrote it and the interpreter can do so without believing that this is how things actually stand.²¹ Whether the note was in fact written by Frank or Daniel, the author intended the audience to interpret as though the note was written by Daniel.

The same explanation can be given of cases in which the audience has (to all intents and purposes) perfect knowledge of the actual context. Suppose that students see Frank write 'I am out to lunch' and attach the note to Daniel's office door. In this case the audience is at least capable of interpreting the sentence on the note relative to the actual context, but the mere ability to do so doesn't require that the audience exercise this ability. Rather, the interpreter might carry out just the procedure they apply in the other cases, the main difference being that in this case the interpreter is fully aware that the proposition they take to be communicated is not what is said by the note.

These cases are not restricted to written language. The same result holds for the case in which Frank is in Daniel's office and says, in his impeccable Daniel impression "I am going out to lunch". If the students believe that Daniel is speaking, they will naturally interpret Frank's utterance relative to a context which, according to the standard account of 'I', determines the proposition that Daniel is going out to lunch. Even if the students know that

²¹ We can appeal to such a convention or intention without, as Corazza et al. (2002, 10-11) do, giving it a semantic role in the determination of what is said. This is more like the *pre-semantic* role noted in Perry (1998, 593).

Frank is the one speaking, they may still be able to recognise his intention and evaluate his utterance relative to such a context.

Nor are these cases restricted to the term 'I'. The central feature of these cases is only that the audience interprets relative to contexts other than the one that actually accompanies the utterance. Predelli (1998a,b) presents a range of cases intended to dislodge the traditional view that the referents of indexicals are determined by the actual context in which they are uttered, that is, the view that 'I' refers to the speaker or writer, 'now' refers to the time of utterance, and 'here' refers to the location of utterance. Predelli suggests instead, that the references of these terms are determined by the speaker's intentions: 'I', 'here', and 'now' refer to whichever individual, location, and time the speaker intends them to refer to.

“Consider the anecdote of Jones, who expects his wife to come home at six, and writes ‘I am not here now’ at four, with the intention of informing Mrs. Jones that he is away from home at six - or, if you prefer, imagine that he records ‘I am not here now’ on a tape, expecting his wife to activate the tape-recorder upon her arrival. However, Jones's wife is late, and she only reads the message (or turns on the tape-recorder) at ten. Clearly the vicissitudes of Mrs. Jones do not affect the content of Jones's message. Intuitively, such content is to be established with respect to the time of intended decoding, and not with respect to the time when decoding actually took place.” Predelli (1998 a, 110)

Predelli takes it to be clear that Mrs Jones's lateness cannot affect the content of the note but this assumption begs the question against Predelli's opponent, who argues that the reference of 'now' is the time of utterance and that Mrs Jones lateness affects the time at which the utterance occurs. The opponent in question can be represented by Sidelle (1991), who suggests that recorded messages are means of deferring utterances until the time at which they are interpreted. The context relevant to determining the content of Jones's note is therefore the context in which it is read by Mrs Jones. If Sidelle is correct, then “the vicissitudes of Mrs Jones” are clearly very relevant to the content of Jones's message. If Mrs Jones arrives home later than Mr Jones expected, the content of 'now' will be other than Mr Jones intended.

Perhaps, however, Predelli's meaning is not this abrupt denial of Sidelle's view, but the very plausible claim that Mrs Jones's lateness will not affect Mrs Jones's interpretation of Mr

Jones's communicative intention. According to Sidelle's suggestion, the note says that Mr Jones is not in their home at ten o'clock, but Mrs Jones, knowing that she has arrived later than expected, will take Mr Jones to intend to communicate, not what is said by the note according to the traditional view, but what would have been said by the note in the context Mr Jones intended it to be read in. This much, however, even the defender of the traditional account can accept, given that Mrs Jones's knowledge of Mr Jones's expectations can be expected to raise to salience the context in which Mrs Jones arrives home at six o'clock, as Mr Jones expected. Interpreting relative to this context, Mrs Jones will take Mr Jones to intend to communicate, not what was actually said, but what would have been said in the salient context according to the traditional view. We can therefore achieve the desired result without abandoning the traditional view of 'now' in favour of Predelli's alternative proposal.

The point extends to other cases that Predelli presents in support of his view. Take cases of the historical present. For example, the following utterance made in 1996:

"It is 1796. Napoleon, now commander of the French troops in Italy, defeats the Sardinian forces and turns against Austria." (1998 a, 112)

Predelli correctly notes that the author is surely "not conveying the false information that Napoleon is the commander of the French troops in 1996". Certainly not, but again, the traditional view is committed only to the position that the note *says* that Napoleon is the commander of the French troops in 1996, not to the position that this is what audiences will take the speaker as intending to communicate. The traditional view is not committed to this because the view has to recognise that contexts other than the actual context might be raised to salience, leading the audience to interpret the speaker as intending to communicate, not what they actually say in the actual context, but what they would have said in another context. As Predelli notes, the speaker "narrating Napoleon's situation from the point of view of 1796" and this obvious pretence, explicitly signalled by the false "It is 1796" raises to salience the context of that year, relative to which any use of 'now' picks out, according to the traditional view, a time in 1796. According to the traditional view, the speaker literally says the false proposition that Napoleon is the commander of the French troops in 1996 but it is consistent, and to my mind extremely plausible, to maintain that the speaker intends to communicate

some proposition other than what they say, specifically the true (I assume) proposition that Napoleon is the commander of the French troops in 1796.

Similar explanations are available of Predelli's (1998a, 113-114) cases involving 'here' and 'you'. Again, Predelli notes that these terms can often be used to communicate something other than what, according to the traditional view, is said in the actual context. Again, however, there is independent motivation for the view that speakers often communicate what they would have said in some context other than the actual context, which they raise to salience either by mistaken belief on the part of the speaker or their audience, or an imaginative pretence of either side.

While these replies are not sufficient to show the falsity of either the view of Corazza, Fish, and Gorvett, or Predelli, they undermine some of the cases used to support their views, which are explained at least as well on the traditional competitor to their accounts. When arguing against a certain view of what is said, it is well to keep in mind that one's opponent has contextual salience among their tools. Note also that there is nothing in principle that reserves the tool for the defender of the traditional story about indexicals. For many different view of the context that determines what is said, there may be cases in which audiences are intended to interpret relative to another context entirely.²²

²² The two cases collapse on Predelli's view, as what is said is determined by the context the speaker intends as relevant. Talk about 'the salient context' might more precisely be put in terms of 'the salient type of context'. Speakers will often not care precisely which context the audience interprets relative to, so long as it will determine the correct proposition. Depending on how finely we distinguish between contexts, this might be as true of standard cases as it is of the cases in which a nonstandard context is salient.

Chapter 2: Underdetermination arguments

Schiffer against the ‘naïve’ hidden-indexical theory of descriptions

In the last few decades, the problem of underdetermination has been presented by a number of different authors, who collectively dismiss a variety of linguistic and philosophical theories.²³ Perhaps the most prolific advocate of underdetermination arguments is Schiffer. We will focus here on his argument against the ‘naïve’ hidden-indexical theory of descriptions, according to which the sentence uttered by an articulation of a definite description like ‘The *F* is *G*’ contains an indexical not pronounced at the level of what is articulated (hence ‘hidden’) that refers to a contextually-determined property *H*.²⁴ The uttered sentence is then given a standard Russellian analysis, so that what is said is true if and only if there is one and only one thing which satisfies *F* and *H*, and it also satisfies *G*.²⁵

Different values for the hidden indexical determine different propositions. If the hidden-indexical theory is correct, then what is said by an utterance of the form ‘The *F* is *G*’

²³ Schiffer presents variations of the argument against an array of views. For argument against the “*hidden indexical description theory of indexicals*” see Schiffer (1995, 117), for argument against the “description theory of indexicals” see Schiffer (1981, 76-78), for argument against the hidden-indexical theory of belief reports see Schiffer (1992, 512- 514; 1994, 286-287; and 1995, 109-112), and see Schiffer (1992, 516) for argument against the description theory of incomplete definite descriptions. Similar arguments have been presented by a number of other authors. Wettstein (1981), whose arguments we will get to shortly, argues against syntactic ellipsis analyses of definite descriptions and notes that a similar style of argument is used by Donnellan (1968, 204, footnote 5) to suggest that incomplete definite descriptions such as ‘the table’ are often to be read referentially. Wettstein’s arguments against syntactic ellipsis analyses are echoed in Recanati (1996, 449). Donnellan (1966, 294, footnote 10) raises similar points in arguing against Strawson’s view (1954, 230) that we can amend the statements of mistaken speakers in accordance with their ‘guessed intentions’. Borg (2002, 493-495) raises the problem in arguing against a descriptive analysis of ‘deferred’ demonstratives. Kripke (1977, 255) alludes to the problem of underdetermination as an objection to Russell (1905), and the suggestion is disputed by Bach (1981, 239), who agrees with Donnellan (1968). Reimer (1992) presses the problem of underdetermination against Neale’s (1990) “explicit” approach to incomplete definite descriptions.

²⁴ Schiffer (1995, 114) opts to present the theory in terms of the meaning rule: “Utter “The *F* is *G*” only if there is a property *H* such that you mean that the *F* and *H* is *G*”. The presentation is peculiar for a number of reasons. Firstly, the specified rule prohibits non-literal speech. Secondly, it doesn’t distinguish the hidden-indexical account from a syntactic ellipsis account. Thirdly, it yields the bizarre result when we substitute in the description “the guy”, which is used in Schiffer’s example, that one should utter the “The guy’s drunk” only if there is property *H* such that you mean that the guy and *H* is drunk. It is unclear what one means if they mean that, e.g., the guy and author of *Smells and Ticks* is drunk. The presentation in the body of the text mimics Schiffer’s presentation of another hidden-indexical theory offered in the same paper: the hidden-indexical theory of belief reports. Schiffer (1995, 108-109). I focus on the hidden-indexical theory of descriptions, rather than belief reports, because the Pergola example he uses to illustrate the argument against the hidden-indexical theory of descriptions is superior to the Bardot example he uses to illustrate the argument against the hidden-indexical theory of belief reports in that it offers at least some degree of context.

²⁵ See Russell (1905) for Russell’s original statement of the view.

depends on the value of the hidden-indexical H. The audience can know what is said by an utterance of the form ‘The F is G’, therefore, only if they can know the value of the hidden indexical H. There are, however, very few cases in which it is plausible to suppose that this is possible. Schiffer (1995: 114-115) illustrates the problem of underdetermination by way of the following example:

“Imagine that you and I are in the audience awaiting a talk by the eminent philosopher Ferdinand Pergola. The professor arrives, and you say to me ‘I’ll be damned! The guy’s drunk’. Even before your utterance it was mutually evident to us that we had knowledge of the professor under numerous shared definite descriptions—*the author of Smells and Tickle*s, *the only man within sight wearing a yellow jacket and red golf pants*, *the man we are waiting to hear*, *the man now staggering up to the podium*, and the list, in any realistic situation, will go on and on.”

In such a situation, says Schiffer, he cannot see how the audience could identify “any one definite description, however complex, as *the one* that figured in the proposition you asserted”.²⁶ Given the context Schiffer describes, there are a number of different but equally viable candidates for the value of the hidden-indexical – one property for each of the descriptions we share. Yet, while equally viable in the context, each of these candidates for the value of the hidden-indexical determines a different candidate for what is said. If the value of the indexical is the property of *being the author of Smells and Tickle*s, then what is said is that there is one and only one guy who authored *Smells and Tickle*s and he is drunk. If the value of the indexical is the property of *being the man we are waiting to hear*, then what is said is that there is one and only one guy who is the man we are waiting to hear and he is drunk.²⁷ These two candidates for what is said are distinct, given that one may be true while the other is false, as is the case in the world

²⁶ Schiffer uses ‘asserted’, rather than ‘said’. For Schiffer, saying that p entails meaning that p. We do not share that assumption here. The Encoding Model from the previous Chapter allows the use of terms like ‘said’ to describe propositions that are not meant. Given that this is a case of literal speech, however, the Encoding Model will endorse the claim that the speaker means what they say in the Pergola case.

²⁷ It is of course consistent for the value of the hidden-indexical to be the property of *being the author of Smells and Tickle*s and *being the man we are waiting to hear* but this is just another candidate for the value of the hidden-indexical; one that is no more supported by the available evidence than any other. This is what Schiffer means by saying that no description, “however complex”, clearly supplies the value for the hidden-indexical.

in which there are two distinct guys, one who is drunk and authored *Smells and Ticks*, and one who is sober and is the man we are waiting to hear.

These candidates for what is said are on a par in terms of the usual suspects like complexity, availability, informativeness, and contextual salience. Neither of these candidates is significantly more complex than the other, as opposed to the proposition *that there is one and only one guy who is the author of Smells and Ticks, the man within sight wearing a yellow jacket and red golf pants, the man we are waiting to hear, and the man now staggering up to the podium*, which is plausibly so complex as to not be worth considering as a candidate for what is said. In part due to their comparable simplicity, these candidates are equally available to audiences. Without some additional context, there is no reason to suppose that one of these candidates should be more readily available to the mind of the interpreter. Similarly, neither has been raised to a position of special contextual salience. The level of salience of each candidate is purely a result of the fact that they are derived from properties known to be distinctive of Pergola. Although the candidates encode different propositions, they seem broadly similar as regards the ‘amount’ of information encoded. An alternative that does markedly poorly in this respect is the proposition that the guy who is drunk is drunk. Finally, although the candidates encode different propositions, either proposition would be adequately informative for the task at hand, that is, picking out Pergola and predicating a property of him.

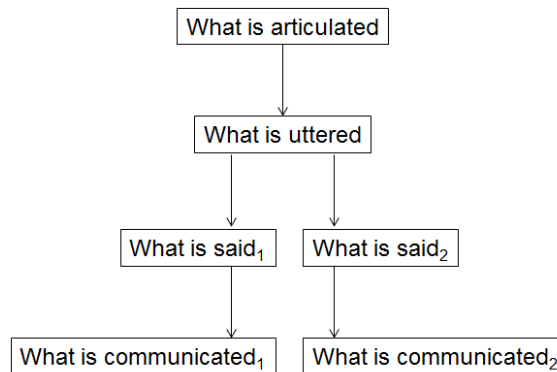
There may of course be contexts in which one of these candidates is clearly superior to the other candidates for what is said. Perhaps we have just been talking about Pergola as ‘the author of *Smells and Ticks*’ for example, in which case the most contextually salient and accessible interpretation may be that according to which you say that the guy who authored *Smells and Ticks* is drunk. Schiffer’s argument, however, requires only that there can be cases in which a number of different propositions are equally viable candidates for what was said.

The audience cannot know what is said, yet Schiffer assumes, as per the Encoding Model, that the speaker means what they said: “there is some contextually-salient property *H* such that in uttering “The guy’s drunk”, you implicitly [said], and thus meant, *that the H guy was drink*, where *the H guy* was one of those shared definite descriptions.”²⁸ If the audience cannot know the value of the hidden-indexical, they cannot know what is said and, by the Epistemic

²⁸ Schiffer uses ‘stated’, rather than ‘said’. See the previous footnote.

Content Constraint, cannot know what the speaker intended to communicate. If the value of the hidden-indexical is the property of *being the author of Smells and Ticks*, then the speaker says and intends to communicate that the guy who authored *Smells and Ticks* is drunk. If the value of the hidden-indexical is the property of *being the man we are waiting to hear*, then the speaker says and intends to communicate that the guy who is the man we are waiting to hear is drunk.

The situation can be represented diagrammatically:



According to this diagram, there is a single sentence uttered in the context: ‘The H guy is drunk’, where ‘H’ is a hidden-indexical. Given that there are two equally viable candidates for the value of the hidden-indexical, however, there are two candidates for what is said, each of which determines a different candidate for what is communicated. The context admits of two interpretive paths, each leading to a different communicated content. If the audience cannot know which path to take (that is, if the audience cannot know what was said) then they cannot know what the speaker intended to communicate.

If the speaker intends to communicate what they say, then the audience cannot know what is said and consequently cannot understand the speaker, as “it ought to be a consequence of any acceptable theory of meaning that if a speaker, speaking literally, meant p in uttering u , then one who understood his utterance perfectly well took him to have meant p in uttering u .”²⁹ Schiffer further assumes that the audience understood the speaker “perfectly well”, leading him to conclude that the speaker “*could not* mean what the naïve theory requires”, that

²⁹ Schiffer allows that something may be successfully communicated by the utterance. “To be sure,” says Schiffer (1995, 115-116) “the speaker in our example was also using “the guy” referentially, and thus meant an object-dependent proposition involving Pergola”. If the speaker did intend to communicate such a proposition, Schiffer’s argument allows that the intention can be successful, but see the discussion of the Problem of the Many in the following Chapter for reasons to doubt this.

is, the speaker could not have intended to communicate any proposition of the form required by the naïve hidden-indexical theory of descriptions. If the hidden-indexical theory were correct, Schiffer concludes, the speaker cannot mean what they say by use of the definite description. As the argument rests on little more than the assumption of a very limited stock of knowledge shared between the speaker and their audience, the conclusion will generalise to the vast majority (and perhaps all) uses of definite descriptions.

Schiffer's indeterminate reference

According to Schiffer's argument, if the naïve hidden-indexical theory of descriptions were true, then we would not be able to know what speakers intend to communicate when the speaker and their audience share even very modest common knowledge. Schiffer therefore concludes that naïve hidden-indexical theory of descriptions should be rejected. Schiffer (1995, 115) suggests that underdetermination is avoided (and the naïveté of the view eliminated) by allowing the speaker to make "an *indeterminate* statement." The view is introduced by way of the following example:

"Suppose you call Ernie Lepore in New Brunswick and ask him where Jerry Fodor is. 'He's here', Ernie replies. To what does the utterance of 'here' refer? To New Brunswick? To Rutgers University? To Douglass Campus? To Davison Hall? To Ernie's office? The example is underdescribed, but even if I fully describe it, there need not be a definite answer. Almost certainly, Ernie's utterance of 'here' does not refer to some definite region of space. The word is being used to make a vague or indeterminate reference. (112)

Ernie hasn't "definitely" referred to anywhere in particular. As what is said is determined by Ernie's reference, he hasn't definitely said anything. On the assumption that he means what he says, therefore, "there is no proposition Ernie definitely meant". The most that can be said is that he "*sort-of-meant*, or vaguely meant, all the propositions that could be used to precisify his indeterminate statement". (112-113)

Extending the analysis, Schiffer says that "In the Pergola example, you did not definitely mean any general proposition in uttering 'The guy is drunk', but that you *sort-of-meant*, or vaguely meant, several general propositions, one for each definite description that

could be used to sharpen what you vaguely meant. Schiffer suggests supervaluation over all precisifications to determine the truth value of the utterance. In this context, “truth is super-truth”³⁰ and an utterance is true “just in case it is true under every admissible sharpening of what you meant, false just in case it is false under every such admissible sharpening, and neither true nor false if it is true under some admissible sharpenings while false under others.” (115)

According to Schiffer, the speaker did not mean any of the candidates for what is said, but sort-of-meant, or vaguely meant, each one of them. There are several ways we might interpret these remarks.³¹ According to one account, the indeterminacy of what is meant follows the indeterminacy of what is said. This follows from the Encoding Model. If the speaker says that P_1 then the speaker intends to communicate that P_1 . If the speaker says that P_2 then the speaker does not intend to communicate that P_1 but rather that P_2 . If it is indeterminate whether the speaker says that P_1 , then it is indeterminate whether the speaker intends to communicate that P_1 .

This is, however, more a restatement of the problem, than a solution. A solution has to tell us how the audience can come to know what is meant in the face of this indeterminacy. If it is indeterminate which of a number of candidates is said, then the audience cannot come to know what was said and consequently cannot know what was meant. If, on the other hand, the audience can know what was meant, then what was meant cannot be indeterminate.³² If it is neither true nor false that P_1 is meant, then the audience cannot know that P_1 was said and the audience cannot know that P_1 was not said.

Perhaps simple restatement of the problem can be avoided by interpreting Schiffer as suggesting that what is said may be not only epistemically indeterminate, in that the evidence available to the audience fails to determine what is said, but indeterminate in some deeper sense. One way to make sense of the claim is via the ‘truth-maker gaps’ of Greenough (2008).

³⁰ In the language of Fine (1975)

³¹ As those of Blackburn (1988, 271), who is cited by Schiffer and recommends the similar response that “there is some more or less vague *class* of propositions, each determined by some expansion of the description” and the speaker “asserts the truth of the propositions in the resulting class (or, perhaps, of some weighted majority of these)”.

³² Another suggestion that may, perhaps, be extracted from Schiffer’s words is that nothing is said in the Pergola case. There are a number of candidates that are equally supported by the context, each of which would communicate the same information. This is the solution that I shall defend in Chapter 4.

‘ P_1 is said or P_1 is not said’, where P_1 is one of the candidates for what is said, is true by bivalence, but its truth is ungrounded in that there is no fact of the matter either way. It is neither a fact that P_1 is said, nor a fact that P_1 is not said, and the same holds for the rest of the candidates. As the speaker means what they say, but it is indeterminate what is said, so it is indeterminate what is meant. ‘ P_1 is meant or P_1 is not meant’ is true but there is no fact of the matter either way.

This position, however, is not consistent with the rest of Schiffer’s commitments. According to Schiffer (115), the speaker meant that p in uttering u only if “one who understood his utterance perfectly well took him to have meant p in uttering u ”. The interpreter in the Pergola case can understand the speaker’s utterance perfectly well but cannot, according to Schiffer, take the speaker to have meant any of the candidates for what is said. Contrary to the suggestion that what is said is indeterminate in that it is neither definitely said nor definitely not said, the speaker *definitely* did not mean any of these candidates, because the audience *definitely* needn’t take the speaker to have meant any of these candidates in order to understand the speaker.³³

Finally, we might interpret Schiffer as suggesting that what is said and meant is entirely determinate, but that the speaker says and communicates a vague proposition. This makes sense of the analogy with Ernie Lepore’s use of ‘here’. In using the term, Ernie is not referring to some “definite region of space” but to some vague region that admits of various precisifications. Extending this proposal to the hidden indexical theory of descriptions, we can say that underdetermination in the Pergola case arose from the assumption that you refer to some particular property by your use of the hidden indexical, when in fact you referred to some vague property. The problem of underdetermination is avoided because there is a vague candidate for what is said that fits the bill better than the precise candidates we have been considering heretofore.³⁴

We might question whether the analogy extends cleanly from the vague spaces that we are acquainted with to the vague properties that this response requires, but for our purposes it will be sufficient to note that the appeal to vagueness doesn’t help us to avoid the problem of

³³ This response comes from Buchanan (2010, 355).

³⁴ These ‘precise’ candidates, being defined using natural language, will certainly inherit some vagueness. The suggestion is that they are less vague, or vague in a different way, from the best candidate for what is said.

underdetermination, even in the case of ‘here’. The problem concerns higher-order vagueness.³⁵ There are many precise candidates for the reference of ‘here’, but there are also many vague candidates. Ernie’s reference is vague, right enough, but just how vague? How far do the borderline cases extend? Different answers to this question yield different vague candidates for what is said and the problem of underdetermination rears its head once more. Candidates for Ernie’s vague reference include the vague space which counts the other side of his office window as a borderline case of ‘here’, and the vague space which doesn’t count that as a borderline case. We know that Fodor inhabits one of these spaces if and only if he inhabits the other (given that he isn’t plastered against the window) so either is an equally viable candidate for the location Ernie is referring to, yet Ernie says, in therefore intends to communicate, different propositions depending on which of these vague spaces he is referring to. Vagueness should make an appearance in any complete theory of interpretation but it doesn’t help us to resolve the problem of underdetermination.

All in all, Schiffer’s talk of indeterminate reference, however interpreted, fails to provide a means of avoiding the problem. If the speaker means some vague proposition, the audience cannot know which vague proposition is meant. If what is meant is indeterminate, then the audience cannot know what is meant. In Chapter 4, I shall explain how we can retain determinate speaker-meaning in the face of indeterminacy as to what is said. For now, we can consider some further examples of the problem.

Wettstein against Russell’s analysis of definite descriptions

Another early proponent of the problem of underdetermination is Wettstein (1981, 244), who argues against Russell’s theory of definite descriptions, according to which “when one utters ‘The F is G’, however the description is used, the same proposition is [said]: there is one and only one F and it is G.”³⁶ *Prima facie* counterexamples to this theory are abundant in the form

³⁵ Schiffer (1995, footnote 10) explicitly ignores higher-order vagueness, given that he wants to pose alternative problems for the view.

³⁶ Wettstein uses ‘asserted’, rather than ‘said’. No damage is done by the substitution. Unless Wettstein intends Russell’s theory to prohibit sarcasm and figurative talk, he must be thinking of what is asserted along the lines of what is said. If it is to be at all plausible, this description of Russell’s theory must be restricted to sentences devoid of context-sensitive terms. As Stanley and Gendler Szabó (2000, footnote 23) point out, Wettstein’s argument is flawed as “an argument that Russell’s theory of descriptions fails” (Wettstein (1981, 245)) because he fails to consider other responses to the *prima facie* problem. Our interest, however, is whether Wettstein’s argument

of incomplete definite descriptions such as ‘the table’ which clearly fail to denote a unique object. According to the Russellian theory, by uttering ‘the table is covered with books’ one says that there is one and only one table and it is covered with books. Given that there are obviously a great many tables in the world, this proposition is clearly false, yet utterances of this sentence can be used to communicate truths.

Wettstein considers the *syntactic ellipsis*³⁷ response to this *prima facie* problem according to which an articulation of ‘the table is covered with books’ should be interpreted as an utterance of a sentence in which the domain of quantification is specified explicitly, such as ‘the table *in room 209 of Camden Hall at t₁* is covered with books’. If there is indeed one and only one table in room 209 at t₁, and it is indeed covered with books, then this sentence says something true, according to Russell’s analysis of definite descriptions.

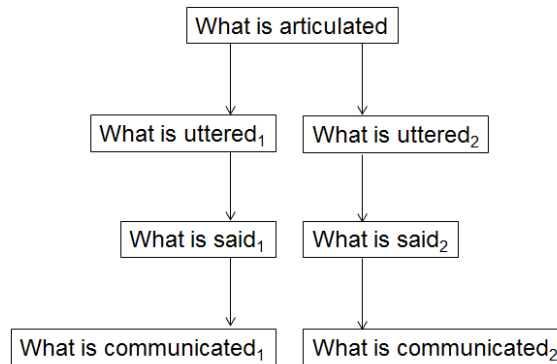
Against this response to the *prima facie* counterexamples, Wettstein (1981, 246) poses the problem of underdetermination. There are any number of “nonsynonymous, uniquely denoting descriptions (for example ‘the table in room 209 of Camden Hall at t₁’, ‘the table at which the author of *The Persistence of Objects* is sitting at t₁,’ etc.)” by which the table could be picked out in any realistic context. In most contexts, therefore, there will be a number of candidates for the sentence uttered – one for each description by which the table can be identified. If these descriptions were synonymous, then underdetermination of what is uttered wouldn’t entail underdetermination of what is said – despite a number of candidates for what is uttered, there would be a single candidate for what is said – but “Since these more complete descriptions are not synonymous, it follows that each time we replace the indefinite definite description ‘the table’ with a different one of these ‘Russellian’ descriptions, it would seem that we obtain an expression for a *different proposition*, one that gets a different analysis via the theory of descriptions” and one that constitutes a different candidate for what is said.

According to the syntactic ellipsis defence of the Russellian analysis, there are a number of candidates for what is said. This leads Wettstein (1981, 246) to ask “which of these more complete (or Russellian) descriptions (or conjunction of such descriptions) is *the correct one*, the one that actually captures what the speaker intended by his use of the indefinite

succeeds as an argument against a syntactic ellipsis defence of Russell’s theory, rather than whether it refutes Russell’s theory entirely.

³⁷ This terminology comes from Stanley and Gendler Szabó (2000).

definite description ‘the table’.” Here, Wettstein clearly endorses the Encoding Model, according to which the basic case is one in which the speaker says the proposition that they intend to communicate, and the audience can therefore know what the speaker intended to communicate only if they can know what the speaker said. As the audience cannot know what is said, the audience cannot know what is meant. If the speaker uttered ‘the table in room 209 of Camden Hall at t_1 is covered with books’, then according to the Russellian account, the speaker said and intended to communicate that there is one and only one table in room 209 of Camden Hall at t_1 and it is covered with books. If the speaker uttered ‘the table at which the author of *The Persistence of Objects* is sitting at t_1 is covered with books’ then the speaker said and intended to communicate that there is one and only one table at which the author of *The Persistence of Objects* is sitting at t_1 and it is covered with books. These propositions are clearly distinct, given that one can be true and the other false, as is the case in a world that contains two tables, one that is covered with books and is the only table at which the author of *The Persistence of Objects* is sitting at t_1 , and another that is not covered with books and is the only table in room 209 of Camden Hall at t_1 . If the audience is unable to know which of these sentences was uttered, therefore, they are unable to know what was said and meant. The situation can be represented diagrammatically



According to this diagram there are two candidates for the sentence uttered, given the sentence articulated and the surrounding context, each of which determines a different candidate for what is said, in turn determining different candidates for what is meant. If the audience cannot know which path to take, that is, if they cannot know which sentence was uttered, then the audience cannot know what the speaker intended to communicate.

Wettstein (1981, 247) takes it as “quite clearly wrong to suppose that, in many such cases, the circumstances of utterance put the listener in a position to select some one of these non-equivalent descriptions as the correct one, the one that actually captures what the speaker intended.” This assertion can be supported, as we did in discussing Schiffer’s Pergola case, by noting that there needn’t be any significant difference between these candidates in terms of complexity, availability, contextual salience, or informativeness. Wettstein’s concerns (1981, 247) go beyond this epistemic worry, however. If utterances of definite descriptions are elliptical, then the audience cannot know what is uttered, said, or meant, and communication cannot succeed, but moreover “the indefinite definite descriptions we actually utter are not elliptical for the uniquely denoting descriptions that Russell’s theory requires” because the speaker’s communicative intention is not appropriately captured by any one of the candidates for what is uttered, to the exclusion of the others. Asking the speaker which of the candidates they intended to utter will leave them at a loss and unable to decide, say, that they intended to refer to the table as ‘the table in room 209 of Camden Hall at t_1 ’ *as opposed to* ‘the table at which the author of *The Persistence of Objects* is sitting at t_1 ’ and it is “implausible in the extreme to suppose that in fact one of these descriptions captures what the speaker intended but that we cannot, even with the help of the speaker himself, come to know which description that is.” Even on the assumption that something is said, the audience cannot have intended to communicate it. “How then” asks Wettstein (1981, 248) are we to account for the fact that in such cases, a non-defective speech act is performed, a fully determinate assertion is made?”

There is a possible answer to this question that Wettstein doesn’t consider. Suppose that the speaker implicates what they intend to communicate. If the speaker intends to implicate some proposition of the form specified by the Russellian account, then the problem simply reoccurs. The speaker intended to communicate what they implicated but there are a number of different candidates for what is implicated that are equally appropriate in the context: the proposition that the only table at which the author of *The Persistence of Objects* is sitting at t_1 is covered with books and the proposition that the only table in room 209 of Camden Hall at t_1 is covered with books.

Blackburn against Wettstein's referential account

Wettstein's alternative view of definite description begins by endorsing Donnellan's (1966) distinction between referential and attributive uses of definite descriptions. Key to this distinction is the claim that an utterance such as 'The murderer of Smith is insane' has two distinct interpretations. On an attributive reading, this sentence has the truth conditions that Russell's analysis suggests; it is true if and only if there is one and only one murderer of Smith and they are insane.³⁸ On the referential reading, however, 'the murderer of Smith' refers directly to some individual who is presupposed to have murdered Smith. If this individual is Jones, then the speaker says that Jones is insane. The difference between these readings comes out when we consider the situation in which, unbeknownst to us, Smith was murdered not by Jones, but by Atkinson. Atkinson is insane, while Jones is not. On the attributive reading, the utterance is true as there is in fact one and only one murderer of Smith (that is, Atkinson) and they are insane. On the referential reading, however, the utterance is false in the same circumstances. Atkinson, the unique murderer of Smith, might be insane but Jones is not.

Suppose that a detective enters a crime scene to see Smith's dead body. The detective exclaims 'The murderer is insane!' Applying Russell's analysis directly to the sentence articulated, what is said is false as there is no unique murderer in the world. If, on the other hand, we interpret according to the syntactic ellipsis analysis, we face the problem of underdetermination as there may be many properties that we presuppose the murderer to instantiate, such as the property of having murdered Smith and the property of having murdered the guy on Smith's couch (who we take to be Smith). In such a context there need be no justification for taking the speaker to have uttered 'The murderer of Smith is insane' rather than 'The murderer of the guy on Smith's couch is insane', yet what is said varies, according to Russell's analysis of definite descriptions, depending on which is uttered. Although the context cannot determine a particular sentence that captures what the speaker means to communicate, Wettstein (1981, 248-249) points out that it can plausibly determine a referent for 'the murderer'. If the detective uses the description referentially to refer to a particular murderer, their utterance can be interpreted as though they used a proper name that

³⁸ This is not to say that Russell's theory gives the correct analysis of *all* attributively used definite descriptions.

referred to the murderer in question.³⁹ Supposing that the murderer is Jones, the detective says that Jones is insane.

It needn't be the case, however, that we can read the detective's use of 'the murderer' referentially. Supposing that the detective has no idea as to who murdered Smith and so cannot be using the definite description referentially,⁴⁰ the most natural interpretation of the detective's utterance is the attributive reading according to which the detective has said something true whoever turns out to have murdered Smith, so long as that murderer is insane. Applying Russell's analysis directly on 'The murderer is insane', the detective says that there is one and only one murderer and that murderer is insane, which is clearly false in a world with more than one murderer. As we have already seen, the syntactic ellipsis theory according to which the words articulated must be supplemented by additional descriptive material leads to underdetermination of what is uttered and said.

Wettstein's (1981, 251-252) proposed solution is that the words articulated are indeed elliptical but require completion, not with further descriptive material, but with referential material.⁴¹ According to this line, there may still be a number of candidates for the sentence uttered, such as 'The murderer of *Smith* is insane' and 'The murderer of *the person on the couch* is insane', where the italicised terms are read referentially.⁴² Context may fail to determine unique

³⁹ Wettstein assumes that proper names and referentially used definite descriptions are not subject to the problem of underdetermination, but see Chapter 3 for reasons to doubt this. The *Problem of the Many*, structurally similar to the problem of underdetermination, fuels the problem of underdetermination by showing that there are many candidates for what is said by an utterance that includes a proper name or a referentially used definite description.


⁴⁰ There is the logical possibility of taking the detective as using 'the murderer' referentially to refer to whoever in fact murdered Smith, despite the fact that the detective has no idea who this is. In that case, the detective says something true no matter who murdered Smith, not because they say the same thing whoever murdered Smith and that thing can be true no matter who murdered Smith, but because they say something different depending on who murdered Smith. If it turns out that Jones murdered Smith, then the detective said that Jones is insane; if it turns out that Atkinson murdered Smith, then the detective said that Atkinson is insane. This response is of no help in explaining the possibility of successful communication, however. If the speaker refers to whomever murdered Smith, then the audience cannot know what was said unless they know who murdered Smith. If the audience doesn't know who murdered Smith, they cannot know what was said and consequently cannot know what the speaker meant. As Wettstein (1981, 251) puts it "One fully understands the proposition without having any idea who murdered Smith ... Indeed, if understanding such a proposition required knowledge of which item was the murderer, the speaker himself would, at least in many cases, not understand his own utterance, for in attributive cases the speaker often has no belief about who the murderer is."

⁴¹ Wettstein isn't entirely explicit about offering a syntactic ellipsis analysis but it is a very natural way of cashing out Wettstein's talk of "implicit reference". Whether or not this is a variation of the syntactic ellipsis theory, Blackburn (1988), whose arguments we will come to shortly, shows that it is subject to the problem of underdetermination.

⁴² The comparison of three natural language sentences in Wettstein (1981, 252) suggests that he takes the elided material to belong to natural language. An alternative to these natural language sentences is to suppose that the

completing material for the definite description but it can at least determine the referent of the completing material. In this case, we have two distinct candidates for the sentence uttered, but as the only function of the completing material is to pick out a referent and the referent is the same whichever sentence is uttered, the same thing would be said either way: that there is one and only one murderer *of Smith* and that murderer is insane. What is said by the use of a sentence is a function of the sentence's structure and the referents of its terms in the context. As the referent of 'Smith' is identical to the referent of 'the person on the couch', the same thing would be said whichever sentence was uttered.

The situation can again be represented diagrammatically.



As before, there are at least two candidates for the sentence uttered, as there are at least two referential terms that would do for the elided material, given the sentence articulated and the surrounding context. Unlike the previous case, however, the same thing would be said either way, given that context determines the same reference whichever the elided term. Despite not knowing what is uttered (indeed, despite nothing being uttered, according to Wettstein's test of quizzing the speaker) the audience can know what is said: it is what is said on either of the plausible interpretations. Although there are two interpretive paths open to the audience, either will lead them to what the speaker intended to communicate.

As Blackburn (1988, 276) notes, however, Wettstein's view is also subject to the problem of underdetermination. In the context, there may be a number of

elided material is simply a variable, so that the sentence uttered by an attributive utterance of 'The murderer is insane' is always 'The murderer of x is insane'. Whichever way we go, the view still faces the problem of underdetermination, as we shall see in discussing Clapp's objection to Stanley and Gendler Szabó's hidden-indexical analysis of quantifier domain restriction.

equally viable candidates for the elided referential term, each of which refers to a different object and so determines a different candidate for the proposition said. Suppose the detective says ‘The murderer is insane’ given a context in which everyone knows that the murderer came in through the bathroom window and murdered Smith. Wettstein says that the articulation is elliptical for a longer sentence containing a referential term, but both the sentence ‘The murderer of *Smith* is insane’ and the sentence ‘The murderer who came in through *the bathroom window* is insane’, where the italicised terms are used referentially, are equally viable in the context. What is said varies depending on which of these candidates is uttered. If it turns out, for example, that there were two different murderers – one who came in through the window and murdered Stephenson and another who came in through the door and murdered Smith – it may be true that the murderer of Smith is insane and false that the murderer who came in through the window is insane.

We might think that the special connection between murderers and their victims forces us to take the elided material to refer to their victim but of course murderers can murder more than one person. Given a context in which we know that the same person murdered Smith and Stephenson, there are at least two candidates for the sentence uttered – ‘The murderer of Smith is insane’ and ‘The murderer of Stephenson is insane’ – each of which determines a different candidate for the proposition said, given that the description as a whole is being used attributively and only the embedded proper name is used referentially. The problem is only magnified when we consider other terms. For ‘the table’, there is no equivalent of a victim that we should refer to preferentially. If we know that the table on a particular rug is identical to a table below a particular chandelier, then we have at least two candidates for the sentence uttered by an attributive utterance of ‘The table is covered with books’: ‘The table beneath *that chandelier* is covered with books’ and ‘The table on *that rug* is covered with books’. Each of these candidates for the sentence determines a different candidate for the proposition said. The audience cannot know what the speaker said and, by the Epistemic Content Constraint, cannot know what the speaker intended to communicate.

Stanley and Gendler Szabó against a syntactic ellipsis account of quantifier domain restriction

Stanley and Gendler Szabó (2000, 237-239) extend Wettstein's concerns to argue against syntactic ellipsis analyses of quantifier domain restriction generally. The problem of quantifier domain restriction is the problem of explaining how it is that speakers are able to communicate different propositions by articulation of the same quantified sentence in different contexts. According to Russell's analysis, 'The table is covered with books' is an existentially quantified sentence and it is clear that an utterance of this sentence can communicate, in one context, that the table in Room 209 of Camden Hall is covered with books and, in another context, that the table at which the author of *The Persistence of Objects* is sitting is covered with books. Stanley and Gendler Szabó illustrate the problem through universal, rather than existential, quantification.

Suppose that Lisa went to the shop to buy bottles for Max's homemade beer. Lisa utters "Every bottle is empty" thereby communicating that every bottle she just bought is empty. How is this achieved? According to a syntactic ellipsis analysis, Lisa should be interpreted as having uttered a sentence in which the domain of quantification is explicitly restricted. Given the proposition that Lisa communicates, an obvious candidate for the sentence uttered is 'Every bottle I just bought is empty' where 'I just bought' is an unpronounced predicate that restricts the domain of quantification. As Stanley and Gendler Szabó (2000, 237-238) point out, however, "There are very few cases where there is a single plausible candidate for the role of the domain restricting predicate." An equally viable candidate is the expression 'I purchased recently'. As we noted in the discussion of Wettstein's referential account, underdetermination of what is uttered might not seem particularly problematic when all candidates for what is uttered are synonymous, in which case the audience may be able to know what is said, even if they cannot know what is uttered. However, the candidates need not even be restricted to synonymous sentences. The completing material might just as well be 'you are looking for', which determines an alternative candidate for what is said.

Stanley and Gendler Szabó (200, 238) consider that there is a single plausible principle that breaks the stalemate between various candidates for the sentence uttered: take the sentence uttered to include a demonstrative predicate, as in 'Every bottle that is one of those is

empty'. Against this suggestion, they point out that this strategy is not available to an audience whose limited understanding of English does not include the word 'those'. Such an audience is not capable of knowing what is uttered or said when the sentence in question includes the word 'those', yet, they claim, such an audience is still capable of knowing what is meant.

If the syntactic ellipsis view of quantifier domain restriction is correct, audiences are regularly unable to know what is uttered or said by utterances of quantified sentences. According to the Epistemic Content Constraint, therefore, the audience cannot know what is meant. Stanley and Gendler Szabó (2000, 239) agree, saying that, if we assume the syntactic ellipsis view, Max "cannot know what sentence was uttered and consequently cannot know in the normal way what proposition was meant."⁴³ As Max has no way to identify the sentence uttered, the proposition said, or the proposition meant, the pair conclude against the syntactic ellipsis analysis of quantifier domain restriction.

There may of course be cases in which the context weighs decisively in favour of one of these candidates (as when Lisa is responding to the question 'Do the bottles you just bought need to be emptied?') but the argument requires only that there are cases in which an utterance is interpretable despite underdetermination of the sentence uttered. To get underdetermination, we need only assume that the speaker and their audience share a very modest stock of information about the bottles. If this prevents the audience from knowing what is the speaker intended to communicate, then audiences are unable to know what speakers intend to communicate in the vast majority of cases.

Clapp against traditional semantics

Stanley and Gendler Szabó present the problem of underdetermination in a bid to support their own preferred account of quantifier domain restriction. It is a significant, therefore, that the problem of underdetermination re-emerges on their preferred account. Indeed, Clapp (2002) argues that the problem will emerge for any account that assumes the "traditional

⁴³ Stanley and Gendler Szabó don't clearly restrict this claim to cases in which the speaker says what they intend to communicate, saying that "in normal instances of *successful* communication, the hearer who grasps the proposition communicated will also know what sentence was uttered and what proposition was [said] by that sentence on the given occasion" without telling us quite what 'normal instances' amount to.

semantic” paradigm endorsed by Stanley and Gendler Szabó. The paradigm is characterised by the principle of Truth Conditional Compositionality:

“The truth conditions of an utterance are a function of (i) the logical form of the utterance (i.e., the structure of the LF of the utterance), and (ii) the meanings of the words in the utterance (i.e., the semantic values of the terminal nodes of the LF of the utterance).” (232)

Although not part of the Encoding Model, which is consistent with many different accounts of the relation between the meanings of words used and what is said by their utterance, allowing for implicature for example, we have regularly assumed Truth-Conditional Compositionality in our examples. Cases of quantifier domain restriction and nonsentential assertion are *prima facie* counterexamples to this principle, as they appear such that a particular sentence uttered (or, equivalently, a particular logical form) and assignment of referents (or, equivalently, semantic values) determine different truth conditions in different contexts. Clapp considers two ways to respond to the *prima facie* problem: either deny that the logical form is stable by positing different elided material in each context, or deny that the values of the terminal nodes are stable by positing a hidden indexical that takes a different value in each context. Clapp argues, however, that both strategies fall to the problem of underdetermination.⁴⁴

Stanley and Gendler Szabó take underdetermination worries to refute the syntactic ellipsis account of quantifier domain restriction but fail to recognise that their own views face exactly parallel problems. As Clapp (2002, 245-248) notes, the very same underdetermination problem arises for Stanley’s (2000, 404-407) syntactic ellipsis analysis of nonsentential assertion, according to which an articulation of ‘A world famous topologist’, said at a party in response to a quizzical look from a friend who has nodded towards a particular female attracting a lot of attention, is elliptical for ‘She is a world famous topologist’. The phrase articulated is supplemented with additional material to yield a complete sentence uttered. Clapp points out that in any realistic context we will share lots of information about the intended referent, leading to a number of equally viable candidates for the elided material. In addition to

⁴⁴ These two strategies are not exhaustive. Alternatively, a term evident in what is articulated (and so not ‘hidden’) may be context-sensitive, as if ‘every’ in ‘Every bottle is empty’ is an indexical. Given the scope and simplicity of the problem of underdetermination, however, it is very plausible to suppose that it could be levelled against any plausible strategy for defending Truth Conditional Compositionality against these *prima facie* counterexamples.

‘She is’, the elided material could be ‘That woman is’, ‘That person is’, or ‘The person attracting a lot of attention is’. Each of these candidates for the elided material determines a different sentence uttered and proposition said. As the audience cannot know what was said, they cannot know what was meant. As the problem of underdetermination arose from the simple assumption that the speaker and their audience share a modest degree of knowledge about the intended referent, the conclusion generalises to the vast majority of nonsentential assertions, in contradiction of the obvious fact that communication can regularly succeed in cases like this. Clapp concludes that syntactic ellipsis accounts cannot defend Truth Conditional Compositionality against this *prima facie* counterexample.

Having also rejected syntactic ellipsis as an account of quantifier domain restriction, Stanley and Gendler Szabó (2000, 251-253) endorse a hidden-indexical analysis, but this view is also subject to the problem of underdetermination. According to this view, a speaker who articulates ‘Every bottle is empty’ utters the sentence ‘Every bottle *i* is empty’ where ‘*i*’ is an unpronounced hidden indexical. For ease of explication, the pair often talk as though the hidden indexical denotes, relative to a context *C*, the set of objects that constitute the domain of the quantifier in *C*,⁴⁵ but, more precisely, they take the value of the hidden indexical, relative to a context *C*, to be a function from a world (or a world and time) to the set of objects that constitutes the domain of the quantifier at that world (or world and time). This change is required to secure the truth of ‘If there were a few more bottles on the shelf, John would not have purchased every bottle’, said about John who has the habit of buying exactly seventy bottles whenever he goes to the supermarket. There were precisely seventy bottles the last time that John went to the supermarket, so John in fact purchased every bottle. Had there had been any more bottles on the shelf, however, John would not have purchased every bottle, although he might have purchased each of the bottles he purchased in the actual world.

By stipulating that a single sentence (‘Every bottle *i* is empty’) is uttered whenever the sentence ‘Every bottle is empty’ is articulated, Stanley and Gendler Szabó avoid underdetermination of what is uttered but Clapp (2002, 249) notes that the pair succeed “only in relocating the fundamental problem” as underdetermination of what is said. Successful

⁴⁵ Actually, the pair represent the unarticulated constituent as ‘*f(i)*’, where ‘*f()*’ is an indexical that denotes, relative to a context, a function from individuals to sets and ‘*i*’ is a further indexical that denotes, relative to a context, an individual. By applying the function determined by a context *C* to the individual determined by *C*, we derive the domain restriction in *C*. This complication can be safely set to one side here.

linguistic interpretation requires identification of what is said as well as, and perhaps more importantly than, identification of what is uttered. Stanley and Gendler Szabó's analysis of quantifier domain restriction renders the latter fairly easy but at the cost of making the former impossible.⁴⁶

Consider again the case of Lisa and Max. When Max asks about the bottles, Lisa replies 'Every bottle is empty' thereby communicating the proposition that she would have said by uttering 'Every bottle I just bought is empty'. According to Stanley and Gendler Szabó's analysis, the sentence that Lisa utters contains a hidden indexical 'i' which denotes a function from worlds to sets of objects but if we assume that Lisa and Max share even a modest stock of information about the bottles, then there will be various candidates for the value of the indexical and so various candidates for the proposition said.

Suppose it is mutually known that Lisa put all the bottles she just bought in an empty cupboard, so every bottle that Lisa just bought is in the cupboard and every bottle in the cupboard is one that Lisa just bought. The value of 'i' might then be a function from any world W to the set of bottles that Lisa bought at W , or a function from any world W to the set of bottles in the cupboard at W . As there are worlds in which Lisa did not put the bottles in the cupboard, but put them on the shelf instead, these different candidates for the denotation of the indexical determine different candidates for the proposition said. Without any reason to privilege one of these candidates over the other, the audience cannot know what is said. If the speaker intends to communicate what they have said, then the audience cannot know what the speaker intended to communicate. Again, if such a limited stock of shared knowledge makes it impossible for the audience to know what the speaker meant, then it will be impossible in most use of quantified language. Clapp concludes that hidden-indexicality cannot save Truth Conditional Compositionality from *prima facie* counterexamples.

We would not get this problematic result if the denotation of the hidden indexical, in a context, were just a set of objects, in which case we could define the proposition said by 'Every bottle i is empty' in a context C as true if and only if the intersection of the extension of

⁴⁶ Clapp's objection focuses on Ludlow's (1989) hidden indexical analysis of 'Bradley is tall' but Clapp (2002, 245) intends his examination of relatively few cases "to illustrate that any analysis utilizing either the pragmatic ellipsis substrategy or the hidden indexical substrategy will fall" to the problem of underdetermination. Stanley and Gendler Szabó's analysis of quantified expressions is mentioned earlier in his paper.

‘bottle’ with the extension of ‘i’ (in C) is a subset of the extension of ‘empty’. Whether the value of ‘i’ (in C) is the set of things that Lisa bought (bottles and pickles) or the set of things in the cupboard (bottles and crayons), the proposition said is true if and only if the bottles are empty.

As Stanley and Gendler Szabó have argued, however, the value of the indexical relative to a context must be a function from worlds to sets of objects, in which case we can define what is said by ‘Every bottle *i* is empty’ in a context C as true at any possibility W such that the intersection of the extension of ‘bottle’ at W with the extension of ‘i’ (in C) at W is a subset of the extension of ‘empty’ at W. If the value of ‘i’ (in C) is a function from any world W to the set of things that Lisa bought in W, what is said (call it p_1) is true at any world W such that every bottle that Lisa bought in W is empty in W. If, on the other hand, the value of ‘i’ (in C) is a function from any world W to the set of things in the cupboard in W, what is said (call it p_2) is true at any world W such that every bottle in the cupboard in W is empty in W. These truth conditions are distinct. There are worlds that satisfy the truth conditions of p_1 but not p_2 , for example the world in which Lisa, after emptying the bottles she bought, put them on the shelf because the cupboard was already stocked with full bottles.

Although Stanley and Gendler Szabó’s account of quantifier domain restriction avoids underdetermination of what is uttered by stipulating that the sentence Lisa articulates determines a particular sentence uttered, it faces the corresponding, and more central, problem of underdetermination of what is said. Assuming a syntactic ellipsis account of quantifier domain restriction, there are various distinct candidates for the domain-restricting predicate. Assuming a hidden-indexical account, there are various distinct candidates for the value of the domain-restricting indexical. Either way, the audience cannot know what is said and, by the Epistemic Content Constraint, cannot know what the speaker intended to communicate.

As he takes syntactic ellipsis and hidden-indexical views to be refuted by the problem of underdetermination, Clapp (2002, 252) concludes that there is no plausible defence of Truth Conditional Compositionality against *prima facie* counterexamples like quantifier domain restriction and nonsentential assertion that will not fall to the problem of underdetermination. Traditional semantics should therefore be rejected.

Underdetermination without traditional semantics

Clapp (2002, 231) takes the problem of underdetermination to stem from the assumption of traditional semantics and suggests replacing it with “some form of *dynamic semantics*, or *truth conditional pragmatics*” but like Wettstein and Stanley and Gendler Szabó before him, Clapp fails to note that underdetermination problems arise even for his favoured class of theories. Clapp doesn’t define dynamic semantic or truth conditional pragmatics but does offer Bach (1994), Carston (1991), Sperber and Wilson (1986), Travis (1985), Recanati (1996), and Kamp and Reyle (1993) as exemplars. Given that the problem of underdetermination can arise for each of these theorists, Clapp’s rejection of Truth-Conditional Compositionality provides no indication as to how the problem of underdetermination is to be avoided.

Completion and Expansion

Bach (1994), Carston (1988), and Sperber and Wilson (1986) all deny Truth Conditional Compositionality by taking the truth conditions of some utterances to be a function not only of the logical form of an utterance and the meanings of its component expressions, but also of contextual features that drive the processes of *completion* and *expansion* by which truth conditions can be derived. Completion is required when the sentence uttered fails to determine a complete proposition. According to Bach (1994, 127) for example, the structure of ‘Steel is not strong enough’ and the meanings of its component expressions fail to determine a complete proposition. Rather, the sentence determines a propositional radical that requires completion in the form of “some contextually identifiable respect” in which steel is not strong enough (such as *for building a 500-storey building* or *to resist bending by Superman*) to determine a proposition. Expansion occurs when the sentence uttered determines a complete proposition, but some closely related proposition is communicated instead, as when an articulation of ‘I have eaten breakfast’ is taken to communicate what would have been said by an utterance of ‘I have eaten breakfast *today*’.⁴⁷ The result of completion or expansion is an *implicature*.

The truth conditions of an utterance are not determined in accordance with Truth Conditional Compositionality but if anything, this only makes the problem of

⁴⁷ Note that implicatures needn’t always be communicated. If I am speaker sarcastically by saying ‘I have eaten breakfast’, the implicature may be that I have eaten breakfast today, even though I communicate that I haven’t eaten breakfast today.

underdetermination more pressing, allowing for another respect in which the speaker's meaning depends on the context and so another juncture at which the context can be compatible with a number of different candidates for what the speaker intended to communicate.

A terminological issue. Do completion and expansion contribute to what is said by an utterance, or is what is implicated distinct from what is said? For our purposes, it makes no difference, as the problem of underdetermination can be formulated either way. Supposing that completion and expansion contribute to what is said, the Epistemic Content Constraint can hold even in cases of implicature because the speaker intends to communicate what they say. Supposing, on the other hand, that what is implicated is distinct from what is said, the speaker doesn't intend to communicate what they say in cases of completion and expansion. Although the Epistemic Content Constraint does not apply in such cases, we can provide a very similar principle that plays the same role, namely

The Epistemic Implicature Constraint: If the audience cannot know what the speaker implicated, then the audience cannot know what the speaker intended to communicate.

For simplicity, we can assume that completion and expansion contribute to what is said and so avoid this additional, though innocuous, principle.

For Bach's example of 'Steel is not strong enough', we can easily generate cases of underdetermination by assuming even a very limited stock of knowledge shared between speaker and audience. If we all know, for example, that our only assignment is to build a 500-story building, both the completion *for building a 500-story building* and *to complete our only assignment* determine equally plausible, though distinct, truth conditions for the speaker's utterance. If we assume that the only challenge involved in building a 500-story building is making it strong enough to resist bending by Superman (who has taken to bending all 500-story buildings), then either of Bach's suggested completions would determine an equally viable candidate for what is said.

Returning to the issue of quantifier domain restriction, Bach (1994, 140) tell us that 'Everybody is coming to my party' is to be interpreted by expansion. For example, it might be expanded so that the truth conditions of this utterance are given by the proposition that would

be said by an unexpanded utterance of ‘Everybody *in my class* is coming to my party’. Given a context, however, in which we know that I like all and only the people in my class, it might equally well be expanded so that the truth conditions of the utterance are given by the proposition that would be said by an utterance of ‘Everybody *I like* is coming to my party’. In such a case, the audience cannot know what was said and, consequently, cannot know the speaker intended to communicate.

Carston (1991, 168), and Sperber and Wilson (1986, 158) offer a way to distinguish between these two alternative candidates for what is said.⁴⁸ They take interpretation to be guided by the Principle of Relevance: “Every act of ostensive communication communicates a presumption of its own optimal relevance”, that is, such acts communicate the presupposition that what the speaker intends to communicate will generate sufficiently great contextual effects to make it worth the addressee’s while to process it and that the speaker’s articulation was the least costly way, in terms of processing effort, the speaker could have chosen to communicate what they intend to communicate.⁴⁹ An interpretation has a greater contextual effect to the extent that it provides evidence for or against pre-existing assumptions, thereby affecting the audience’s confidence in those assumptions, or it interacts with pre-existing assumptions to yield further assumptions.

During interpretation, the audience must try to find an interpretation that is consistent with the Principle of Relevance, where this means finding an interpretation that “a rational communicator might have expected to be optimally relevant to the addressee”.⁵⁰ What the speaker intends to communicate may not in fact be optimally relevant. If the audience recognises that the speaker is a bore, they may recognise that what they intend to communicate is not as relevant as the speaker thinks. In the best cases, however, the presupposition is accurate and these are the cases we will focus on.

⁴⁸ Neither work endorses Bach’s particular account of ‘Steel is not strong enough’ but offer parallel accounts of other sentences. We will stick with Bach’s example for simplicity.

⁴⁹ Carston (1991, 168) says “hearer’s”, rather than “addressee’s”. I prefer “addressee’s”, given that the hearer might be an eavesdropper, rather than the intended addressee. If the articulation was designed to be optimally relevant when interpreted by one individual, there is no guarantee that it will be optimally relevant when interpreted by someone else, whose prior presuppositions may be entirely different from the intended addressee’s, leading to entirely different contextual effects.

⁵⁰ Sperber and Wilson (1986, 166), quoted in Carston (1991, 169).

We said previously that there may be two equally plausible completions of ‘Steel is not strong enough’ in a given context – *for building a 500-story building* and *to complete our only assignment*, for example – leading to two equally plausible candidates for what is said. Sperber and Wilson (1986, 168), however, argue (and Carston (1991, 169) agrees) that at most one interpretation will be consistent with the Principle of Relevance. Even supposing that both of these distinct candidates for what is said have contextual effects sufficient to make each interpretation worth the speaker’s while, only one can be the interpretation that requires the least processing effort: whichever interpretation occurs first. As the interpretation that occurs to the interpreter first requires the least processing effort, this is the interpretation that is optimally relevant and therefore the interpretation that the interpreter should assume, according to Carston and Sperber and Wilson, captures what the speaker intended to communicate.

As already noted, however, the audience should only settle on an interpretation that a rational communicator might have expected to be optimally relevant to the addressee. As the first interpretation is optimally relevant only because it occurred to the interpreter first, a rational communicator could only have expected it to be optimally relevant if they could have predicted that it would occur to the interpreter first. In this case, however, the speaker could not have had predicted which interpretation would occur first, as both are equally appropriate, and so could not have expected either one to be optimally relevant.

The audience cannot therefore assume that their first interpretation is the one the speaker intended and we are still left with two equally viable candidates for what is said. In fact, it is not at all obvious that either interpretation of what is said need occur to the interpreter first. Both interpretations may occur simultaneously. In such a case, Sperber and Wilson (1986, 169) say that “the addressee will be unable to decide what the informative intention was, and communication will fail.” Whether one of the candidates occurs first, or both occur simultaneously, the audience cannot know what the speaker intended to communicate, and communication will fail.⁵¹

⁵¹ What about the conjunctive account, according to which the most relevant interpretation is the conjunction *Steel is not strong enough for building a 500-story building and to complete our only assignment*? This is certainly less relevant than either of the conjuncts. To derive the conjunction, the audience must derive each of the conjuncts. As the conjunction has no contextual effects beyond either of the conjuncts but necessarily requires greater processing

Austinian semantics

According to the Austinian semantics discussed in Recanati (1996), the truth conditions of an utterance (what is said) are determined by two things. First a set of situations; second a contextually-determined individual situation. As the individual situation is not determined by the structure of the sentence uttered or the referents of its terms, this view denies Truth Conditional Compositionality. An utterance is true if and only if the individual situation is a member of the set. This allows for an analysis of ‘Every bottle is empty’ according to which the meaning of the sentence is the same whenever it is uttered – it refers to a set of situations – while allowing that the truth conditions of utterances of this sentence vary because of variation in the individual situation referred to.

Although this view denies Truth Conditional Compositionality, it still allows for underdetermination of what is said because different situations may be equally salient when the speaker and interpreter share even a modest degree of common knowledge. Suppose we know that every bottle in my apartment is in the kitchen. I say ‘Every bottle is empty’. This sentence refers to a set of situations and my utterance is true if and only if the contextually-determined situation is a member of this set, but equally viable candidates for the individual situation include my apartment and my kitchen. There are therefore two distinct candidates for what is said by my utterance: that my apartment is a member of the set of situations referred to by ‘Every bottle is empty’ or that my kitchen is a member of that same set of situations.

One potential response is that, when faced with a number of situations that seem equally relevant, we should select the largest. Another response suggests that we should always select the smallest. Either response is entirely arbitrary, however. There is in principle nothing about larger or smaller situations that makes them more appropriate. Even if we know that every bottle in the Oval Office is empty, my apartment remains a better candidate for the individual situation referred to than the situation composed of my apartment plus the Oval office. Equally, even if every beer in my kitchen is in the half of my kitchen closest to the floor, my kitchen remains a better candidate for the individual situation referred to than the lower half of my kitchen.

effort than either conjunct, no rational speaker could have expected the conjunctive interpretation to be optimally relevant to the addressee.

The S-View

Travis (1985, 187) defends the S-View, according to which “typically, an (e.g.) English expression is such that, with its meaning (unambiguously) fixed, there are a variety of distinct (perhaps better: distinguishable) things to be said in using it on some production of it or other”. The view is supported by a limited number of cases that are intended to generalise very broadly. For our purposes, we will focus on one example that reoccurs throughout Travis (1985), that is, the sentence ‘The kettle is black’, which, according to Travis, may be true of a particular kettle at some time, when uttered in one context, and false of the very same kettle at the very same time when uttered in another context. In the former context, what is said may be that the kettle is black on the inside, while in the second context, what is said may be that the kettle is black on the outside, one of which may be true and the other false of the very same kettle at the very same time.

Travis’s discussion is deeply interesting because it pre-empts a great many positions that would later be distinguished more precisely. For the same reason, however, it is not entirely clear that a unified view is being discussed throughout. Here, we can note some of the more prominent candidates for Travis’s S-View, all of which remain subject to the problem of underdetermination.

One position can be raised and set aside in short order. Travis (1985, 190) says that “an incomplete thought is the most that is determined by what ... words mean in English” which is clearly reminiscent of the position of Bach (1994) that the meanings of some complete sentences are associated only with propositional radicals that require completion for the sentence to express a complete thought or proposition. Perhaps the meaning of ‘The kettle is black’ determines only a radical that delivers a proposition only when supplemented with some particular way of being black, e.g. being black on the inside or being black on the outside. As noted above, such a view denies Truth Conditional Compositionality but remains subject to underdetermination of what is said.

At points, however, Travis suggests (whether intentionally or not) a family of views compatible with Truth Conditional Compositionality, according to which the reference of some articulated expressions vary (whether due to context-sensitivity, or something else)

radically from context to context.⁵² As Travis (1985, 217) says “in the case of ‘is black’ one element in fixing a complete thought, for any given speaking, might be a proper specification of what being black is to be understood to come to, for the purposes of what is being said on that speaking” just as “for any speaking of the name ‘George’ a specification may be called for as to which (purported) George is to be understood to have been named.” In the case of ‘George’, the difficulty is figuring out the individual being referred to by ‘George’. The corresponding difficulty in the case of ‘is black’ is figuring out what property is being referred to. By drawing an analogy between predicates like ‘is black’ and those paradigms of articulations that vary in reference from context to context, the proper names, Travis suggests that we must appeal to context in determining the referent of an occurrence of ‘is black’, which is entirely consistent with Truth-Conditional Compositionality.

The view that the reference of an articulated term varies from context to context, however, is consistent with a number of more precise views: that the articulated term is context-sensitive, for example, or that it is shorthand for a more complex linguistic expression, as in cases of syntactic ellipsis.⁵³ In addition, we can explain the variation as an ambiguity *of* language, as opposed to ambiguity *in* language. Even when we know that someone speaks English on two different occasions, there is a sense in which we might still not know whether they spoke the same language on each occasion. That is, identifying someone as speaking English, or even a much more specific dialect, radically underdetermines what language they are speaking. A language requires, minimally, a pairing of lexical items and their meanings, but there are many ways of pairing lexical items and meanings that are consistent with speaking any dialect. I find this view attractive, in particular, for the case of proper names. When I talk about George by saying ‘George is hungry’, I don’t take myself to be using a context-sensitive or ambiguous term that refers to George on this occasion but might refer to someone else

⁵² Travis (1985) repeatedly tells us to hold references fixed across contexts, but there is no suggestion that he thinks of predicates as having ‘references’ in the relevant sense. Even if the requirement to hold references fixed precludes a role for context in determining what is said on the basis of what is uttered, it is consistent with a role for context in determining what was uttered on the basis of what was articulated, as in a syntactic ellipsis view.

⁵³ Travis (1985, 203) does point out that ‘is black’, for example, is used to predicate a certain colour and that anything more than this – that the object is black on the inside or on the outside – “is not to be expected from the meanings of the *words*” but this is consistent with a context-sensitivity account, as the rest of the work is done by the context. Compare ‘I’. The meaning of this paradigm of context-sensitivity ensures only that it refers to speakers. More than this – the identity of the particular speaker it refers on a particular occasion – is not provided by the meaning of ‘I’ alone, but only in combination with the context.

(someone also called ‘George’) on another occasion, but as using a univocal, context-insensitive term that refers to George in every context. In the language I am speaking on that occasion, ‘George’ is a lexical item that refers to George. When I articulate the same sentence to refer to another George, I am speaking a subtly different language in which ‘George’ is a lexical item referring to a different individual. According to this line, what is said by using a term is not ambiguous or context-sensitive in the language spoken on each occasion. The term is univocal on each occasion because the language spoken varies from occasion to occasion.⁵⁴

A further possibility compatible with Truth Conditional Compositionality, suggested by Predelli, is also suggested by Travis (1985, 206, footnote 31): “One might detect, e.g., one thing which is said in all speaking of ‘The kettle is black.’, and which, furthermore, gets evaluated as true or false, but which may be true as said in some speaking, not true as said in others ... Perhaps for a given occasion, we could be told or shown what it would be reasonable to take being black to come to”. According to one interpretation of Predelli (2005, 142-149), the truth value of a sentence may shift, not because of a shift in the truth-conditions⁵⁵ of the sentence, or because the world has changed, but because our way of categorising the world has changed. Given a context in which the interior of the kettle is relevant to our interests, blackness on the interior of the kettle may be sufficient for us to categorise the kettle as *black*; given a context in which the exterior of the kettle is relevant to our interests, blackness on the exterior of the kettle may be sufficient for us to categorise the kettle as *black*. Given two contexts in which our interests, and so our categorisation of the world, are different, therefore, the truth value of a sentence may vary even though its truth conditions remain stable.

In short, by citing Travis (1985), Clapp provides no indication as to how abandoning Truth Conditional Compositionality in favour of a dynamic semantic or truth conditional pragmatic approach can avoid the paradox of underdetermination. On perhaps the most prominent interpretation of Travis’s S-View, it is a truth conditional pragmatic approach akin to that defended in Bach (1994) and both are subject to underdetermination for the very same reasons. On other interpretations, it is not even inconsistent with Truth Conditional Compositionality.

⁵⁴ See, for example, Lewis (1975).

⁵⁵ Or, in Predelli’s terminology, “t-distribution”.

Discourse Representation Theory

Kamp and Reyle (1993) provide a thorough account of Discourse Representation Theory (DRT). Distinctive of DRT is its representation of the meaning of an uttered sentence. While we have generally represented the meanings of linguistic expressions by associating them with referents, which compose to yield the meaning of complete sentences, in DRT expressions are primarily associated with rules for the construction of Discourse Representation Structures (DRSs) – mental representations of the state of a discourse. DRT is, however, entirely consistent with the Extended Model introduced in Chapter 1. Although an uttered sentence is associated with a DRS rather than truth conditions, in the first instance, the truth conditions of the uttered sentence can be defined in terms of that DRS. This truth condition defines what is said by the utterance of that sentence. On the assumption that the speaker means what they say, therefore, we have the Epistemic Content Constraint. Although we can define the truth conditions of an utterance in terms of its effect on the DRS, this effect, and so the truth conditions, are partially determined by features of any pre-existing DRS. As these features of the existing DRS needn't affect the referents of the terms in the sentence uttered, we have the denial of Truth Conditional Compositionality.

Despite denying Truth Conditional Compositionality, however, DRT can be confronted with the very same problem of underdetermination that confronted theories which endorse Truth Conditional Compositionality. Take Stanley and Gendler Szabó's case. Lisa utters 'Every bottle is empty'. According to Kamp and Reyle (1993, Chapter 2.2) the truth conditions of this sentence are exactly as we assumed: the sentence is true at a world W (or in a model) if and only if the set of bottles at W is a subset of the set of empty things in W . In other words – and, following more directly from Kamp and Reyle's translation of the DRS associated with 'Every bottle is empty'⁵⁶ as $\forall x(\text{bottle}(x) \rightarrow \text{empty}(x))$ – the sentence is true at W if and only if everything in W is either not a bottle, or is an empty bottle; if and only if, that is, every bottle is empty.

Lisa doesn't want to communicate this proposition. She is well aware that there are plenty of non-empty bottles around. Rather, Lisa wants to communicate that every bottle she

⁵⁶ The account of 'Every bottle is empty' is extracted from Kamp and Reyle's (1993, account of 'Every professor who knows German owns Buddenbrooks' as $\exists yz(y = \text{German} \ \& \ z = \text{Buddenbrooks} \ \& \ \forall x((\text{professor}(x) \ \& \ \text{knows}(x,y)) \rightarrow \text{owns}(x,z)))$ '.

just bought is empty, remaining neutral about the emptiness or otherwise of the rest of the bottles in existence. More than that, Lisa is successful in doing so. This is precisely the problem of quantifier domain restriction that Stanley and Gendler Szabó were concerned to solve. In responding to this problem, both the syntactic ellipsis and hidden-indexical analyses are open to the DRT theorist, and each will lead to the problem of underdetermination in the ways already described.

For an articulation of ‘Every bottle is empty’, context may restrict the domain of quantification by providing additional linguistic material or by providing an interpretation for a ‘hidden’ indexical in the sentence uttered. The syntactic ellipsis analysis faces precisely the problem it faced for traditional linguistic theories. Given a situation in which the very same bottles are known to constitute the bottles in the cupboard and the bottles that Lisa bought today, we have two equally good candidates for the sentence uttered, each of which determines a different truth condition and proposition said. As the audience cannot know the truth condition of the utterance, they cannot know what the speaker intends to communicate, given the Epistemic Content Constraint.

The problem for the hidden-indexical analysis is also fundamentally the same as the problem that arose in the context of traditional linguistics. According to this view, the sentence uttered by an articulation of ‘Every bottle is empty’ is in fact the sentence ‘Every bottle *i* is empty’, where ‘*i*’ is an indexical that should be interpreted differently depending on the context. In the context of traditional linguistics, we said that the reference of the indexical varies with the context. In this context a more theory-neutral hidden-indexical analysis says that the rule for the interpretation of the indexical varies with context. In one context, the indexical should be interpreted in similar fashion to the expression ‘in the cupboard’ and in another context the same expression should be interpreted in similar fashion to the expression ‘Lisa purchased today’. Given a context, however, in which either interpretation of the indexical is equally warranted, we face the problem of underdetermination. Depending on which way the utterance is interpreted, the truth conditions of the utterance will differ. If the audience cannot know which interpretation to take, then they cannot know what is said and, consequently, cannot know what is meant.

As we shall see, DRT has the resources to provide a solution to the problem of underdetermination. Not, however, because it denies Truth Conditional Compositionality, but

because it has the resources to model local equivalence between different candidates for what is said. In Chapter 4, this solution will be presented in terms of a Stalnakerian framework, which more conspicuously represents the role of presupposition and is consistent with Truth Conditional Compositionality.

The source of the problem

In this Chapter, we have seen several hypotheses as to the source of underdetermination. Wettstein (1981) argued that problem arose from a failure to allow for the use of directly-referential expressions in a syntactic ellipsis analysis. As Blackburn (1988) made clear, however, the problem of underdetermination emerges once again on Wettstein's account, showing that he was mistaken as to the ultimate source of the problem. Stanley and Gendler Szabó laid the blame with the syntactic ellipsis analyses more generally, yet the problem arises in essentially the same form for their alternative hidden-indexical analysis. Noticing this, and that applications of the hidden-indexical and syntactic ellipsis analyses regularly assume Truth Conditional Compositionality, Clapp lays the blame with this more foundational principle, yet, as we have seen, the problem of underdetermination arises even for theories which reject Truth Conditional Compositionality, showing that it cannot be the source of the problem.

We have surveyed only a small sample of the literature that poses the problem of underdetermination and have restricted discussion to a limited number of linguistic constructions and hypotheses. In particular, we have focused on two types of linguistic construction, namely definite descriptions such as 'The guy's drunk' and universally quantified sentences such as 'Every bottle is empty', and on two types of linguistic hypothesis, namely hidden-indexical and syntactic ellipsis accounts. It is worth stating again that the aim of this Thesis is not to endorse any particular linguistic hypothesis as the correct account of the constructions at issue. Rather, the aim is to motivate an alternative to the Encoding Model according to which underdetermination of what is said is not necessarily problematic.

As such, my interest here lies less in the actual extent of underdetermination in English and more in the extent of underdetermination in theory, that is, less in the extent of

underdetermination according to the correct linguistic account of English⁵⁷ and more in the extent of underdetermination in proposed accounts, whether ultimately correct or not. So far, we have shown only how the problem arises for definite descriptions and universally quantified expressions. In the next Chapter, we will see that the Problem of the Many can be leveraged in extending the problem of underdetermination to referential terms such as proper names and indexicals.

In Chapter 4 it will be argued that the phenomenon of underdetermination is not in itself problematic, but appears so due to the mistaken assumption of the Encoding Model, which encouraged us to view communication on the model of passing a package of content from speaker to audience in the form of a sentence. On this model, a failure to know what was in the package can be expected to derail communication, hence the Epistemic Content Constraint. Rather, we should view communication less on the pass-the-parcel model and instead recognise that the role of the speaker is to guide their audience down an interpretative path that leads to what the speaker intends to communicate. On this model, there need be no proposition that is said, only a number of candidates for what is said, each of which leads to the same propositions being communicated.

We have already hinted at this alternative model in this Chapter, when we noted that the underdetermination of what is uttered needn't pose a problem for communication if the same thing is said by each of the candidates for what is uttered. What may not be obvious is how this model can be extended to the underdetermination of what is said. If the speaker speaks literally, without introducing any implicature, then surely the speaker means that *p* if they say that *p*. Indeed. It is consistent, however, for the saying of a proposition to communicate more than what is said. If we have two candidates for what is said, *p* and *q*, it is true that *p* is meant if *p* is said, while *q* is meant if *q* is said. Yet, in some contexts, it may also be true that *q* is meant if *p* is said and *p* is meant if *q* is said. Given that each of the candidates, *p* and *q*, is communicated whichever is said, the audience can therefore know what the speaker intends to communicate, even if they cannot know whether *p* or *q* is said.

⁵⁷ A question, by the way, which cannot be answered until we have the correct linguistic account of English. The project of giving such an account is complicated by the illegitimate restriction on the domain of viable linguistic theory by faulty arguments from underdetermination.

Chapter 3: The paradoxes of underdetermination

Epistemic underdetermination

Of what is said

Schiffer argues against the hidden-indexical theory of descriptions by way of the underdetermination of what is said. We can present the argument in the form of a paradox: a set of mutually inconsistent, though individually acceptable propositions. We begin by assuming an articulation in a context. In Schiffer's argument, this was an articulation of 'The guy's drunk' in the context of some fairly limited common knowledge of Ferdinand Pergola. For our purposes, it will be sufficient to limit this shared knowledge to two facts about Pergola: that he is the author of *Smells and Tickle*s, and that he is the man we are waiting to hear. Successful communication seems a clear possibility, which minimally requires

1. The audience can know what the speaker intended to communicate.

Whatever the speaker intends to communicate, the audience can know what it is. Interactions like this – articulations of definite descriptions in the context of a very modest stock of knowledge shared between the speaker and their audience – happen all the time. If successful communication is impossible in this case, this sceptical result will generalise to a great many, perhaps the majority, of articulations of definite descriptions.

As the case is one of literal speech, rather than implicature or the like, the Encoding Model assumes that the speaker says what they mean. In other words:

2. There is some proposition p such that the speaker says that p and, for any proposition q , if q is said then the speaker intends to communicate that q .

Now we introduce a theory of what is said by the articulated sentence. In Schiffer's argument, this was the hidden-indexical theory of definite descriptions, according to which the sentence uttered by an articulation of 'The guy's drunk' includes a hidden-indexical that refers to some contextually determined property. The problem is that the evidence available to the audience is equally consistent with two conflicting candidates for the value of the hidden-indexical. As what is said is determined by the value of the hidden-indexical, given an assignment of values

to the rest of the terms in the sentence uttered, if there are two candidates for the value of the hidden-indexical consistent with the audience's evidence, then there are two candidates for what is said consistent with the audience's evidence. If the evidence available to the audience were to settle the question of what is said, then it would settle the question of the value of the hidden-indexical. In Schiffer's Pergola case, we took the candidates for what is said to be the proposition that the guy who is the author of *Smells and Tickles* is drunk (call this P_1) and the proposition that the man we are waiting to hear is drunk (call this P_2). As P_1 and P_2 are equally viable candidates for what is said, given the evidence available to the audience,

3. The audience cannot know what is said.

There are two mutually-exclusive options:

- a) P_1 is said
- b) P_2 is said

Why are these options mutually exclusive? Perhaps the most prominent version of the Encoding Model is one that assumes sentences to be associated, in context, with at most one proposition said, although there is no reason to suppose in principle that we cannot construct a view within that model, according to which speakers can package multiple propositions into the same utterance.⁵⁸ Given that we are working within the hidden-indexical theory of descriptions, however, these options are certainly mutually exclusive, as what is said is determined by the value of the hidden-indexical and there is no value of the right type (the right type being properties) that will result in both candidates being said.⁵⁹ Given 2, if a) obtains, then the speaker intended to communicate P_1 ; if b) obtains, then the speaker intended to communicate P_2 . As the audience cannot know whether a) obtains or b) obtains, the audience cannot know whether the speaker intended to communicate P_1 or P_2 .

We therefore have the Epistemic Content Constraint:

⁵⁸ See Chapter 6 for further discussion of this option.

⁵⁹ The value of the indexical could be the property of *being the author of Smells and Tickles and the man we are waiting to hear*, but even in that case only one proposition is said, the proposition that the guy who is the author of *Smells and Tickles* and the man we are waiting to hear is drunk, which is neither P_1 nor P_2 .

4. If the audience cannot know what the speaker has said, then the audience cannot know what the speaker intended to communicate.

From 3 and 4, therefore

5. The audience cannot know what the speaker intended to communicate.

By apparently sound reasoning from individually acceptable premises, we reach a contradiction between 1 and 5. The authors we encountered in the previous Chapter tended to respond to the paradox by rejecting 3, and so by rejecting any theory according to which the audience cannot know what is said in the case at issue.

Of what is uttered

Wettstein argues for the underdetermination of what is said by way of the underdetermination of what is uttered. We can present the argument in the form of a paradox: a set of mutually inconsistent, though individually acceptable propositions. We begin by assuming an articulation in a context. In Wettstein's argument, this was an articulation of 'The table is covered with books' in the context of some fairly limited common knowledge of the table in question: that it is the only table at which the author of *The Persistence of Objects* is sitting at time *t*, and that it is the only table in room 209 of Camden Hall at *t*. Successful communication seems a clear possibility, which minimally requires

1. The audience can know what the speaker intended to communicate.

Interactions like this – articulations of definite descriptions in the context of a very modest stock of knowledge shared between the speaker and their audience – happen all the time. If successful communication is impossible in this case, this sceptical result will generalise to many, if not all, articulations of definite descriptions.

As the case is one of literal speech, rather than implicature or the like, the Encoding Model assumes that the speaker says what they mean. In other words:

2. There is some proposition *p* such that the speaker says that *p* and, for any proposition *q*, if *q* is said then the speaker intends to communicate that *q*.

Now we introduce a theory of what is uttered. In Wettstein's argument, this was the syntactic ellipsis theory of definite descriptions, according to which the sentence uttered by an articulation of 'The table is covered with books' includes a contextually-determined domain-restricting predicate. The problem is that the evidence provided by the context is equally consistent with two non-synonymous candidates for this linguistic material: 'at which the author of the Persistence of Objects is sitting at t ' and 'in room 209 of Camden Hall at t '.

Stanley and Gendler Szabó seem to take there to be at least a *prima facie* problem even when the candidates for what is uttered are synonymous in that they can only be used, in the context, to say the same thing. This is because they seemingly endorse the Epistemic Utterance Constraint independently of the Epistemic Content Constraint.

The Epistemic Utterance Constraint: If the audience cannot know what is uttered, then the audience cannot know what the speaker intends to communicate.

We, on the other hand, endorse the Epistemic Utterance Constraint only in so far as it follows from the Epistemic Expression Constraint, with the result that we take synonymous candidates for what is uttered to pose no significant issue, in agreement with Wettstein. If the audience is to know what is said, then they must be able to know what is uttered when, as in our simplified example, the only candidates for what is uttered are non-synonymous. In a situation like this, if the audience cannot know what is uttered, then the audience cannot know what is said. Supposing that the audience knows what is said, the audience can know which of the two candidates was said: whichever of the two would, in the actual context, say what was actually said.

Given that there are two equally viable, though non-synonymous, candidates for what is said that are consistent with the evidence available to the audience,

3. The audience cannot know what was said.

There are, again, two mutually-exclusive options:

- a) P_1 is said
- b) P_2 is said

These options are mutually exclusive because what is said is, according to the syntactic ellipsis account, determined by the domain-restricting predicate and we have assumed that there are only two available in the context, neither of which will determine both of these propositions.⁶⁰ Given 1, if a) obtains, then the speaker intended to communicate P_1 ; if b) obtains, then the speaker intended to communicate P_2 . As the audience cannot know whether a) obtains or b) obtains, the audience cannot know whether the speaker intended to communicate P_1 or P_2 .

We therefore have the Epistemic Content Constraint:

4. If the audience cannot know what the speaker has said, then the audience cannot know what the speaker intended to communicate.

From 3 and 4, therefore

5. The audience cannot know what the speaker intended to communicate.

By apparently sound reasoning from individually acceptable premises, we reach a contradiction between 1 and 5. Wettstein, Stanley and Gendler Szabó, and Clapp unite in solving the paradox by rejecting the third premise and the syntactic ellipsis view of definite descriptions from which it follows.

Beyond language

Benacerraf: numbers are not sets

There is an analogy to be drawn between the underdetermination arguments that we encounter in the context of linguistics and those we encounter elsewhere. Clapp (2002) notes a similarity to Benacerraf's (1965) argument that numbers cannot be sets. Benacerraf's argument begins by noting that there are different set-theoretic analyses of the natural numbers. According to one account, the natural numbers 1,2,3 ... are identical to the sets $\{\emptyset\}$,

⁶⁰ If the two predicates are available in the context, we might think their conjunction is also available. This cannot be a general truth, as it would lead to an infinite number of recursively defined candidates. I take it as clear that 'at which the author of the Persistence of Objects is sitting at t and in room 209 of Camden Hall at t and at which the author of the Persistence of Objects is sitting at t' is a worse candidate for the domain-restricting predicate than either of the simple candidates we have been considering. Even if the predicate were 'at which the author of the Persistence of Objects is sitting at t and in room 209 of Camden Hall at t', neither P_1 nor P_2 is said if this predicate determines the sentence uttered, but their conjunction, which is a third candidate for what is said.

$\{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} \dots$ while according to another account, they are identical to the sets $\{\emptyset\}, \{\{\emptyset\}\}, \{\{\{\emptyset\}\}\} \dots$

To simplify matters, let us consider just the number 3. What is it? We have two candidates on the table, that the number three is $\{\{\{\emptyset\}\}\}$ and that the number three is $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$. The problem, according to Benacerraf (1965, 62) is that “there seems to be little to choose among the accounts” as “Any purpose we may have in giving an account of the notion of number and of the individual numbers, other than the question-begging one of proving of the right set of sets that *it* is the set of numbers, will be equally well (or badly) served by any one of the” set-theoretic definitions of the number 3. Equally good, these accounts may be, but Benacerraf recognises that they are inconsistent. The number 2, for example, cannot both be a member of the number 3 and fail to be a member of the number 3, yet 2 is a member of 3 according to the first account, and fails to be a member of 3 according to the second.

This establishes an epistemic underdetermination problem. There are two inconsistent candidates for the correct set-theoretic account of the number 3 and none of the evidence available to us can settle the issue of which is correct. This epistemic problem is strangely weak, however, allowing that one of the two competing accounts of the number 3 might be correct and that we could know which one is correct if only we had more evidence. Benacerraf’s conclusion is stronger: the number 3 is not a set at all. The key premise in establishing this stronger conclusion is, as Benacerraf (1965, 67) puts it, that “all the evidence is in; if no decision is possible on the basis of it, none will ever be possible”.

We can distinguish three options:⁶¹

- a) Both accounts are correct.
- b) Exactly one account is correct.
- c) Neither account is correct.

The first possibility is absurd. b) is seriously problematic. If one account is correct, we can never know which. All the evidence is in, and it is inconclusive between the two accounts

⁶¹ Benacerraf (1965, 56)

Benacerraf (1965, 62) considers it “hardly tenable” that the answer to the question is fundamentally unknowable and therefore concludes that c) neither account is correct.⁶²

The Problem of the Many

Benacerraf’s problem is structurally similar to the problem of the one and the many.⁶³ We ordinarily take it to be possible for there to a room that contains one table and no more. A table, like any material object, changes its composition over time. Material, which once formed part of the table, becomes detached and is part of the table no longer. Take one such piece of the table, which was at one point firmly part of the table, but has now begun to detach and will, at some point in the future, clearly not be part of the table. Depending on whether we take this piece to be part of the table *now*, between the past time at which it was clearly part of the table and the future time at which it is clearly not, we have two candidates for the objects that compose the table: one object which includes the problematic piece, and one that does not. Call these T_1 and T_2 . Now, each of these collections of particles seems to make an equally good table. Either will hold up my coffee, hurt me if I collide with it, and so on for all the standard properties of tables. Here, we are not making a claim about the evidence available to us at some particular time, but about the totality of the evidence. As Benacerraf put it, all the evidence is in.

The Problem of the Many fuels the problem of underdetermination.⁶⁴ Wettstein, for example, takes the problem of underdetermination to dissolve when definite descriptions are used referentially. When using ‘The table’ referentially in ‘The table is covered with books’, the thought goes, there is no room for underdetermination of what is said because the definite description simply refers to some particular object. The Problem of the Many raises the question ‘Which object?’ There are at least two candidates for the table in question (T_1 and T_2) meaning that there are two candidates for what is said, even when the speaker intends ‘The table’ to refer to a particular object: the proposition that T_1 is covered with books, and the proposition that T_2 is covered with books.

⁶² We might doubt whether we should be able to know the nature of the number 3, but the analogous claim is much stronger in the arguments to come.

⁶³ Cf. Unger (1980).

⁶⁴ Thanks here to Patrick Greenough.

Schiffer similarly takes it to be unproblematic for an utterance of ‘The guy is drunk’ to communicate the object-dependent proposition that Pergola is drunk, but the Problem of the Many applies to persons as much as to tables. By focusing on some problematic part of Pergola, such as a piece of skin that was once firmly part of Pergola and will at some point in the future be clearly detached from Pergola, we can distinguish two candidates for the object *Pergola*, call them $Pergola_1$ and $Pergola_2$, one of which includes the problematic Pergola-part, the other of which does not. By distinguishing two candidate Pergolas, we distinguish two candidates for the object-dependent proposition that the speaker intends to communicate: the proposition that $Pergola_1$ is drunk and the proposition that $Pergola_2$ is drunk.

We have three possible responses to the Problem of the Many:

- a) Both T_1 and T_2 are tables.
- b) Exactly one of T_1 and T_2 is a table.
- c) Neither T_1 nor T_2 is a table.

These possibilities are formally very similar to Benacerraf’s. The difference is that it is far from clear which of these possibilities is superior. a) is inconsistent with our assumption that there can be such a thing as a room containing exactly one table. Whenever we think there is just one table, there are actually many. c) is inconsistent with the same assumption. Whenever we think there is a table, there are in fact none. b) makes tables a mystery. All the evidence that will ever be available is available now. If we can’t know now which of T_1 and T_2 is a table, then we can never know.

Non-epistemic underdetermination

The underdetermination in our first two arguments was *epistemic*. We argued only that the evidence available to the audience is consistent with two inconsistent candidates for what is said and therefore that the audience cannot know what was said. This is consistent with there being a fact of the matter as to what is said and with the possibility that the audience would know what was said if only they had more evidence. Stanley and Gendler Szabó (2000) rest with the epistemic version of the paradox. For them, the contradiction is focused on the question of whether a subject can come to know what was uttered by a particular articulation.

On the one hand, “in normal instances of successful communication, the hearer who grasps the proposition communicated will also know what sentence was uttered” (231) but if the syntactic ellipsis account of quantifier domain restriction is correct, then “the hearer cannot know what sentence was uttered and consequently cannot know in the normal way what proposition was meant.” (239)

Some of the other authors we have considered, however, suggest the non-epistemic conclusion. Schiffer (1995, 115) presents the epistemic considerations en route to framing the non-epistemic version of the puzzle. Having presented a number of candidates for what is said by an utterance of ‘The guy’s drunk’, he considers that, imagining himself as the intended audience, he cannot see how he could have identified any particular proposition as the one that was said. Here we have the epistemic consideration that the evidence available to the audience underdetermines what is said, which Schiffer leverages to support the conclusion that nothing is said. If something was said, then the speaker must have meant it.⁶⁵ If the speaker meant something, then the audience must be able to know what it was, given that they understand the speaker. If something was said, therefore, and the audience understands the speaker, the audience must be able to know what it was. Given that the audience understands the speaker yet cannot know, for any *p* that is a candidate for what was said, that the speaker meant that *p*, nothing was said. In this way, Schiffer moves to establish the non-epistemic version of the paradox by way of epistemic considerations.

Wettstein (1981, 246-247) clearly distinguishes the epistemic and non-epistemic considerations. In his problem cases, there are a number of candidates for the unarticulated description that features in the sentence uttered. He begins with the epistemic version of the argument when he asserts that “it is quite clearly wrong to suppose that, in many such cases, the circumstances of utterance put the listener in a position to select some one of these non-equivalent descriptions as the correct one, the one that actually captures what the speaker intended” but moves swiftly on to the problem he considers “more important” that “none of these Russellian descriptions is *the correct one*”. At this point, he notes “our concern is no longer

⁶⁵ Schiffer reserves terms like ‘said’ for propositions that are consistent with the context-invariant meaning of a sentence and that are meant. We are using a different terminology, yet given that this is a case of literal speech, our terminologies are equivalent here.

merely the epistemic one”. In a similar fashion to Schiffer (1995), Wettstein uses the epistemic considerations as a stepping stone to framing the non-epistemic argument.

We can return to our previous arguments and apply Benacerraf’s reasoning to extend our conclusions from the epistemic to the non-epistemic. We have seen that, if the hidden-indexical theory of descriptions is correct, the evidence available to the audience is insufficient for them to know what the speaker said in Schiffer’s Pergola case. This leaves open the possibility that something was in fact said and that the audience could come to know what was said if only they had more evidence. Suppose that the audience acquires all the evidence they can, even asking the speaker what they meant. We might think that even in this case, what is said will remain underdetermined. There simply is no evidence that could settle the matter one way or another.

Again, we have three options:

- a) Both accounts are correct.
- b) Exactly one account is correct.
- c) Neither account is correct.

a) is absurd, according to the hidden-indexical theory of definite descriptions. What is said is determined by the value assigned to the hidden indexical, and the indexical can only be assigned properties, but there is no property we can assign as the value of the indexical that would deliver two propositions said. Looking at this another way, there are two candidates for what is said: that the guy who authored *Smells and Ticks* is drunk, and that the guy who is the man we are waiting to hear is drunk. If the first is said, then the value of the hidden-indexical is the property of being the author of *Smells and Ticks*; if the second is said then the value of the hidden-indexical is the property of being the man we are waiting to hear. The value of the indexical cannot be both.⁶⁶

If b) then there is a fact as to what was said, but it is fundamentally unknowable. No matter what evidence the audience gathers, they cannot know what the speaker said. Wettstein (1986, 247) considers this a very strange results, saying that it is “implausible in the extreme” to

⁶⁶ Although the value of the indexical could be the property of *being the author of Smells and Ticks and the man we are waiting to hear*, but then what is said would be the proposition that the guy who is the author of *Smells and Ticks* and the man we are waiting to hear is drunk, which is a third and distinct candidate for what is said.

suppose that the speaker intended to communicate one of the candidates in particular but that we cannot “even with the help of the speaker himself” come to know what it was.⁶⁷

We are left with c). Both accounts cannot be correct, as they are inconsistent, yet it cannot be that one of these accounts is correct and the other is not, so both are incorrect. According to the hidden-indexical theory, only one of the candidates could be said. Neither of these candidates is said, however, so, according to the hidden-indexical theory, nothing is said. The point extends easily to the syntactic ellipsis theory.

Of what is said

In framing a non-epistemic version of the problem of underdetermination for what is said, we can begin the argument in the same way, by assuming an articulation. Again, we can illustrate through Schiffer’s Pergola case in which a speaker articulated ‘The guy’s drunk’ in the context of some fairly limited common knowledge about Ferdinand Pergola. Again, for our purposes, it will be sufficient to limit this shared knowledge to two facts about Pergola: that he is the author of *Smells and Ticks*, and that he is the man we are waiting to hear.

As the case is one of literal speech, rather than implicature or the like, the Encoding Model assumes that the speaker says what they mean. In other words:

1. There is some proposition *p* such that the speaker says that *p* and, for any proposition *q*, if *q* is said then the speaker intends to communicate that *q*.

The speaker said something and they intended to communicate what they said. Recall, however, that there are a number of distinct candidates for what is said, given the hidden-indexical theory of descriptions. The epistemic problem of underdetermination rested on the claim that the audience cannot know which of two candidates was said, which is consistent with the possibility that something was said. If this is the case, then there must be some evidence unavailable to the audience which determines what is said. Yet, it seems as though the audience has all of the evidence that could be relevant to determination of what is said.

⁶⁷ The full quote: “Surely it is implausible in the extreme to suppose that in fact one of these descriptions captures what the speaker intended but that we cannot, even with the help of the speaker himself, come to know which description that is.” Here, Wettstein is considering various non-synonymous candidates for what was uttered, but the point works just as well for different candidates for what is said.

If the audience lacks some evidence relevant to determining what is said, what evidence do they lack? Perhaps the best suggestion along these lines is that the speaker's intentions determine what is said. All the evidence available to the audience is inconclusive as to what is said, but that is because the speaker's intentions are not part of the evidence available to the audience. It is, however, unclear why the speaker would have the intention to say some particular proposition, given that the audience cannot identify the proposition that the speaker intends. Further, it seems that speakers' intentions, if they truly do determine what is said in these cases, are opaque to the speaker themselves, who won't usually recognise themselves as having meant either of the candidates for what is said to the exclusion of the others.

If we assume that the evidence available to the audience exhausts the evidence relevant to determining what is said, then

2. If something is said, then many things are said.

The two candidates for what is said are the best, and equally good, candidates for what is said, given the evidence available to the audience, which is all the evidence relevant to deciding what is said. If something is said, then it has to be one of these candidates, given that they are the best, but if one of these candidates is said, then so is the other, given that they are equally good. If one of them meets the conditions for being said, then so does the other. Many things cannot be said, however, according to the hidden-indexical theory of descriptions. The proposition said is a function of the value assigned to the hidden indexical and no value of the appropriate type (the type being properties, or functions from objects to truth values, for the hidden-indexical theories we have encountered) will deliver more than one proposition.

3. At most one thing is said.

From 2 and 3 therefore

4. Nothing is said.

By apparently sound reasoning from individually acceptable premises, we reach a contradiction between 1 and 4. If the hidden-indexical theory of descriptions is correct, therefore, then there is usually no proposition that the speaker says and means. The majority of uses of definite descriptions are therefore more like cases of implicature than cases of literal speech, which may

be seen in itself as an unpalatable conclusion. Indeed, these cases are more like the unusual cases of implicature, such as exploitations of the Maxims of Manner, in which communication is achieved despite nothing being said. Rather than take this option, the authors we surveyed in Chapter 2 opt to reject any linguistic theory which allows for underdetermination. Chapter 4 will instead motivate an alternative model, which allows for underdetermination of what is said, even in cases of literal speech.

Of what is uttered

We can again reach the same conclusion by consideration of what is uttered. Again, we can illustrate through Wettstein's case involving an articulation of 'The table is covered with books' in the context of some fairly limited common knowledge of the table in question: that it is the only table at which the author of *The Persistence of Objects* is sitting at time *t*, and that it is the only table in room 209 of Camden Hall at *t*.

As the case is one of literal speech, rather than implicature or the like, the Encoding Model assumes that the speaker says what they mean. In other words:

1. There is some proposition *p* such that the speaker says that *p* and, for any proposition *q*, if *q* is said then the speaker intends to communicate that *q*.

The speaker said something and they meant what they said. Recall, however, that here are a number of distinct candidates for what is said, given the syntactic ellipsis view of definite descriptions. The epistemic problem of underdetermination rested on the claim that the audience cannot know which of two candidates was said, which is consistent with the claim that something was said, although the audience cannot know what it was. If this is the case, then there must be some evidence unavailable to the audience which determines what is said. Yet, it seems as though the audience has all of the evidence that could be relevant to determination of what is said.

If we assume that the evidence available to the audience exhausts the evidence relevant to determining what is said, then

2. If something is said, then many things are said.

The two candidates for what is said are the best, and equally good, candidates for what is said, given the evidence available to the audience, which is all the evidence relevant to deciding what is said. If something is said, then it has to be one of these candidates, given that they are the best, but if one of these candidates is said, then so is the other, given that they are equally good. If one of them meets the conditions for being said, then so does the other. Many things cannot be said, however, according to the syntactic ellipsis account of definite descriptions. The proposition said is a function of the contextually-determined domain-restricting predicate and no predicate will deliver more than one proposition.⁶⁸

3. At most one thing is said.

From 2 and 3 therefore

4. Nothing is said.

By apparently sound reasoning from individually acceptable premises, we reach a contradiction between 1 and 4. Again, this shows that, if the syntactic ellipsis view is correct, there is often no proposition that the speaker says. While it may be unpalatable to represent the majority of definite descriptions as saying nothing and so more akin to some of the more unusual cases of implicature, such as exploitations of the Maxims of Manner mentioned in Chapter 1, the following Chapter will explain how communication can be achieved despite the non-epistemic underdetermination of what is said, while sharply distinguishing between these cases and cases of implicature.

⁶⁸ Right enough, there is the predicate ‘at which the author of the Persistence of Objects is sitting at *t* and in room 209 of Camden Hall at *t*’ but that determines one more complex sentence, rather than a pair of sentences. We have been assuming, for simplicity, that there are only two equally viable candidates for what is uttered. Given that the conjunction of these candidates is fairly simple, yet more informative, it might seem a clearly superior candidate. Given that, in many realistic situations, speakers and audiences will share a great deal more common knowledge, however, we will not always be able to conjoin the simple candidates without forming a horribly over-complex conjunction.

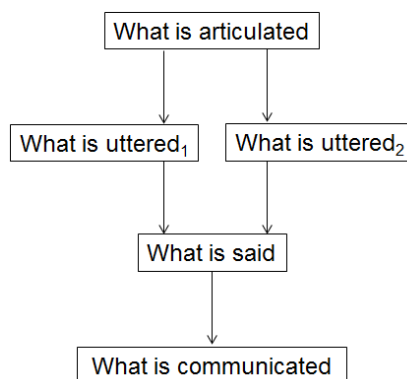
Chapter 4: Responding to the paradoxes

Embracing underdetermination

In Chapter 2, we saw that various authors have presented the problem of underdetermination in aid of rejecting some linguistic theory. Taken together, these arguments severely limit the domain of viable linguistic theory; a surprising result, given that the theories in question are independently plausible. In Chapter 3, we saw that the problem extends even to referential terms such as proper names. In contrast to the approach taken by the authors surveyed in Chapter 2, this Chapter proposes to solve the supposed problem of underdetermination by denying that underdetermination is problematic.

Cases of underdetermination are taken at face value. If the *prima facie* plausible hidden-indexical and syntactic ellipsis analyses that we have considered are correct, then there are various mundane cases in which an audience can know what a speaker intends to communicate, despite being unable to know what is said. The paradoxes of the previous section are avoided, not by denying any linguistic hypothesis that leads to underdetermination of what is said, as previous authors have done, but by denying the adequacy of the Encoding Model, according to which communication centres (at least in cases of literal speech) around the saying of some proposition that the speaker intends to communicate.

This Chapter will present an alternative according to which the saying of any proposition is inessential to communication. Unlike the Encoding Model, this alternative model recognises that there might be various distinct ways of interpreting an utterance, each of which leads to precisely what the speaker intended to communicate. We got a taste of this model in Chapter 2 when we noted that the underdetermination of what is uttered is unproblematic when the same thing would be said by all the candidates for what is uttered.




To avoid underdetermination of what is uttered, however, we must be able to explain how different candidates for what is said can communicate the same information. If the speaker intends to communicate some implicature derived from what is said, it is easy to see how different candidates for what is said would communicate the same proposition. You can, for example, implicate that I am a poor philosopher by either complimenting my handwriting or by complimenting my vocabulary. Even if I cannot know what you say, in the right context I can know what you mean if the candidates for what is said are the proposition that my handwriting is excellent and the proposition that my vocabulary is excellent.

It is, however, more difficult to see how the same result could be achieved in cases of literal speech, which we previously took to support the Epistemic Content Constraint:

The Epistemic Content Constraint: If the audience cannot know what the speaker has said, then the audience cannot know what the speaker intended to communicate.

According to the Encoding Model, if *p* is said, then the speaker meant that *p*. So far so good. Where we go wrong, however, is in assuming that, if *p* is said, then *p* is the only proposition meant by the speaker. On the contrary, when the saying of a proposition communicates the proposition said, it also communicates propositions in addition. Given a context, therefore, in which the best candidates for what is said are *p* and *q* but the saying of *p* would communicate that *p* and that *q*, and the saying of *q* would communicate that *q* and that *p*, the audience needn't know what the speaker has said to know what they intended to communicate.



The next section will explain how it is that the saying of a proposition can, in cases of literal speech, communicate propositions beyond what is said. We will then observe the phenomenon of *local equivalence*. When two propositions are locally equivalent, the same thing is communicated whichever is said. As we shall see, the cases from Chapter 2 are characterised by the local equivalence of each candidate for what is said, explaining how communication can be entirely successful when what is said is underdetermined.

What is said and what is communicated: one to many

Consider the following pair of cases. In each case, one fairground worker utters to another, referring to the same person in each case, “John is 135cm tall”. In the first case, the rules of the fairground stipulate that all and only those over 130cm tall are permitted to ride. In the second case, the rules stipulate that all and only those over 140cm tall are permitted to ride. In each case communication succeeds, but the speaker intends to communicate something different in each case. In the first case, the speaker intends to communicate that John is tall enough to ride, while in the second case, the speaker intends to communicate that John is not tall enough to ride. An audience who understands the utterance might naturally report it by saying “I was told that John was/wasn’t permitted to ride” and only upon reflection (if at all) recognise that this was not quite what the speaker said by the words uttered. In each case, I take it that the speaker also communicates that John is 135cm tall, but the audience fails to understand the full significance of the speaker’s utterance if this is all they take the speaker to mean.

On the assumption that a single proposition is said, many propositions are communicated. The propositions communicated vary between the two cases, even though what is said is held fixed. Depending on your view of proper names and height attributions, you might be drawn to quibble with the example on this point but as relevantly similar cases could be constructed for any sentence whatsoever, it will be a simple matter for you to amend the case if you think that what is uttered or said varies between these two cases.

We have, of course, allowed for cases in which the speaker communicates many propositions by the saying a single proposition. By saying that I have excellent handwriting, you might communicate that I am a lousy philosopher, in addition to communicating what you said. The thought was that speakers standardly mean precisely what their words mean but the audience will sometimes be expected to work out some further implicature. Against this suggestion, notice that divergence in what is communicated is perfectly adequately explained by entailment, without recourse to implicature. Different propositions are entailed by what is said in each of the fairground cases and these entailments are themselves communicated.⁶⁹ How can the entailments of a proposition differ while the proposition itself is held fixed? The answer is that different things are taken for granted in each case and what is entailed, in a particular context, by what is said depends on which premises are taken for granted in the context. I am assuming that the fairground workers know the ride's height restrictions and take them for granted in their fairground dealings. If all and only those over 130cm are permitted to ride and John is 135cm tall, then John is permitted to ride. If all and only those over 140cm are permitted to ride and John is 135cm, then John is not permitted to ride.

Grice notes several characteristics of implicatures that distinguish them from entailments.⁷⁰ Grice (1989: 31) tells us that conversational implicatures should be calculable from his maxims of conversation as an inference to the best explanation. In the fairground case, inference to the best explanation plays no role. The difference in what is communicated is fully explained by deductively valid entailments. Nor need the Maxims feature in an

⁶⁹ This is not to say that all entailments of what is said are communicated. Many of them will be common ground already, and we cannot communicate anything that is already taken for granted. Such is the case for all the necessary truths, for example.

⁷⁰ By 'implicatures' I mean 'conversational implicatures'. I take it as clear that these are not cases of generalised (Grice (1989, 37-38)) or conventional (Grice (1989, 25-26)) implicature. The words used do not generally, or by convention, carry any particular implicature.

explanation of the case, when the speaker and their audience share the required presuppositions.⁷¹

Implicatures are cancellable, while entailments are not. “I have five children. I don’t have only five though” can be acceptable, because the proposition that the speaker has five children *and no more* is implicated, rather than entailed, by what is said. On the other hand, “I have five children. I don’t have three though” is never acceptable (except perhaps when it is being used to introduce an implicature; perhaps the implicature that I don’t really care how many children I have) because the proposition that the speaker has three children is not implicated but entailed by what was said by the first sentence. The different propositions communicated in the fairground cases are not cancellable. If the speaker had uttered, in the first case, “John is 135cm tall. John isn’t permitted to ride though” the audience could be expected to be at a complete loss: if John is 135cm tall, then John *must* be permitted to ride! If something is communicated by this utterance, it is not due to the cancellation but the introduction of an implicature; perhaps an implicature to the effect that the audience is mistaken as to the rules of the ride.⁷²

Perhaps most importantly, Grice (1989, 39) tells us that “Since the truth of a conversational implicature is not required by the truth of what is said (what is said may be true – what is implicated may be false), the implicature is not carried by what is said, but only by the saying of what is said”. Here, Grice clearly distinguishes between implicatures and entailments. In the fairground case, the different propositions communicated are required by the truth of what is said. At least from the perspective of the participants (but entirely objectively if the rules of the park are known, rather than only presupposed, by the conversational participants), it must be true that John is/isn’t permitted to ride, given the truth of what is said.⁷³ Although there are possibilities at which what the speaker says is true and yet John is not permitted to ride (it is possible, that is, for the rules of the park to be different from the way the speaker and

⁷¹ As we found a role for the Maxims in the determination of what is uttered and said, we can find a role for them in presupposition accommodation. Supposing that the audience fails to presuppose the rules of the ride, the audience might be able to work out the speaker’s presuppositions by recognising that the speaker violates the Maxims of Conversation unless they are presupposing the rules of the ride.

⁷² Grice (1989, 39) also allows for implicatures to be contextually cancellable, but this applies only to generalised implicatures.

⁷³ Grice (1989, 39) also mentions “nondetachability” as a feature of implicatures. As entailments are clearly nondetachable, this criterion cannot help us to distinguish between entailments and implicatures.

their audience take the rules to be) all such possibilities are excluded from consideration by the speaker's presuppositions. The speaker's only concern in our example is ensuring that the park's rules, as they stand, are adhered to, not assessing the rules themselves.

This raises an important point, which is that successful interpreters needn't always engage in an additional interpretive process when what the speaker intends to communicate diverges from what is said. In the fairground case, the speaker's meaning falls out of what they say, not because they say exactly what they intend to communicate, or because the audience derives what is meant by resolving an implicature, but because the speaker's meaning is entailed by what they say, given the presuppositions in effect. Given the audience's presupposition that all and only those over 140cm are permitted to ride, the truth of the proposition that John is 135cm requires the truth of the proposition that John is not permitted to ride. In representing the one proposition as true by saying it explicitly, the other proposition is thereby represented as true. This explains the data about cancellation. In saying one proposition, an entailed proposition is represented as true, so the entailment can only be cancelled if what is said is cancelled, leaving the assertion empty, or if the speaker implicates the falsity of the audience's presuppositions.

Although implicature doesn't feature in the fairground cases, we can note the same phenomenon in cases where implicature is all important. Alter the case so that the proposition that John is 135cm tall is implicated, rather than said. In such cases, the speaker intends to communicate what they implicate in each case, yet we still see divergence between what is implicated and what is communicated as a function of the presuppositions in effect. Depending on the height restrictions the audience presupposes, they may correctly take the speaker to mean that John is, or that John is not, permitted to ride. So long as the speaker intended the audience to interpret relative to these presuppositions, and formulated their articulation appropriately, communication will be successful.

The divergence of what is meant and what is said is particularly clear in the fairground cases because the speaker and their audience presuppose that all and only those who are over 130cm/140cm are permitted to ride. The point is entirely general, however. Speakers never intend to communicate only what they say. Given that utterances always take place against a background of presuppositions, the relation between the propositions said and propositions communicated is always one-to-many. By saying that *p*, while presupposing that *q*, therefore, a

speaker communicates that $p \& q$. The communicated content of an utterance is never exhausted by what is said.

Local equivalence

If the saying of a proposition communicates multiple different propositions, then it doesn't follow from a multiplicity of candidates for what is said that the audience cannot know what the speaker intended to communicate, even on the assumption that the speaker intended to communicate what they say. Supposing two candidates for what is said, p and q , it is true that the speaker intends to communicate that p if p is said, and true that the speaker intends to communicate that q if q is said, but this observation fails to establish the Epistemic Content Constraint. What we require in addition are the claims that, if p is said, then the speaker did not intend to communicate that q and that, if q is said, then the speaker did not intend to communicate that p . This claim seems very natural when we view communication along the lines of the Encoding Model, where we view communication as the packing and unpacking of the content that the speaker intends to communicate, as any content not packed into the sentence will not standardly be communicated. Given that what is meant or communicated comes apart from what is said, however, there is the possibility that the same propositions would be communicated by each of the candidates for what is said, given the presuppositions in effect in the context.

Stalnaker (1999: 49) notes as much in pointing out that in a context where we both presuppose that my neighbour is an adult male, what is communicated is identical whether I say that my neighbour is unmarried or that my neighbour is a bachelor.⁷⁴ As in the fairground cases, what is communicated diverges from what is said. By saying that my neighbour is unmarried when it is presupposed that he is an adult male, I communicate, not only that he is unmarried, but also that he is a bachelor. If my neighbour is an adult male, then if my

⁷⁴ There is a potential conflict between what Stalnaker says in this passage and what he says in his later work. In this passage, Stalnaker could be read as suggesting that the conversational situation is identical whether I articulate 'My neighbour is a bachelor' or 'My neighbour is unmarried'. Against this, Stalnaker (1999: 86) suggests that the conversational situation is different whenever a different sentence is articulated. This potential conflict can be avoided, however, by considering a situation in which there are two different interpretations of the proposition said by a single articulated sentence, such as when I say 'My neighbour is too', given a context in which the properties of *being unmarried* and of *being a bachelor* are equally salient.

neighbour is unmarried, my neighbour is a bachelor.⁷⁵ Indeed, the reverse is also true: if my neighbour is an adult male, then if my neighbour is a bachelor, my neighbour is unmarried.⁷⁶ Whichever of these propositions is said, therefore, both are communicated. Given the context, these two propositions are, I shall say, *locally equivalent* in that the same thing would be communicated, whichever of these propositions were said.⁷⁷ Whether I say that my neighbour is unmarried, or I say that my neighbour is a bachelor, the same thing is communicated: that my (adult male) neighbour is unmarried (and so a bachelor). Given a context in which two propositions are mutually entailing, then, *ceteris paribus*, the same thing is communicated whichever is said.⁷⁸

Stalnaker's framework for assertion offers one way to model local equivalence. What is taken for granted or *presupposed* in a particular conversation determines a set of possibilities known as the *context set*, composed of all and only those possibilities at which everything presupposed is true. Presupposition is an attitude distinct from knowledge and belief, although conversational participants often presuppose whatever they take to be commonly known or believed by everyone involved in the conversation. The constitutive role of assertion is to discriminate between the possibilities in the context set by adding what is said to what is presupposed, thereby eliminating from the context set all those possibilities at which the proposition said is false.⁷⁹

The speaker's communicative intention is not to distinguish between absolutely all possibilities (even if sense can be made of unrestricted quantification over possibilities) but

⁷⁵ Perhaps there are in fact more restrictions on bachelorhood (is the Pope a bachelor?) but we can join Stalnaker in ignoring these subtleties for the sake of the example.

⁷⁶ In the fairground case, the two propositions communicated are asymmetric in that the saying of one communicates the other but the saying of the other does not communicate the first one. If all and only those over 130cm tall are permitted to ride, then John is permitted to ride if John is 135cm tall but John is not necessarily 135cm tall if he is permitted to ride.

⁷⁷ Stalnaker (1999, 49) says that they communicate the same "*increment of information*".

⁷⁸ The *ceteris paribus* clause rules out, for example, the situation in which one is said in a way too complicated for the audience to comprehend. In that situation, nothing might be communicated when that proposition is said.

⁷⁹ This is not quite true in more complex cases. When the speaker is being sarcastic, for example, the successful interpreter will add the negation of what is said to the common ground. More generally, assertions add to what is presupposed the proposition that Stanley and Gendler Szabó called 'what is communicated'. There is a complication, however. Stanley and Gendler Szabó use the term 'what is communicated' to denote the proposition derived by resolution of sarcasm, as well as whatever the speaker intends ultimately to convey. These two must be distinguished because the speaker can convey different information by sarcastically implying the same proposition against different presuppositional backgrounds, just as they can convey different information by saying the same proposition against different presuppositional backgrounds. In what follows, I will continue to use 'what is communicated' to denote the message ultimately conveyed by the speaker.

rather to distinguish between the possibilities in the context set. What they intend to communicate is not therefore captured simply by what is said, but by the effect of what is said on the context set; that's why it is unclear what I intend to communicate by saying something that is already presupposed and so eliminates no possibilities from the context set. In Stalnaker's bachelor case, the speaker could impose the same effect on the context set by saying either of two different propositions, so what is communicated would be identical either way. If all possibilities in the context set are such that my neighbour is an adult male, the proposition that my neighbour is unmarried and the proposition that my neighbour is a bachelor have the very same effect on the context set: elimination of all possibilities at which my (adult male) neighbour is married, leaving only those at which he is unmarried (and therefore a bachelor).⁸⁰

In Stalnaker's framework, we can see the propositions communicated as those that were not true at all possibilities in the context set before the update, but are true at all possibilities in the context set after the update.⁸¹ In this case, both the proposition that my neighbour is unmarried and the proposition that my neighbour is a bachelor are communicated.⁸² Note, however, that the communication of multiple propositions requires no more cognitive work on the part of the speaker or the audience than the communication of a single proposition. The additional propositions communicated emerge in the incorporation of what is said into the context set, a process that is no more complicated for the fact that many propositions are communicated.


Stalnaker's framework allows for a very picturesque representation of local equivalence. Two propositions p and q are locally equivalent, relative to a context set C , if and only if $p \cap C = q \cap C$,⁸³ that is, only if they make the same cut in the context set, as pictured below:

⁸⁰ Communication doesn't rely on the acceptance of what is said, but what is said is incorporated into the background presuppositions only if it is accepted, so what is communicated is a result of considering the effect of incorporating what is said into the background presuppositions, rather than actually doing so.

⁸¹ Note that this need not include presuppositions introduced by the articulation of the sentence uttered, such as the presupposition that the speaker uttered a particular sentence, which are not, according to Stalnaker (1999, 86) part of the "ESSENTIAL effect" of assertion. See Chapter 6 for further discussion.

⁸² Representing what is communicated by two propositions is a matter of choice on the part of the theorist. We might as well say that a single conjunctive proposition is communicated.

⁸³ Here, the propositions p and q are represented as sets of possibilities. Specifically, each proposition is represented by the set of possibilities at which it is true. What about more fine-grained conceptions of propositions, according to which propositions may be individuated by properties beyond their modal truth-value profiles? According to one sort of view, two propositions can be distinguished by their constituents, even if they



In this diagram, the square represents a context set, while the dotted rectangle and solid rectangle represent different propositions that might be said. As the intersection of each proposition with the context set is precisely the same, these propositions are locally equivalent and the same propositions are communicated, whichever is said.

The observation of local equivalence provides a further argument for the one-to-many relation between propositions said and propositions communicated. Suppose two propositions p and q are locally equivalent in a context. As what is communicated is determined by the effect of a proposition on the context set, the same thing is communicated by saying p and by saying q . Saying q communicates that q while saying p communicates that p , so saying either p or q will communicate that p and that q . Further, relative to any presupposition r , the proposition that p and the proposition that $p \& r$ are locally equivalent, so the saying of any proposition, relative to any presupposition, communicates something in addition to what is said.

have the same truth value at all possible worlds. According to this line, two candidates for what is said may communicate different things, despite having the same truth value, because they differ in terms of their constituents. The proposition that my neighbour is unmarried and the proposition that my neighbour is a bachelor may take the same truth value at all worlds that interest us but, for example, these propositions may communicate different things by making different properties available for future reference. One proposition is partly constituted by the property of *unmarriedness*, thereby making this property available for further reference, while the other is partly constituted by the property of *bachelorhood*, thereby making this distinct property available for further reference. In what follows, I maintain that differences between propositions are relevant to what is communicated by saying those propositions only if they entail a difference in the modal profiles of those propositions. It is consistent with this view, however, to allow that different propositions may have different collateral effects, such as making different objects available for future reference.

So far we have been less than fully explicit as to what is communicated. I argued that what is communicated in the fairground case includes a proposition entailed by what is said and that, when two propositions are locally equivalent, both are communicated by either one of them. One way to accommodate these observations is to allow that p is communicated if and only if p is true at every possibility in the intersection of what is said with the context set. According to this view, however, we communicate what is already presupposed, which includes everything that we have already communicated.

A superior account is that p is communicated if and only if p is true at every possibility in the intersection of what is said with the context set, but not true at every possibility in the context set. We avoid the problem of communicating everything that is presupposed but face what may be a *prima facie* problematic result when we combine this account with Stalnaker's (1999, 86) claim that speakers presuppositions include, not only propositions about the subject matter of their conversation, but also about the conversational situation in which they are discussing the subject matter, such as "that a speaker is speaking, saying the words he is saying in the way he is saying them". If these include, for example, the presupposition that the conversation is taking place in a particular language, then a speaker of English and a speaker of German communicate different propositions by saying synonymous sentences. The English speaker presupposes that the language being spoken is English. Call this ' q '. The German speaker presupposes that the language being spoken is German. Call this ' r '. For any proposition p that the speakers communicate by their utterances of their synonymous sentences, therefore, the speaker of English will communicate also communicate the proposition that p and q , while the German speaker will communicate that p and r .

There are a number of potential responses to this potential problem. The first is to accept the conclusion but to deny that it is problematic. The speakers of different languages communicate some different things but they also communicate many of the same things, according to the above account of what is communicated. When comparing utterances spoken in different languages, we find it easy to abstract from those things communicated which were parasitic on presuppositions about the conversational situation and those that were not, although the details of how this is done are negotiable. Perhaps we might define 'what is communicated' so as to eliminate the problematic propositions and so allow for a perfect

correspondence of ‘what is communicated’ in English and German, but perhaps doing so introduces a notion that fails to cut at the joints of Stalnaker’s account.

An alternative is to abandon the letter of Stalnaker’s account for one according to which presuppositions about the conversational situation are represented separately from presuppositions about the subject matter of the conversation. The view presented in Lewis (1979), for example, is amenable to Stalnaker’s suggestion that we represent presuppositions by sets of possible worlds, but there is no requirement that presuppositions about the conversational situation and subject matter of the conversation be represented by the same entry on the “conversational scorecard”. There seems, in principle, no reason why such a more fine-grained representation of the conversational context cannot capture all phenomena that can be captured according to the letter of Stalnaker’s account.

Although there isn’t much to be gained here by surveying all the possible options, it is worth noting that the alternative to the Encoding Model that is being defended here is not tied to Stalnaker’s framework and local equivalence can be modelled in other frameworks. According to a file change model of discourse of the sort found in Heim (1983a, 226), the state of a discourse can be represented by a system of files. What is said by ‘She is old’ is different from what is said by ‘The baker is old’ as there are systems of files that each of these sentences will update differently. In a discourse where there are two distinct files, one and only one concerning a female and the other and only the other concerning a baker, ‘She is old’ will add ‘is old’ to the female’s file, and ‘The baker is old’ will add ‘is old’ to the baker’s file. Given a discourse in which there is just one file, however, and that file concerns a female baker, both would have the same effect: to update that file with ‘is old’. Relative to such a discourse, the different file change potentials of these two sentences are locally equivalent. I stick with Stalnaker’s framework only because the equivalence of propositions is captured so conspicuously, in the identity of the intersection of each proposition with the context set.

Let’s revisit the cases from Chapter 2 to see how the candidates for what is said are locally equivalent, given the right presuppositions, and that the assumption of local equivalence is central to the persuasiveness of the argument.

Responding to the underdetermination arguments

Schiffer against the ‘naive’ hidden-indexical theory of descriptions

Consider Schiffer’s Pergola case. You and I are waiting to hear a talk by eminent philosopher Professor Ferdinand Pergola. As the Professor stumbles up to the stage, you say to me “I’ll be damned! The guy’s drunk”. According to the hidden-indexical theory of descriptions, you utter a sentence containing an indexical that refers to some contextually-determined property H. What you say is true if and only if there is a single individual that is a guy and H, and that individual is drunk.

Underdetermination looms because “Even before your utterance it was mutually evident to us that we had knowledge of the professor under numerous shared definite descriptions” such as ‘the eminent author of *Smells and Tickles*’, and ‘the man we are waiting to hear’, leading to a number of candidates for the property H and therefore what is said. That this common knowledge be mutually evident is not in itself sufficient for an underdetermination problem, however. Even if we all know that Pergola authored *Smells and Tickles*, there may, for example, be a convention that we don’t treat it as common knowledge in our conversations; perhaps it was an embarrassing book that all engaged in the discipline would prefer to forget. At any rate, we need something stronger than common knowledge to get the problem going. There must be numerous shared definite descriptions that we could use in the conversational context to refer to Pergola, that is, it must not only be commonly known but presupposed that Pergola satisfies numerous definite descriptions.

If it is presupposed that these descriptions pick out the same individual (that is, that the eminent author of *Smells and Tickles* is the man we are waiting to hear) then it is clear why underdetermination looms. I cannot take you to refer to any of these properties in particular, as each would pick out Pergola equally well relative to our background assumptions, but the very same presupposition also ensures that the candidates for what is said are locally equivalent.

If the eminent author of *Smells and Tickles* is the man we are waiting to hear, then the eminent author of *Smells and Tickles* is drunk if and only if the man we are waiting to hear is drunk. At every possibility in the context set there is an individual who is both the eminent author of *Smells and Tickles* and the man we are waiting to hear. Each of the candidates for

the proposition said therefore recommends elimination of the same possibilities from the context set: all those at which it is not the case that the eminent author of *Smells and Ticks* (who is the only man we are waiting to hear) is drunk. Whether the speaker said that the eminent author of *Smells and Ticks* is drunk, or that the man we are waiting to hear is drunk, both propositions would be communicated. Although what is said may be underdetermined because speaker and interpreter presuppose that their referents satisfy multiple descriptions, this presupposition ensures that the candidates are locally equivalent and so that the audience will know what the speaker means to communicate, whichever of the viable interpretations they take.

Schiffer (1995, 115) notes that he cannot see how the audience “could have identified any one definite description, however complex, as *the one* that figured in the proposition you asserted.” I agree. If Schiffer’s usage of ‘asserted’ is equivalent to my usage of ‘said’, the audience cannot know which of various propositions was asserted, whether that is because nothing is asserted, as per the non-epistemic underdetermination argument from Chapter 3, or because something is asserted but the audience lacks the evidence relevant to figuring out what it was. The audience doesn’t need to know which of the locally equivalent candidates was actually asserted in order to know what the speaker intended to communicate. If, on the other hand, Schiffer’s usage of ‘asserted’ is equivalent to my usage of ‘meant’ or ‘intended to communicate’, then I again agree with Schiffer. The audience cannot identify *the one* definite description that figured in the proposition asserted because there is no one such description. The speaker intends their utterance to communicate several propositions, each figuring its own descriptive content.

To see how the assumption of local equivalence of the candidates is central to setting up the problem of underdetermination in the first place, consider the case when the two ‘candidates’ for what is said would have a different effect on the context set, thereby communicating different propositions. Suppose, for example, that all possibilities in the context set are such that the man we are waiting to hear is the man who is currently staggering up to the stage but only some possibilities in the context set are such that the author of *Smells and Ticks* is the man we are waiting to hear. We are open, that is, to the possibility, but do not presuppose, that the man we waiting to hear is the author of *Smells and Ticks*.

If we now try to pose the problem of underdetermination with our original candidates for what is presupposed, it falls flat, admitting of a clear solution. Did the speaker say that the author of *Smells and Ticks* is drunk or that the man we waiting to hear is drunk? Clearly, between these two candidates, the proposition that the man we are waiting to hear is drunk is a far better candidate for what is said, given the context in which your comment is clearly directed at the man who is currently staggering up to the podium. Rather, if we want to pose the problem of underdetermination in this context, the best candidates for what is said are the proposition that the man we are waiting to hear is drunk and the proposition that the man staggering up to the podium is drunk. Given that these propositions are locally equivalent in the context set, however, each will communicate the same information.

If we simply define the context so that there are two equally viable candidates for what is said that are locally non-equivalent, then we undermine the first premise of the epistemic argument from Chapter 3 and the problem is again avoided. Suppose, the context just outlined, for example, in which the proposition that the author of *Smells and Ticks* is drunk and the proposition that the man we waiting to hear is drunk partitions the context set in different ways, thereby communicating different propositions. Now simply stipulate that these are the best candidates for what is said. For these propositions to be the best candidates for what is said, it must be no more plausible that the speaker intends to communicate one of these propositions than the other, in which case, the theoretically neutral verdict is that the audience cannot know what the speaker intends to communicate. In this context, it is not only unclear what proposition you have said, according to the hidden-indexical theory of descriptions, but unclear, on any theory, what guy you are talking about and what you intend to communicate.⁸⁴

Wettstein against Russell's analysis of definite descriptions

Wettstein (1981) presented arguments against a syntactic ellipsis defence of Russell's theory of descriptions, according to which an articulation containing an incomplete description such as 'The table is covered with books' should be interpreted as an utterance of a sentence in which

⁸⁴ We might be tempted to suggest that you intend to communicate the object-dependent proposition that Pergola is drunk but in that case there is a clearly superior candidate for what is said, according to the hidden-indexical account: the proposition that the guy who is identical to Pergola is drunk.

the incomplete description is expanded so as to be uniquely denoting; for example, ‘The table in room 209 of Camden Hall at t_1 is covered with books’.

The problem, according to Wettstein (1981, 246) is that “the table the speaker has in mind can be more fully described in any number of ways, by the use of any number of non-synonymous, uniquely denoting descriptions”. While there might be any number of true descriptions of the table, however, this doesn’t necessarily pose an underdetermination problem because these descriptions might not all be consistent with what is presupposed by the speaker and their audience. If we mistakenly presuppose, for example, that the author of *The Persistence of Objects* is not sitting at the table in room 209 of Camden Hall, then these two descriptions cannot both be candidates for the elided material, even if they are in fact both true of the same table. The problem arises only when there are two non-synonymous descriptions that are both available for picking out the relevant table relative to our presuppositions. The most obvious such case is that in which we presuppose that the author of *The Persistence of Objects* is not sitting at the table in room 209 of Camden Hall.

If we presuppose that, at t_1 , the author of *The Persistence of Objects* is sitting at the only table in room 209 of Camden Hall, the candidates for the sentence uttered include ‘The table in room 209 of Camden Hall at t_1 is covered with books’ and ‘The table at which the author of *The Persistence of Objects* is sitting at t_1 is covered with books’. These candidates are non-synonymous in that what is said varies depending on which of these sentences the speaker uttered. If these sentences are equally viable candidates for what the speaker uttered, therefore, there are at least two distinct candidates for what the speaker said. Given that there are a number of candidates for what is said, Wettstein (1986, 247) asks “how are we to account for the fact that the speech act is not at all indefinite or indeterminate?” the suggestion being that there is no way to explain how, as he puts it (1981, 248), “a fully determinate assertion is made” in the face of indeterminacy as to what is said.

The answer is that, while the sentence articulated and surrounding context fail to determine what is said, they successfully restrict the candidates for what is said to propositions that communicate the same information. Given the two candidates for what is uttered, there are two candidates for what is said, according to the Russellian analysis; one that is true if and only if there is one and only one table at which the author of *The Persistence of Objects* is sitting at t_1 and it is covered with books, the other true if and only if there is one and only one table in

room 209 of Camden Hall at t_1 and it is covered with books. As it is possible for these propositions to vary in truth value, they are distinct. One possibility in which they differ in truth value, for example, is the situation in which there are two distinct tables, one being the only table at which the author of the Persistence of Objects is sitting at t_1 , the other being the only table in room 209 of Camden Hall at t_1 , only one of which is covered with books. If any such possibilities feature in the context set, then these two candidates for what is said would communicate different propositions.

No such possibilities feature in the context set, however, if the speaker and their audience presuppose that there is a particular table that is both, at t_1 , the table at which the author of the Persistence of Objects is sitting and the only table in room 209 of Camden Hall. Given this presupposition, the table at which the author of *The Persistence of Objects* is sitting at t_1 is covered with books if and only if the table in room 209 of Camden Hall at t_1 is covered with books. At every possibility in the context set, there is a single table which, at t_1 , is both the only table in room 209 of Camden Hall and is the only table at which the author of *The Persistence of Objects* is sitting. Whichever of the candidates was said, therefore, the very same possibilities would be eliminated from the context set – all those at which the table at which the author of *The Persistence of Objects* is sitting at t_1 (which, at any possibility in the context set is the table in room 209 of Camden hall at t_1) is not covered with books. Contrary to Wettstein's assumption, we can account for the determinacy of the speech act while denying that anything is said, because all the best candidates for what is said may communicate the same propositions.

Blackburn against Wettstein's referential account

After mistakenly taking the Russellian account to be refuted by the problem of underdetermination, Wettstein instead defends an analysis of definite descriptions according to which they must be completed with some directly referential term. Blackburn notes, however, that the problem of underdetermination re-emerges on this account. Given a context in which we know that the very same table is below a particular chandelier and above a particular rug, candidates for the sentence uttered by an articulation of 'The table is covered with books' include 'The table below *that chandelier* is covered with books' and 'The table above *that rug* is covered with books' where the italicised terms are taken to be directly referential. While the motivation for Wettstein's view was mistaken, Blackburn's argument against the view is

mistaken for the very same reason. While these candidates may have different truth values at possibilities outwith the context set, they can take the same truth value at all possibilities within the context set, if our common knowledge is presupposed. If the table under *that chandelier* is the table above *that rug* then the table under *that chandelier* is covered with books if and only if the table above *that rug* is covered with books. Each of the candidates for what is said therefore communicates the same propositions, by eliminating all and only those possibilities in the context set at which the table (which is under a certain chandelier and above a certain rug) is not covered with books. Having eliminated these possibilities, the speaker has introduced two new presuppositions into the context set. It is now true at all worlds in the context set that the table below *that chandelier* is covered with books and that the table above *that rug* is covered with books.

Stanley and Gendler Szabó against syntactic ellipsis

Stanley and Gendler Szabó's case of underdetermination is essentially the same as Wettstein's, extended to pose a problem for quantifier domain restriction outwith the context of Russell's theory. We are offered a syntactic ellipsis analysis of 'Every bottle is empty' according to which the sentence uttered by an articulation of 'Every bottle is empty' includes an unarticulated domain-restricting predicate. Stanley and Gendler Szabó note, however, that there are various different predicates, synonymous and, more importantly for our purposes, non-synonymous that are equally well-suited to play the role of the unarticulated predicate. Stanley and Gendler Szabó (2000, 238) conclude that "the hearer cannot know what sentence was uttered and consequently cannot know in the normal way what proposition was meant".

While Stanley and Gendler Szabó don't tell us quite what the 'normal way' amounts to, I take these cases of underdetermination to stand as a significant obstacle to the assumption they consider "immensely plausible ... that in normal instance of *successful* communication, the hearer who grasps the proposition communicated will also know what sentence was uttered and what proposition was [said] by that sentence on the given occasion" (231) If the speaker and their audience make the right presuppositions, then the audience can grasp the propositions communicated without knowing what was uttered or said. Stanley and Gendler Szabó might respond that this is an abnormal case but it's hard to see what could support this

objection while the correct linguistic theory of English is still a matter of debate. Given the state of the debate, it seems entirely possible that the normal case is the one just described.

In this case, the presupposition required to secure local equivalence of the candidates is the presupposition that ‘I purchased today’ (as spoken by Lisa) and ‘in the cupboard’ apply to all and only the very same bottles, that is that every bottle Lisa purchased on the day is in the cupboard. Given this presupposition, what is uttered (and so what is said) is underdetermined because each of the candidates for the domain-restricting predicate would pick out the same bottles relative to any possibility consistent with Lisa and Max’s presuppositions. For the very same reason, however, the candidates for what is uttered are locally equivalent. If the bottles in the cupboard are identical to the bottles that Lisa purchased today, then every bottle in the cupboard is empty if and only if every bottle Lisa purchased today is empty. Every possibility in the context set is such that there is a particular collection of bottles that constitutes both every bottle that Lisa just bought and every bottle in the cupboard, so each of the candidates for what is said again recommends the same update to the context, that is, elimination of all possibilities at which it is not the case that every bottle that Lisa purchased today (i.e. every bottle in the cupboard) is empty.

Clapp against traditional semantics

Clapp (2002) brings these arguments together to argue against traditional semantics, characterised by

“Truth Conditional Compositionality: The truth conditions of an utterance are a function of (i) the logical form of the utterance (i.e., the structure of the LF of the utterance), and (ii) the meanings of the words in the utterance (i.e., the semantic values of the terminal nodes of the LF of the utterance).”

As what is said is often underdetermined, whether we adopt a syntactic ellipsis or hidden-indexical analysis, Clapp concludes that the truth conditions of some utterances cannot be a function of logical form and meaning. According to the present view, however, these problems show only that there are cases in which utterances lack truth conditions, that is, in which nothing is said. This is consistent with Truth Conditional Compositionality, which, on a standard analysis, tells us only that *if* something is said, it is a function of LF and semantic

value. Although these utterances lack truth conditions in the sense defined by Truth Conditional Compositionality, however, they determine the propositions that the speaker intends to communicate. If, on the other hand, Truth Conditional Compositionality is intended to require, of the utterances in our examples, that they are associated with a unique truth condition as a function of their logical form and the meaning of the words, then we agree with Clapp that the intended principle is false.

Underdetermination without traditional semantics

Completion and expansion

It is worth noting that Truth-Conditional Compositionality is not part of the Extended Model in Chapter 1, although we assumed it for the sake of simplicity in several examples. Rather, we allowed that what is said might be derivable only on the basis of completion or expansion, as presented in Bach (1994). This is one of the papers cited by Clapp (231 and footnote 2) as an example of the sort of “*dynamic semantics* or *truth conditional pragmatics*” that should replace the ‘traditional semantics’ characterised by the acceptance of Truth Conditional Compositionality. Bach (1994), Carston (1991), and Sperber and Wilson (1986) all reject Truth Conditional Compositionality by allowing that what is said by an utterance depends not only on the structure of the sentence uttered and the semantic values of its terms, but also by features of context that drive the processes of completion and expansion. ‘Steel is not strong enough’ requires, according to Bach (1994, 127) completion by “some contextually identifiable respect” in which steel is not strong enough. Bach suggests *for building a 500-storey building* and *to resist bending by Superman*. Even allowing for this process of completion, we can again manufacture cases of underdetermination by allowing for a modest degree of knowledge shared between the speaker and their interpreter. When they presuppose, for example, that steel is strong enough for building a 500-story building if and only if it is strong enough to resist bending by Superman (perhaps Superman has taken to bending all buildings over a certain height) then there needn’t be any justification for privileging one of these candidates as giving the truth conditions of the utterance.

This is precisely the problem of underdetermination that faced the traditional semantic theories we looked at previously and is unproblematic for precisely the same reason. By

articulating ‘Steel is not strong enough’ in this context, the speaker has not done enough to ensure that something is said, but they have done enough to restrict all the candidates for what is said to locally equivalent propositions. It is presupposed that Steel is not strong enough for building a 500-story building if and only if it is not strong enough to resist bending by Superman. At every possibility in the context set, therefore, either steel is strong enough for building a 500-story building and steel is strong enough to resist bending by Superman, or Steel is not strong enough either for building a 500-story building and steel is not strong enough for resist bending by Superman. Whether the speaker said that steel is not strong enough for building a 500-story building or said that steel is not strong enough to resist bending by Superman, the same thing would be communicated as the very same possibilities would be eliminated from the context set: all and only those at which steel is strong enough for building a 500-story building (that is, all and only those at which steel is strong enough to resist bending by Superman).

Austinian Semantics

We saw underdetermination issues arise also for the Austinian proposal in Recanati (1996), according to which what is said by an utterance is not captured simply by the meaning of the sentence uttered (against Truth Conditional Compositionality) but also by a contextually-determined situation. What is said is true if and only if the contextually-determined situation is of the type determined by the meaning of the sentence uttered. The suggested case was one in which I say ‘Every bottle is empty’ in a context in which we know that every bottle in my apartment is in my kitchen. There are therefore two equally viable candidates for the contextually-determined situation – my apartment and my kitchen – leading to two equally viable candidates for what I said: that every bottle in my kitchen is empty and that every bottle in my apartment is empty.

The solution, again, lies in the recognition that what is said determines what the speaker intends to communicate only in conjunction with their presuppositions. I may not have done enough to ensure that something is said by my utterance, but I have done enough to ensure that the same thing is communicated by each of the candidates for what is said by my utterance. If all the bottles in my kitchen constitute every bottle in my apartment, then every bottle in the apartment is empty if and only if every bottle in my kitchen is empty. By

saying either proposition, therefore, we eliminate the very same possibilities from the context set – all those at which every bottle in my apartment (that is, every bottle in my kitchen) is empty – thereby communicating both that every bottle in my apartment is empty and that every bottle in my kitchen is empty. Even if nothing is said, as per the non-epistemic underdetermination argument from Chapter 3, the speaker can successfully communicate both propositions by restricting the viable candidates for what is said to these propositions. Whichever of these viable interpretations the audience takes, the speaker will successfully communicate the same things.

Discourse Representation Theory

Finally, we considered Discourse Representation Theory (DRT), as described in Kamp and Reyle (1993).⁸⁵ Despite some significant differences from the reference-based semantics that we assumed for the sake of most of our examples, DRT faces precisely the problem of quantifier domain restriction that Stanley and Gendler Szabó were concerned to solve. Both the syntactic ellipsis and hidden-indexical strategies are open to DRT and both accounts lead to underdetermination of what is said. Again, however, the underdetermination is unproblematic relative to the right assumptions.

Clapp is correct that DRT can avoid the problem of underdetermination. This is not, however, because DRT can avoid underdetermination, or because DRT denies Truth Conditional Compositionality but because DRT has the potential to model existing conversational presuppositions as features of an existing DRS against which utterances can be communicated. If the existing DRS contains the information that, for all x , x is a bottle Lisa bought yesterday if and only if x is a bottle in the cupboard, then the truth conditions of the resultant DRS will be identical whether the speaker says that every bottle Lisa bought is empty, or that every bottle in the cupboard is empty.

How to build an underdetermination argument

The common feature of the underdetermination arguments addressed in Chapter 2 is a number of candidates for what is said that are true at all and only the same possibilities in the

⁸⁵ We also considered the S-View of Travis (1985), also cited by Clapp (2002), but concluded that the view was too vague to be engaged with in much depth.

context set, but vary in truth value at possibilities outwith the context set. In order to build an underdetermination argument, we can identify a candidate for what is said and vary its modal profile outwith the context set.

In Chapter 3, the Problem of the Many was leveraged to extend the problem of underdetermination to the referents of singular terms. ‘Pergola’ refers to some person called Pergola, but we noted that there are several different candidates for the physical object that constitutes Pergola, depending on whether we, say, take this particular hair to be a part of Pergola or not. We have two candidates for Pergola and so two candidates for what is said by ‘Pergola is drunk’: that Pergola_1 is drunk, and that Pergola_2 is drunk. The audience cannot know what is said. If we follow the response of the authors from Chapter 2, we might conclude that ‘Pergola’ is not, after all a referential term. While these candidates for what is said are different propositions that take different truth values at various possibilities, however, they coincide in truth value at every possibility in the context set. Pergola_1 is drunk if and only if Pergola_2 is drunk. If one of these Pergolas could be drunk without the other being drunk, then we would have a criterion by which to distinguish them. The Problem of the Many would then not apply and only one of these propositions could be a candidate for what is said.

We can extend the problem of underdetermination to any term. Take the very first example from Chapter 1: an utterance of ‘John runs’. A general formula for identifying alternative candidates for what is said by a sentence is to consider a far-fetched possibility already excluded from the context set and ask whether the sentence is true. The answers ‘yes’ and ‘no’ identify two different candidates for what is said. In this case, we might ask whether ‘John runs’ is true if John has tentacles for legs, which he uses to propel himself swiftly over the ground. Does John run in this situation? I, for one, do not know, but the positive and the negative answer identify two different candidates for what is said.

The audience cannot know what was said by an utterance of ‘John runs’. Again, following the response from Chapter 2 might lead us to reject the view of ‘runs’ according to which it denotes a function from objects to propositions. There are at least two candidates for this function – the one which delivers a proposition true at the tentacle world, and one which delivers a proposition that is false at the tentacle worlds – and therefore two candidates for what is said. As the audience cannot know what is said, the Epistemic Content Constraint demands that the audience cannot know what the speaker intended to communicate.

Thankfully, audiences do not have to distinguish between these candidates in order to know what the speaker intends to communicate, as both of these candidates for what is said are true at all and only the same possibilities in the context set, given the presupposition that John has legs for legs, rather than tentacles.

Chapter 5: An account of the referential/attributive distinction

The referential/attributive distinction

Donnellan (1966) describes two uses of definite descriptions, such as ‘the murderer of Smith is insane’. His initial characterisation is as follows,

“A speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so. A speaker who uses a definite description referentially in an assertion, on the other hand, uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing.” Donnellan (1966, 285)

Donnellan illustrates attributive use by the following case. Suppose that we come across Smith’s murdered body. From the condition of Smith’s corpse, you infer that the killer must have been a lunatic. You say to me, “The murderer of Smith is insane”. Assume (although this is not essential to the phenomenon) that we have no suspects. There is no particular individual that you mean to draw my attention to by using the term ‘the murderer of Smith’. You intend to talk about the murderer of Smith, whoever that may be, because the only evidence for your assertion is the manner of the killing, which could have been committed by anyone. You said something true if it turns out that Jones murdered Smith and is insane, but likewise you said something true if it turns out that Atkinson or Stephenson, or anyone else, murdered Smith, provided that they are insane. In this case, you could not have said what you meant by replacing the definite description with a proper name such as ‘Jones’ or ‘Atkinson’. As Donnellan (1966, 285) puts it, “the definite description might be said to occur essentially, for the speaker wishes to assert something about whatever or whoever fits that description ... the attribute of being the so-and-so is all important”.

Compare a referential use of the same description. Jones stands in the dock accused of murder and we are observing his trial. From his erratic behaviour, you infer that he is a lunatic. You say to me “The murderer of Smith is insane”. In contrast to the previous case, there is a particular individual that you have in mind and intend to draw my attention to, and the evidence for your claim is the behaviour of that particular individual: Jones. Unlike the attributive case, you were right only if it turns out that Jones is insane; if Jones is sane, then

you were wrong, even if it turns out that Atkinson murdered Smith and is insane. When used referentially, the definite description actually used is often inessential to the speech act. You could have said what you meant (although the nuances would not have been exactly the same) by using a proper name, another description, a pronoun, or a demonstrative. As Donnellan (1966, 285) puts it, in the referential case “the definite description is merely one tool for doing a certain job-calling attention to a person or thing-and in general any other device for doing the same job, another description or a name, would do as well.”

This fits well with Donnellan's initial characterisation of the distinction. The role of a referentially used definite description is to allow one's audience to identify the object about which one wishes to state something or other, and any term that the audience takes to refer to the same individual can do this just as well. In the illustrative case, I take 'Jones' and 'the murderer of Smith' to pick out the same individual, so it matters not which term you use.

This brief discussion suggests a criterion by which referential and attributive uses of definite descriptions can be distinguished,

- 1) The truth conditions of an utterance may be unaffected by replacing a referentially used definite description with a term (a proper name, pronoun, demonstrative, or description) that denotes the same object.

This makes sense of Donnellan's claim that other terms could “do as well” as a referential use of a definite description. For one term to do just as well as another, all else being equal, that term should not alter the truth conditions of the utterance.

Donnellan (1966, 287-8) notes another way in which attributive and referential uses come apart,

- 2) If an utterance contains a definite description that correctly applies to nothing, the utterance can be used to state something of a particular individual only if the description is used referentially.⁸⁶

⁸⁶ This is not to say that an attributive use cannot be successful in any way under these conditions, only that they cannot succeed in their characteristic role. If the speaker's intention is just to say something false, then they may succeed.

Suppose that Smith actually died in a freak accident after a brutal (but non-fatal) beating, and so that ‘the murderer of Smith’ correctly applies to nothing. There is no murderer of Smith, but if I have used the description referentially, then I may have successfully predicated insanity of someone by saying “The murderer of Smith is insane”, successfully told you to bring me someone by saying “Bring me the murderer of Smith”, or successfully asked you about someone by saying “Is the murderer of Smith insane?”. The role of attributive uses is to pick out whatever satisfies the description. If nothing satisfies the description, this role cannot possibly be fulfilled. Referential uses, on the other hand, are intended to bring some particular individual to the attention of one’s audience, and this intention may well be fulfilled even if that individual in fact fails to satisfy the description.

According to some interpretations, 2) does not do justice to Donnellan’s view. Donnellan did not only think that referential uses of empty definite descriptions could be used to state things; he thought they could be strictly and literally true. Donnellan (1966, 298) says, for example, that “Using a definite description referentially, a speaker may say something true even though the description correctly applies to nothing.” Donnellan expands upon this claim, however, by telling us that “The sense in which he may say something true is the sense in which he may say something true *about someone or something*” (my italics) and that “This sense is ... an interesting one that needs investigation” which suggests that this is not the sense in which we ordinarily talk about the truth and falsity of assertions. Similar hedges⁸⁷ are found throughout Donnellan (1966) and leave room for an interpretation according to which Donnellan is committed to nothing stronger than the view that one may make a true claim by a

⁸⁷ “As we have seen, it is possible for the linguistic purpose of the speech act to be accomplished in such a case even though nothing fits the description; it is possible to say something true or to ask a question that gets answered or to issue a command that gets obeyed. For when the definite description is used referentially, *one’s audience may succeed in seeing to what one refers* even though neither it nor anything else fits the description.” (p. 292)

“But where the definite description is used referentially, something true may well have been said. It is possible that something true was said *of the person or thing referred to*.” (p. 293)

“If we think about what the speaker said about the person he referred to, then there is no reason to suppose he has not said something true or false *about him*” (p. 302)

All italics my own. A notable exception follows on p. 302 when Donnellan says simply that “when a speaker uses a definite description referentially he may have stated something true or false even if nothing fits the description.”

referential use, even though nothing in fact satisfies the description. He need not be committed to the much stronger view⁸⁸ that the sentence one uses is strictly and literally true.

According to Donnellan (1966, 297) attributive and referential uses of definite descriptions are not to be distinguished by positing an ambiguity:

“The grammatical structure of the sentence seems to me to be the same whether the description is used referentially or attributively, that is, it is not syntactically ambiguous. Nor does it seem at all attractive to suppose an ambiguity in the meaning of the words; it does not appear to be semantically ambiguous.”

Rather, Donnellan (1966, 288) suggests that a further difference between referential and attributive uses of definite descriptions lies in the speaker’s presuppositions,

- 3) “when a definite description is used referentially, not only is there in some sense a presupposition or implication that someone or something fits the description, as there is also in the attributive use, but there is a quite different presupposition; the speaker presupposes of some particular someone or something that he or it fits the description.”

When you used the description ‘the murderer of Smith’ attributively, there was no particular individual who you presupposed to have murdered Smith, although you presupposed that there was a murderer. When using the description referentially, however, you presupposed not only that someone or other was the murderer of Smith, but that *Jones* was the murderer of Smith. If it turned out that the lunatic Atkinson murdered Smith, you got something wrong in the referential case, but not the attributive.

It is important to distinguish what is conversationally presupposed from what either the speaker or their audience believes. We often presuppose things that we do not believe, and probably always refrain from presupposing things that we do believe. Suppose the judge is considering Jones’s sentence. Healthy murderers go to prison, while the mentally ill go to a secure hospital. The judge and all those present in the courtroom know, on the basis of overwhelming evidence, that Jones murdered Smith. The judge says “The murderer of Smith is

⁸⁸ Which Reimer (1998, 94), Wettstein (1981, 241-244), and Soames (1994, 153-154) take to be problematic.

insane; we know that from the condition of the body. 'The evidence that Jones committed the crime is overwhelming; therefore I conclude that Jones is insane.' The judge's reasoning is valid, but would beg the question if they were using the description 'The murderer of Smith' referentially, concluding that Jones is insane from a premise that is equivalent (given 1) to the premise that Jones is insane. The judge uses the definite description attributively, failing to presuppose that Jones is the murderer, despite believing that Jones is the murderer.⁸⁹ Now suppose that Smith's family know that Smith committed suicide, and so that no one fits the description 'the murderer of Smith', but that Smith's suicide was an attempt to frame Jones for murder. To help the plan succeed, Smith's family refer to Jones as 'The murderer'. Here the definite description is used referentially, so the family presuppose that Jones is the murderer, despite knowing that Jones is not the murderer.⁹⁰

The problem for Russell

Donnellan takes referential uses to pose a substantial problem for Russell's (1905) analysis of definite descriptions. According to Russell's analysis, by uttering a (context-insensitive) sentence of the form 'The F is G', the speaker says that there is one and only one F and that F is G. Assuming that Russell's semantic analysis of definite descriptions cannot be extended to proper names,⁹¹ something different is said by an utterance of 'The murderer is insane' and an utterance of 'Jones is insane', which is commonly thought to be the 'singular proposition' that is true if and only if Jones is insane. As it is possible that Jones is sane, but the unique murderer of Smith is not, these propositions are distinct. If Jones is sane, but the lunatic Atkinson murdered Smith, then it is false that Jones is insane but true that the unique murderer of Smith is insane. We will assume that the truth conditions of utterances involving

⁸⁹ Such examples not only help to distinguish presupposition from belief, but show how the referential/attribution distinction is relevant to the distinction between valid and invalid reasoning. An example from Rostworowski (2013, 35) shows how the distinction is also relevant to the distinction between informative and uninformative assertion. We have just found out that Smith was murdered by their nephew and I say, triumphantly, "I have suspected from the beginning that the murderer of Smith was a member of his family!" This assertion may be totally uninformative (and a little ridiculous) if the definite description is read referentially, in which case it is equivalent to saying "I have suspected from the beginning that Bobby Smith was a member of Smith's family!" as we might have known from the outset of our investigation that the two were related.

⁹⁰ See Donnellan (1966, 290-291) for similar examples.

⁹¹ The view that Russell's analysis of definite descriptions cannot be extended to proper name, while commonly assumed, is contrary to the view expressed in Russell (1998) that "Common words, even proper names, are usually really descriptions".

proper names, pronouns, and indexicals are given by such *singular* or *object-dependent* propositions.

It may appear that Russell's analysis cannot apply to all uses of definite descriptions. According to 1, the same thing is said by utterances of referentially used definite descriptions and by utterances of proper names. To account for the phenomena, we may think that we need to restrict Russell's analysis of definite descriptions to their attributive uses. This Chapter will show how Russell's analysis can accommodate referential uses. The key is, once again, to recognise that communicated content is not determined simply by what is said, but by what is said in combination with the presuppositions, allowing that different things may be communicated by referential and attributive uses of definite descriptions, even though the same thing is said.

A second objection can be easily set aside. According to Donnellan (1966, 291), while “Both the attributive and the referential use of definite descriptions seem to carry a presupposition or implication that there is something which fits the description ... the reasons for the existence of the presupposition or implication are different in the two cases.”⁹² When a definite description is used referentially, the speaker's goal is to draw the attention of their audience to a particular person or thing that they want to say something about. Although this intention might be realised even if nothing satisfies the description, one's audience is more likely to identify the intended referent if it actually fits the description, so using a definite description referentially indicates the presupposition that the intended referent satisfies the definite description. As presupposition is closed under logical consequence, this presupposition entails the weaker presupposition that something or other satisfies the definite description.

When a definite description is used attributively, on the other hand, the presupposition that something or other satisfies the definite description does not stem from the stronger presupposition that something in particular satisfies the definite description. In attributive cases, the speaker's intention is to say something about whoever or whatever uniquely satisfies the definite description. Unlike the referential case, this intention cannot possibly be realised if

⁹² This Chapter will focus on presupposition, rather than implication, but the two are easily connected. By producing an utterance that presupposes that *p*, I can implicate that *p*.

nothing satisfies the description. Intention requires the presupposition of potential success, so the speaker must presuppose that something or other uniquely satisfies the description.

According to Russell, ‘The φ is ψ ’ logically entails ‘There exists one and only one φ ’. Donnellan (1969, 292-3) objects that “Russell’s theory does not show - what is true of the referential use - that the implication that something is the φ comes from the more specific implication that what is being referred to is the φ ”. This objection simply misunderstands the purpose of Russell’s analysis, which is a semantic account of the truth conditions of sentences, not an account of the origin of speaker presuppositions. A more serious objection would be that Russell’s view is inconsistent with the facts about the origin of presuppositions in the referential case. Russell’s view is consistent with the facts about presupposition, however. It may be that ‘The φ is ψ ’ entails ‘There exists one and only one φ ’ and yet that speakers presuppose that something or other satisfies ‘ φ ’ because they presuppose that something in particular satisfies ‘ φ ’. This specific presupposition explains both the more general presupposition and why the speaker uses a sentence which entails that more general presupposition.

The account of the referential/attributional distinction offered in the next section differs substantially from Stalnaker’s (1999, 40-44) account, although both are presented in terms of the same broad framework. One significant difference is that Stalnaker’s account sanctions what he terms a “pragmatic ambiguity” in which a univocal rule governing the use of a term can be deployed in two different ways, either in the derivation of content or merely of truth value. Such ambiguities play no role in the account presented in the following section; interpretation follows the same procedure whether a definite description is used referentially or attributionally, the only difference being the presuppositions entertained by the interpreter, as stated in 3. Later, when we discuss presupposition failure, we will see a further difference between Stalnaker’s view and the view presented here, specifically that the view presented here is more amenable to a unified explanation of non-catastrophic presupposition failure.

A Russellian account of the distinction

This section will utilise the account of local equivalence outlined in Chapter 4 to explain how Russell’s account of definite descriptions can accommodate referential uses of definite

descriptions. Recall the three distinguishing features of referential usage that we extracted from Donnellan (1966):

- 1) The truth conditions of an utterance may be unaffected by replacing a referentially used definite description with a term (a proper name, pronoun, demonstrative, or description) that denotes the same object.
- 2) If an utterance contains a definite description that correctly applies to nothing, the utterance can be used to predicate something of a particular individual only if the description is used referentially.
- 3) “when a definite description is used referentially, not only is there in some sense a presupposition or implication that someone or something fits the description, as there is also in the attributive use, but there is a quite different presupposition; the speaker presupposes of some particular someone or something that he or it fits the description.”

According to Russell's analysis, the same thing is said by a (context-insensitive) description, whether used referentially or attributively. Referential and attributive uses are associated with different presuppositions, however, and it is this feature of the distinction which explains the other two.

According to 3, the appropriate use of a definite description requires the presupposition that something fits the description. A speaker who uses the definite description ‘the φ ’ attributively by saying ‘The φ is ψ ’ presupposes that there is one and only one thing that is φ , but doesn't presuppose of any particular thing (of Jones, or of Atkinson say) that it is the unique φ . Given Stalnaker's account of presupposition, all possibilities in their context set are such that there is a unique φ , but different things may fill that role in different possibilities; at one possibility within their context set, Jones may be the unique φ , while at another possibility within their context set, Atkinson may be the unique φ .

Referential usage requires an additional presupposition. A speaker who uses ‘the φ ’ referentially by saying ‘the φ is ψ ’, presupposes not only that something is the unique φ , but that a particular thing – Jones, say – is the unique φ . All possibilities in the speaker's context set are

therefore such that there is a unique φ , but unlike the attributive case there is a particular object x such that x is the unique φ at every possibility in the context set.

1) requires that the truth conditions of an utterance may be unaffected by replacing a referentially used description with another term that denotes the same object. Assuming a Russellian analysis of definite descriptions and a direct reference account of proper names, demonstratives, and pronouns, substituting a definite description for a proper name always changes the truth conditions of the utterance. Thus, Russell's view seems to be inconsistent with Donnellan's primary observation concerning referential uses of definite descriptions.

The appearance is illusory, however, arising from equivocal use of 'the truth conditions of the utterance'. One and the same utterance may be associated with different truth conditions at the levels of what is said and what is meant. If 1) says that there is no difference between what is said by the use of a sentence containing a proper name and the use of a corresponding sentence in which the proper name has been replaced by a referentially-used description, then we beg the question as to the correct account of the referential/attribution distinction by assuming a semantic explanation from the outset. The assumption is not only question begging, but independently implausible. If what is said by 'The murderer of Smith is insane' is identical to what is said by 'Jones is insane', then what is said by 'The murderer of Smith is insane' can be true even if there is no murderer of Smith, or if the murderer of Smith is perfectly sane.

Rather, what we should take as a datum in theorising about the distinction is the fact that what the speaker communicates may be identical, whether use a sentence containing a proper name or the corresponding sentence in which the proper name has been replaced by a referentially-used description. It follows from Russell's view of descriptions that the same thing is said whenever a (context-insensitive) sentence of the form 'The F is G ' is uttered, but this is consistent with such utterances communicating different propositions. As noted in Chapter 4, the saying of a proposition will always communicate locally equivalent propositions. Given the presuppositions that Donnellan takes as characteristic of the referential use of a definite description, it is easy to show how Russell's unitary account of definite descriptions vindicates 1).

3) encapsulates Donnellan's observation that attributive and referential uses of definite descriptions are characterised by different presuppositions. In the referential murderer case,

the speaker must presuppose that a particular individual – Jones – is the unique murderer of Smith. All possibilities in the context set are therefore such that there is one and only one murderer of Smith, and Jones is that murderer. Possibilities at which someone other than Jones is the unique murderer of Smith are ignored. If Jones is the unique murderer of Smith, then the unique murderer of Smith is insane if and only if Jones is insane. Whether the speaker asserts the proposition that Russell associates with ‘The murderer of Smith is insane’ or the proposition that Jones is insane, which we assume to be semantically associated with ‘Jones is insane’, the very same possibilities will be eliminated from the context set, all those at which Jones (the unique murderer of Smith at every possibility in the context set) is sane. Whichever of these sentences is used, therefore, the same propositions are communicated. Saying that Jones is insane communicates that Jones is insane (when we speak literally). As this proposition is locally equivalent to the Russellian content of ‘The murderer of Smith is insane’, uttering this sentence also communicates that Jones is insane.

2) requires that only referential uses of definite descriptions can be successful (can successfully make statements, issue commands, or ask questions about particular individuals) if nothing satisfies the description used. The fact that attributive uses of definite descriptions cannot be successful in these conditions is clear from what has already been said. The point of attributive use is simply to pick out whatever uniquely satisfies the description. If nothing does, this intention cannot possibly be successful. In the referential case, however, the purpose of the description is just to pick out a particular individual, and this can be achieved even if nothing in fact satisfies the description.

There is a clear way in which you can pick out Jones and issue statements, commands, or questions about him, by using the description ‘The murderer of Smith’, even if no one murdered Smith. You need only presuppose that Jones murdered Smith, thereby satisfying condition 3 of referential usage. In that case (as we have already seen), saying that the murderer of Smith is insane will locate the actual world in the same subset of the context set as saying that Jones is insane. By uttering ‘The murderer of Smith is insane’ while presupposing that Jones is the unique murderer of Smith, you can communicate that Jones is insane. Although the actual world is excluded from the context set by the false presupposition that Jones is guilty, you distinguish between the worlds in the context set on the basis of Jones’s sanity. You

have presupposed something false (that Jones is the murderer) in order to communicate something true (that Jones is insane).⁹³

Yablo (2006, 167-168) suggests that there is a *prima facie* problem here. The purpose of assertion, on Stalnaker's model of discourse, is to partition a set of possibilities into a set that contains the actual world and a set that does not. If the actual world is already excluded from the context set, "it is hard to see how ... the speaker is saying anything about actuality. It's as though I tried to locate Sicily for you by saying that as between North and South Dakota, it's in the North, although truth be told it's not in either Dakota."

We have assumed that the purpose of assertion is to indicate that the actual world is in a particular subset of the context set, but this is in itself a fairly uninteresting aim, unless all members of that subset have some interesting unifying property. If the members of the set have such a property, the real import of the utterance is the indication that the actual world has this property, and the actual world may indeed have the property indicated, whether or not it is a member of the indicated set.⁹⁴

Referential uses of definite descriptions can communicate propositions that predicate properties of particular individuals, even if nothing in fact fits the description used. Not so in the attributive case. In that case, the only proposition communicated was that there is one and only one murderer of Smith and they are insane. If we discover that Smith wasn't murdered, we will be at a loss if we try to identify anyone of whom insanity was predicated.⁹⁵

As the same proposition is said whether the description is used referentially or attributively, what is said is false if Smith wasn't murdered. The difference in the referential case is that there are other interesting propositions available to facilitate reinterpretation of the utterance. In the original context, the speaker communicated precisely what they would have done by uttering the sentence 'Jones is insane'.⁹⁶ The only reason for the speaker to have

⁹³ One way to read the above is as arguing that proper names are equivalent to definite descriptions. Exercise caution regarding this interpretation though. I have not argued that proper names are semantically equivalent to descriptions. On the contrary, I have assumed that Russell's analysis of descriptions cannot be extended to proper names and argued only that these distinct propositions are locally equivalent relative to certain presuppositions.

⁹⁴ To push the Dakota analogy a bit, the locations within South Dakota all have the property of being south of the border between North and South Dakota. If I tell you that Sicily is in South Dakota, one thing I might indicate is that the location of Sicily shares this property, and that is true even though Sicily is not in either of the Dakotas.

⁹⁵ Although see the final section for reasons to doubt this.

⁹⁶ Although the nuances would not have been exactly the same. 'The murderer of Smith', for example, may have the additional effect of introducing the presupposition that there is a unique murderer of Smith.

uttered ‘The murderer of Smith is insane’, rather than the sentence including a proper name, was their presupposition that Jones is the murderer of Smith, which rendered the propositions associated with these sentences locally equivalent. Once this presupposition is shown to be false, however, there is no sense to be made of saying the proposition associated with ‘the murderer of Smith is insane’, which fails to partition any context set when none of its members have a unique murder of Smith. The original utterance did not only communicate that the murderer of Smith is insane, however, but also that Jones is insane and this proposition can partition the members of the new context set.

Such reinterpretation will not always be permissible, of course. In many cases, the evidential basis for all of the propositions communicated will be removed when the falsity of the presupposition is revealed. In the case at issue, however, the evidence for the speaker’s belief that Jones is insane was Jones’s observed behaviour, and this provides evidence for Jones’s insanity whether or not he is the unique murderer of Smith.

By referential use of a definite description, one can communicate truths about Jones even if Jones fails to fit the description you use. Relative the presupposition that Jones is the murderer of Smith, we can communicate that Jones is insane by saying ‘The murderer of Smith is insane’. Jones is innocent, so the actual world is excluded from the context set, but although the propositions communicated by the utterance were determined relative to the possibilities within the context set, they are defined relative to possibilities outwith this set; Jones may be insane, even if he is no murderer. This explains how presupposition failure can be non-catastrophic and a speaker may succeed in making a claim, despite presupposition failure. Even after the recognition of presupposition failure it can be reasonable to take the speaker as committed to some of the claims they made while making a false presupposition.

Strong evidence for this view of communication comes from the obvious pedantry, in many cases, of pointing out the falsity of a presupposition. Sometimes the speaker needn’t care whether the actual world is excluded from the context. All they need care about is whether it shares a property with a particular subset of the context set. If the pressing issue is whether or not insanity exists, it may be totally uninteresting to be told that Jones is not in fact guilty. If, in saying ‘The murderer of Smith is insane’, I communicated that someone is insane (Jones, the man straight ahead, etc.), and this is the only proposition we care about in our discourse, we

needn't care whether or not a presupposition evidentially unconnected to this presupposition is false.

The account of presupposition failure highlights a point of difference between Stalnaker's account of the referential/attributive distinction and that presented here. Yablo (2006, 172-173) complains that Stalnaker has no general account of non-catastrophic presupposition failure. Stalnaker (199, 41-43) tells us that terms are associated with semantic rules; 'you' with the rule that it denotes the addressee, and 'the murderer of Smith' with the rule that it denotes the unique murderer of Smith. The rule associated with the definite description is pragmatically ambiguous in that it can be used in two ways. It might itself feature in what is said, in which case we have an attributive usage, or it might determine a referent who then features in what is said, in which case we have a referential usage. In the referential case, the rule picks out the individual who is presupposed to be the unique murderer of Smith. If this is Jones, then the utterance expresses the proposition that Jones is insane, and this proposition can be true even if Jones is innocent: "The presupposition helps to determine the proposition expressed, but once that proposition is determined, it can stand alone."

Yablo's complaint is that the same explanation cannot extend to other cases of non-catastrophic presupposition failure. An example from Kripke (1977, 14),

"Two people see Smith in the distance and mistake him for Jones. They have a brief colloquy, "What is Jones doing?" "[Jones is] Raking the leaves." "Jones", in the common language of both, is a name of Jones; it *never* names Smith. Yet, in some sense, on this occasion, clearly both participants in this dialogue have referred to Smith, and the second participant has said something true about the man he referred to iff Smith was raking the leaves."

The second participant utters a sentence that is semantically associated with a false proposition (on the assumption that Jones doesn't just happen to be raking leaves as well) and yet they say something true if Smith, the figure in the distance, is raking the leaves. According to Yablo, there is no reading of 'Jones is raking the leaves' on which it is true if Smith is raking the leaves. The sentence does not admit of a pragmatic ambiguity that allows 'Jones' to pick out

Smith, so Stalnaker's previous explanation of non-catastrophic presupposition failure cannot apply here.⁹⁷

Whether or not Stalnaker's explanation of non-catastrophic presupposition failure can explain this case, the account offered here does so easily. Relative to the presupposition that the figure in the distance is Jones, the proposition that Jones is raking the leaves and the proposition that the figure in the distance is raking the leaves are locally equivalent. The speaker therefore communicates, not only that Jones is raking the leaves, but that the figure in the distance is raking the leaves. As Smith is in fact the figure in the distance, the second claim is true only if Smith is raking the leaves. If Jones is not raking the leaves, then the first claim is false, but given the right setup, the second claim may be more important for the purposes of the conversation, in which case the falsity of the first claim is unproblematic and presupposition failure need not be catastrophic. None of this is to say that the utterance wasn't false. One of the propositions communicated was that Jones is raking the leaves, and this is false. But that doesn't prevent the utterance also communicating a true proposition, and that may be sufficient for the overall success of the utterance.

Having now mentioned Yablo (2006), we can break to consider how the account of non-catastrophic presupposition failure offered here compares to the one offered in that paper.

Yablo on non-catastrophic presupposition failure

Yablo (2006, 175) sets himself the task of explaining why some utterances strike us as making claims despite the falsity of their presuppositions. Yablo follows Strawson in assuming that the truth value of a sentence is undefined when its presuppositions are false. 'The murderer of Smith is insane' is not, contrary to Russell's analysis, false if no one murdered Smith, or if he

⁹⁷ There may be a way of applying that explanation, although I don't find it very attractive. Suppose that the term 'Jones' is associated with the rule that the term denotes the unique individual called 'Jones' (or, perhaps, that the term 'Jones_n' is associated with the rule that the term denotes the unique individual called 'Jones_n', where n is an unpronounced qualifier that disambiguates between the names of different Joneses.) In Kripke's example, the pair presuppose that Smith – the actual guy in the distance to whom they are causally related by their perceptual systems – is called 'Jones'. This is false. Smith is called 'Smith' not 'Jones'. The false presupposition, however, allows the rule to determine the proposition that Smith is raking the leaves, which is true. Stalnaker (2006, 278–279) doubts whether all the cases Yablo discusses should be given a uniform treatment. Regardless, the possibility of a uniform treatment is a possibility for my account, even if it is not for Stalnaker's.

was murdered by a pair. In such a circumstance, according to Strawson (1950, 12) the question of its truth and falsity would simply “not arise”.

When a presupposition, π , is false, the proposition semantically associated with that sentence is neither true nor false. Why then do such sentences sometimes strike us as communicating truths or falsehoods? Yablo (2006, 175) answers that the proposition *semantically associated with the sentence* may be neither true or false, but the proposition *asserted* by the sentence may still have a truth value, “S’s semantic content-what in context it *means*-will be a proposition defined only on π -worlds; it is semantic content that determines S’s truth-value. Truth-value intuitions are driven not by what a sentence means, however, but by what it says, it’s asserted content.”

The proposition asserted by a sentence is the proposition that must be added to the sentence’s presupposition in order to secure the truth of the sentence, so $\pi, X \models S$, where π and S are the uttered sentence’s presuppositions and the proposition semantically associated with the sentence, respectively, and X is the sentence’s asserted content. We need to impose some restrictions on X beyond rendering the above argument valid. We won’t go through these restrictions in detail but one will be very important. According to Yablo (2006, 182) π & X should not imply anything stronger than S. We are looking for the asserted content of the sentence uttered, not of some stronger sentence. As S implies π (the sentence is true only if π), this constraint amounts to the requirement that S imply X.

Let’s apply this view to one of our cases. I say ‘The murderer of Smith is insane’ when dismayed by the condition of Smith’s corpse. ‘The murderer of Smith is insane’ presupposes that there is a unique murderer of Smith. If Smith was in fact murdered by a pair, then this presupposition is false and the uttered sentence lacks a truth value. Such presupposition failure needn’t be catastrophic, however. Although the proposition semantically associated with ‘The murderer of Smith is insane’ is undefined, its asserted content may still have a truth value and so the utterance may make an evaluable claim. To find the asserted content, we find the proposition that fills the gap in the argument:

There is a unique murderer of Smith

X

— — —

The murderer of Smith is insane

The proposition that the unique murderer of Smith is insane is ruled out by Yablo's requirement that X have no "trace" of π in it. Replacements for X that meet this requirement are the proposition that some murderer of Smith is insane, or the proposition that all murderers of Smith are insane. In such a case, Yablo prefers to conjoin all of the candidates, so the asserted content of 'The murderer of Smith is insane' is the proposition that some and all murderers of Smith are insane, which may have a truth value even when Smith was murdered by a group.

The view of this Thesis agrees. Relative to any context set representing a discourse in which it is presupposed that there is a unique murderer of Smith, the proposition that the murderer of Smith is insane (the semantic content of 'The murderer of Smith is insane') is locally equivalent to the proposition that some and all murderers of Smith are insane. The set of possibilities at which the murderer of Smith is insane and the set of possibilities at which some and all murderers of Smith are insane intersect with the possibilities at which there is a unique murderer of Smith to yield the same set, the set of possibilities at which there is a unique murderer of Smith and they (that individual who constitutes all of Smith's murderers) are insane.

As every context set representing a discourse state at which it is presupposed that there is a unique murderer of Smith is a subset of the set of possibilities at which there is a unique murderer of Smith, the proposition that some and all murderers of Smith are insane and the proposition that the murderer of Smith is insane are locally equivalent relative to any such context.⁹⁸ By uttering a sentence semantically associated with the proposition that the murderer

⁹⁸ If $p \cap r = q \cap r$, then $p \cap r^* = q \cap r^*$ for any r^* that is a subset of r . $p \cap r = q \cap r$ tells us that any member of r is a member of p if and only if it is a member of q . As all members of any subset of r are members of r , $p \cap r' = q \cap r'$ for any r' that is a subset of r . Just suppose that $p \cap r = q \cap r$ but $p \cap r' \neq q \cap r'$. In that case, there is a member of $p \cap r'$ that is not a member of $q \cap r'$ (or vice versa, but the proof is symmetrical either way). This rogue member must be a member of both p and of r' . Every element of r' is a member of r , so this rogue is a member of $p \cap r$. As the rogue is not a member of $q \cap r'$, but is a member of r' , it isn't a member of q , so it isn't a member of $q \cap r$. The rogue is a member of $p \cap r$ but not a member of $q \cap r$, so $p \cap r \neq q \cap r$.

of Smith is insane, in a context such that it is presupposed that there is a unique murderer of Smith, one communicates that some and all murderers of Smith are insane. As the proposition that there is a unique murderer of Smith is a semantic presupposition of the sentence ‘The murderer of Smith is insane’, it is in effect in any discourse in which the sentence is appropriately uttered, so any utterance of this sentence communicates the proposition that some and all murderers of Smith are insane, and that is a large part of what it means to be the asserted content of the sentence. So far, it seems that we are in agreement.

Now consider referential uses of definite descriptions, as when I say ‘The murderer of Smith is insane’ while observing Jones’s erratic behaviour. Donnellan tells us that referential uses presuppose that a particular individual uniquely satisfies the description, so Yablo takes it that the asserted content of the sentence is the proposition that fills the gap in the following argument,

Jones is the unique murderer of Smith⁹⁹

X

— — —

The murderer of Smith is insane

What we need is the proposition that Jones is insane, and this looks to be consistent with our earlier results. Saying one proposition communicates any locally equivalent proposition. Given the presupposition that Jones is the unique murderer of Smith, the proposition that the murderer of Smith is insane is locally equivalent to the proposition that Jones is insane.

$q \cap r$, contrary to our initial assumption, so $p \cap r = q \cap r$ only if $p \cap r' \neq q \cap r'$ for any r' that is a subset of r . Any context set representing a discourse state in which r is presupposed is a subset of r , so we have our conclusion.

This also runs the other way round. If p and q are locally equivalent relative to any context set representing a state of a discourse in which r is presupposed, then $p \cap r = q \cap r$. Just suppose that $p \cap r' = q \cap r'$ for every r' that is a subset of r , but $p \cap r \neq q \cap r$. Then there is a rogue member of $p \cap r$ that is not a member of $q \cap r$ (or vice versa, but the proof is symmetrical either way). As the rogue is a member of $p \cap r$, it is a member of p and r . As it is not a member of $q \cap r$ but is a member of r , it is not a member of q . The rogue is a member of p and of r , but not of q , so there is some subset of r , r' , that includes the rogue, and the rogue is a member of $p \cap r'$ but not a member of $q \cap r'$, in contradiction of our initial assumption. Two propositions intersect identically with a third if and only if they intersect identically with any context set representing a discourse state at which the third is presupposed.

⁹⁹ I assume this to be the semantic content of ‘That man is the murderer of Smith’, said while pointing at Jones. See Yablo (2006, 180).

The problem with Yablo's account is that identifying the proposition that Jones is insane as the *asserted content* of the sentence requires that definite descriptions be ambiguous or context-sensitive. The proposition semantically associated with a sentence must entail its asserted content. If the asserted content of the referentially used description is the proposition that Jones is insane, then any use of the same sentence entails this proposition. Attributive uses of 'The murderer of Smith is insane' need not entail that Jones is insane, so they must be semantically associated with a different proposition, despite their superficial similarity. An attributive use of the description that doesn't entail that Jones is insane must be semantically associated with a different proposition from the referential use.

Yablo's account delivers this unwelcome consequence because he treats the presupposition that Jones is the murderer of Smith on a par with the presupposition that there is a unique murderer of Smith, a semantic presupposition the truth of which is required for the truth of the sentence uttered. Instead, we can treat the presupposition that Jones is the murderer of Smith as a conversational presupposition, not required for the truth of the sentence. The semantic and asserted contents of 'The murderer of Smith is insane' are the same whether the sentence is used referentially or attributively, but the propositions communicated can vary with the speaker's conversational presuppositions, and need not be entailed by the truth of the sentence.

When there is no unique murderer of Smith, the proposition semantically associated with 'The murderer of Smith is insane' is false if we follow Russell, or undefined if we follow Strawson. Either way, it can still be used to communicate truths, relative to the presupposition that Jones is the unique murderer of Smith. We secure 2) and so have captured all of the key features of referential usage.

Loar on the semantics of singular terms

Before moving on to consider some more complex cases, I want to briefly discuss the account of the referential/attribution distinction offered in Loar (1976). The account bears some similarity to the account offered here, but there are important differences between them. For Loar, the distinction between referential and attributive uses of definite descriptions is that, in attributive cases, the description used is "connected in a certain special way" with the speaker's

intentions; the description the speaker actually uses perfectly captures what they mean. Suppose that it turns out that Smith was not murdered, but died of natural causes after a severe beating. If the description ‘The murderer of Smith’ was used attributively, the speaker will stick to their guns. What they meant was that the murderer of Smith is insane, although that turned out to be false:

“If one utters ‘the F is G’ attributively, then in response to any discovery of a mistake about the identity or existence of the F one would ideally judge, “Nevertheless, what I meant was that the F, whatever it may be, is G.” This is not to say that one would continue to maintain that it is true that the F is G, but merely that that was what one had meant.” Loar (1976, 359)

In contrast, a definite description used referentially fails to capture the speaker’s local intention. Loar argues that “on every normal referential utterance of ‘t is G’ there is some description ‘the H’ such that the speaker could have expressed what he had meant by saying ‘the H is G’, and that that utterance would have been attributive.” Had the speaker used this other expression, rather than the description they actually used, the case would have been attributive. As the speaker doesn’t use the description which best captures their local intentions, however, referential uses are characterised by rephrasing in response to presupposition failure.

Suppose that Jones did not murder Smith. When the speaker is told of this fact, they might respond that they didn’t really mean that the murderer of Smith is insane, but rather that the person standing in front of them is insane. When told that there is no one standing in front of them because they are looking into a mirror, they might respond that what they really meant was that the person they can see while looking straight ahead is insane. This process can be expected to end somewhere. If told that they cannot see anyone while looking straight ahead because they are hallucinating, then there may be no better way to rephrase what they meant. If so, then we have found the description that best captures what they meant, the person they can see while looking straight ahead is insane. This is false (because they are hallucinating) but it is what they meant nonetheless.

Loar’s arguments place a lot of weight on the hypothesised behaviour of speakers in the face of presupposition failure. These are not Loar’s only arguments but my aim here is to

clarify the similarities and differences between our accounts, rather than to decisively reject Loar's view. Speakers will, according to Loar, be willing to rephrase in order to capture what they meant. To a certain extent, I agree. By uttering 'The murderer of Smith is insane', when it is presupposed that the person standing in front of the speaker is the murderer of Smith, the speaker communicates that the person standing in front of them is insane. When it is recognised that the person standing before the speaker is innocent, the speaker may still be willing to communicate that the person before them is insane, and rephrase as 'The person standing in front of me is insane'.

While Loar takes this new utterance to capture what the speaker meant by their original utterance, I take it to capture only some of the elements of what was meant. The speaker may (given the presuppositions that remain in effect) communicate many of the propositions that they intended to communicate by their original utterance, but there are propositions which the original utterance was intended to communicate and the new utterance does not capture, the obvious example being the proposition that the murderer of Smith is insane.

Although we both take presupposition failure as important to a characterisation of referential uses of definite descriptions, Loar and I do not agree that there is some particular description (or even, as suggested at p. 361, some simple conjunction of descriptions) which most accurately captures what the speaker means. One of the features of referential usage is that the speaker communicates many propositions which are, so to speak, lumped together by the speaker's presuppositions. By communicating one proposition, the speaker communicates all the members of a set, and the speaker could use any sentence semantically associated with a member of the set to communicate every member. Whichever of these sentences the speaker chose to utter, the upshot would be the same and the utterance would be no less referential, given that the characteristic presuppositions associated with referential uses are in effect. Giving up on a presupposition can detach one member from the set, in which case the speaker may rephrase their original utterance in order to retain the other members, but this new utterance will not perfectly capture the speaker's original intentions.

Before closing, I want now to briefly consider some more complex cases which provide the opportunity to expand the view already outlined. These cases will put some pressure on Loar's view. On Loar's view, it turns out, it is not clear that there can ever be

attributive uses of definite descriptions. According to my view, whether descriptions are used attributively or referentially is a matter of degree.

Developing the distinction

Having seen how Donnellan's distinction is consistent with Russell's semantics for definite descriptions, I want to consider some 'mixed' cases in which it is unclear whether the utterance satisfies the features of referential usage. Suppose that the speaker and their audience are at a party. They recognise all of the guests except for one person who is the only one wearing a mask. They speak to this person for hours and discuss many topics, but never discover their identity. By the end of the evening, two possibilities remain: the person in the mask is either Smith or Jones.

The speaker says 'The person in the mask is fun', and uses the description to enable their audience to pick out a particular individual – the person in the mask who they spoke to for hours about this and that – and they could just as easily have used another description to do so. From Donnellan's initial characterisation, the case looks more like a referential use than an attributive one. The speaker doesn't know the person's given name, so cannot pick them out by a proper name without coining one, but could certainly have picked them out with a demonstrative, a pronoun, or another description, such as 'The person we spoke to for hours'. According to 1), the case appears referential. Secondly, the speaker's utterance might have successfully predicated a property of some particular individual, even if nobody satisfied the description. If it turns out that no one at the party was wearing a mask, but the person they spoke to all evening had a facial tattoo that was mistaken for a mask, then still the speaker may have successfully picked out a particular individual and predicated a property of them, as required by 2).

These properties are again explained in terms of local equivalence. The speaker and their audience presuppose that the individual they spoke to for hours, etc. is wearing a mask. All possibilities in the context set are therefore such that there is an individual wearing a mask who the pair spoke to for hours and so on for all of the properties they presuppose this person to instantiate. Whether they use the description 'The person wearing a mask' or 'The person we spoke to for hours', they eliminate the same possibilities from the context set. For the same

reason, the utterance might be successful even if the description fails to fit the intended referent. If the person they spoke to all evening had a facial tattoo instead of a mask, the speaker still communicated the proposition that the person they spoke to for hours was fun, and this is true. If the false presupposition is recognised, the audience can go through the kind of reinterpretation that we noted previously by interpreting based on the many propositions communicated other than that semantically associated with the sentence actually used.

Although usage of ‘The person in the mask’ evidenced features 1-2 of referential use, it does not clearly exhibit feature 3. According to 3, a referential use of ‘The person in the mask’ requires that the speaker presuppose, of a particular individual, that they uniquely satisfy the description. If the speaker makes this presupposition, then there is a particular individual who uniquely satisfies the description at all worlds in the context set. This person must either be identical with Smith, or not identical with Smith. Either option threatens to severely complicate what seem to be totally mundane assertions.

Suppose that Smith is the unique person wearing a mask at all possibilities in the context set. This conflicts with the stipulation that, for all the evidence of the speaker and their audience, the person in the mask might be Smith or Jones, and threatens to complicate the analysis of both ‘Smith is the person in the mask’ and ‘Smith is not the person in the mask’. Given the context, both of these utterances should be informative and sensible contributions to the discussion, but if Smith is the unique individual wearing a mask at all possibilities in the context set, then the first utterance is entirely trivial, eliminating no possibilities from the context set. The second utterance is at the other extreme, eliminating all possibilities from the context set and rendering continuation of the discourse impossible. The results are reversed for these sentences if it is presupposed that Smith is not the person in the mask.

This case provides an interesting opportunity to look a bit more closely at condition 3),

“when a definite description is used referentially ... the speaker presupposes of some particular someone or something that he or it fits the description.”

Although it does not seem in the mask case that the speaker presupposes *of Smith* or *of Jones* that either fits the description used, there is clearly a sense in which they do presuppose of some particular someone that they fit the description: the person in the mask who they spoke to for hours about this, that, and the other. I want to suggest that this case is no different in

kind from the referential cases that we considered previously. If there is a difference, then it is a difference only of degree.

One merely apparent difference between the cases is that in the case of the mask, the speaker does not presuppose enough to settle the identity of their intended referent. Even when we refer to someone regularly by a proper name, it is very often consistent with our presuppositions that they turn out to be identical with the referent of another proper name. We had no trouble in our original referential case assuming that the speaker and their audience presupposed of Jones that they were the murderer of Smith, but this may be entirely consistent with the later discovery that Jones and Atkinson (about whom the speaker may presuppose much) are identical.

The use of a proper name in describing the original example naturally leads us to assume that the speaker presupposes as much of their intended referent as we usually presuppose of those with whom we are personally acquainted, while the stipulation that the speaker in the mask case is uncertain as to which proper name they should apply to their intended referent naturally leads us to suppose that they presuppose very little. I have no doubt that there is nothing essential about the connection between referential usage and proper names, however. We can introduce a proper name for an individual that we know almost nothing about, such as 'Jack the Ripper' and should be comfortable saying that the police presuppose this and that *of Jack the Ripper* and refer to him by saying 'The murderer', even though there is a clear sense in which the police have failed to presuppose the identity of Jack the Ripper; it is consistent with what they presuppose that Jack the Ripper might turn out to be Smith or Jones. If we want to distinguish between these cases, we need to look deeper than proper names.

The suggestion I want to make, although I cannot argue for it in any detail, is that presupposing something *of someone* is a matter of degree. It can make sense to say that the police presuppose something *of Jack the Ripper* even when they presuppose relatively little of whoever the murderer was, and presuppose far less than is required to represent their discourse by a set of possibilities all of which include the very same individual. For all we said in our first example, it may be consistent with everything that the speaker presupposes to learn that Jones is identical with someone else that they presuppose more about. This is not only consistent with the letter, but also the spirit, of the case.

Because this notion of presupposing *of someone* features in Donnellan's characterisation of the referential/attributive distinction, the distinction is itself a matter of degree if this notion is; a spectrum with the most clearly attributive cases at one end and the most clearly referential at the other. Not only is the third condition a matter of degree. As we explained the first two characteristics in terms of the third, and satisfying the third is a matter of degree, we should expect to find that the first two are also a matter of degree. This is what we in fact find. The more interesting information is presupposed to hold of the intended referent, the more interesting propositions are locally equivalent to whatever was said, and so the more replacements for the description actually used will preserve the truth value of the utterance and the more propositions will survive presupposition failure.

Eventually, when the descriptions associated with the referent reach a sufficient number,¹⁰⁰ we become comfortable describing the speaker's presuppositions using proper names – of saying, for example, that they presuppose *of Jones* that he is the murderer – and we reach the level of referentiality present in Donnellan's original example, which is suggested by the use of a proper name.

Our final account of the referential/attributive distinction is as two ends of a spectrum that includes all uses of definite descriptions, organised by the amount (and perhaps the type) of information that is presupposed of whoever satisfies the definite description. If this is the case, we should be able to find cases that are more on the attributive side but that share the features of referential usage to some degree. Again, this is what we in fact find. Reconsider attributive usage in the case of Smith's murder. We come across Smith's body and I say 'The murderer of Smith is insane'. As the case was originally described, we have no suspects, but this is consistent with our having some knowledge of the murderer. In fact, it is hard to imagine how one could make such an utterance without presupposing a good deal about the murderer, that whoever murdered Smith was responsible for the condition of Smith's body, for example. Even in this case, therefore, we see some of the features that Donnellan uses to characterise referential usage, although to a lesser degree. What is communicated by the speaker's utterance is unchanged, for example, by replacing the description 'The murderer of Smith' with 'The person responsible for the condition of this body'. Because of this, the

¹⁰⁰ Or perhaps when they include descriptions of the right type. See Loar (1976, 362-365) for a brief examination of what he terms *identifying* descriptions.

speaker communicates, not only that the murderer of Smith is insane, but that the person responsible for the condition of Smith's body is insane, and this fact could aid reinterpretation if we discover the Smith happened to die of natural causes after a brutal beating.¹⁰¹

Such cases pose a problem for the account of attributive usage posed by Loar. According to that account, an attributive utterance is one in which the definite description actually used is the one that best captures the speaker's local intentions. The test for this is to consider what would happen in a case of presupposition failure. If, in response to every potential failure of a presupposition, the speaker would respond that, in any case, what they meant was captured by their original utterance, then the utterance was attributive. In the above case, the speaker might respond to the realisation that Smith was not murdered by saying "Well, what I really meant was that the person responsible for the condition of Smith's body is insane" so the description 'The murderer of Smith' was not used attributively. It seems likely that the same holds for any use of a definite description and therefore that there are no attributive uses of definite descriptions, according to Loar's account. The speaker will always presuppose something about their intended referent, and it is very plausible that for some presupposition failure, more of what the speaker intended to communicate will be retained by using a definite description other than that they actually used. At best, attributive uses of definite descriptions are, according to Loar's account, much rarer than we initially thought and, at worst, impossible.

To press the point, consider just one outlandish case.¹⁰² I am told that you were bitten by a cat and, convinced that owned cats do not bite; I say to you 'The cat that bit you was a stray'. Feel free to alter the example to make it as clearly attributive as you can, while retaining 'The cat'. Cats are necessarily (let us assume) animals. You tell me, correctly, that there are in fact no cats. 50 years ago, you tell me; all cats were removed from Earth and replaced with highly sophisticated machines. When apprised of this shocking fact, I might very reasonably respond "Wow! Well I guess what I meant was that the cat-like thing that bit you was a stray." According to Loar's test, the original utterance was not an attributive use of 'The cat that bit you', but a referential one. What I really meant by the utterance was that the cat-like thing that

¹⁰¹ Donnellan (1968, 209) considers this case a "near miss".

¹⁰² Inspired by Putnam (1967, 660). Note that this is one of Putnam's more moderate examples. As the cats were replaced with robots only recently, and we were previously using the word 'cats' to refer to a type of animal, 'The cat that bit you was a stray' must be false if you were bitten by a robot.

bit you was a stray. I take it that similar examples of outlandish presupposition failure could be constructed for just about any use of a definite description. I further take it that the non-existence (or close to it) of attributive uses of definite descriptions is at least a *prima facie* cost of Loar's analysis.

Chapter 6: Objections and Replies

Can we retain traditional semantics?

I have suggested that the alternative model of communication can avoid the problem of underdetermination and so preserve a broader domain of viable linguistic theory than the Encoding Model. It is part of this alternative model, however, that nothing need be said by an utterance for communication to succeed. In what way, then can the alternative model preserve linguistic theories that associate sentences with propositions said?

The answer is that these linguistic theories may well be true. It may be true, for example, that an utterance of ‘The guy is drunk’ says that the guy who authored *Smells and Ticks* is drunk, given a context in which the hidden-indexical refers to the property of *being the author of Smells and Ticks*. What our alternative model denies is that the contexts in which we actually talk need to determine what is said for communication to succeed. Nor are these linguistic theories redundant in the explanation of communication, as it is in part the conventions of the language that determine the candidates for what is uttered and said.

While we can allow for the possibility that these theories make true predictions about what is uttered or said in certain context, the alternative model also allows for another way to think about linguistic theory. The Encoding Model was conducive to a certain realism about linguistics that, on the alternative model, is a red herring. According to the Encoding Model, we utter sentences and say propositions. The question then is what sentence is uttered and what proposition is said in different cases. The hidden-indexical theory of quantifier domain restriction suggests that the sentence uttered by an articulation of ‘Every bottle is empty’ is ‘Every bottle *i* is empty’, while according to the syntactic ellipsis view, the sentence uttered may be ‘Every bottle in the cupboard is empty’. Thinking in terms of the Encoding Model, these theories make inconsistent claims about the world, in which case at most one can be correct.

According to the alternative model, however, nothing need be uttered or said. Rather than trying to identify the nature of some object in the world – the sentence uttered – these theories should be viewed as identifying different interpretive structures, either of which might be allowable in any given context. The very same result can be accomplished by positing and

resolving syntactic ellipsis or by positing and resolving hidden-indexicality in any particular case. Looking at things this way, these theories raise different interpretive possibilities but make no inconsistent claims.

Revenge problem 1: The underdetermination of what is meant

An initial and instructive revenge problem has it that denying the Utterance and Expression Constraints fails to address the core of the problem, as all the arguments from Chapter 2 can be recast as direct arguments for the indeterminacy of what is meant, rather than by way of the indeterminacy of what is said or uttered.

The point is well made via Buchanan's (2010) argument that speakers needn't mean propositions by their utterances. For the sake of the argument, Buchanan assumes

Content: What a speaker means, or intends to communicate (at least in cases of indicative speech) must be a proposition.

Success: Understanding a speaker's utterance U requires (minimally) entertaining what she meant by U.

The Encoding Model assumed Content, that is, that propositions are the objects of speakers' communicative intentions. Buchanan's (2010, 344-345) intended reading of Content agrees with the Encoding Model that speakers might mean multiple propositions by a single utterance. Regarding Buchanan's Success condition, it is not entirely clear as to what he means by *entertaining* a proposition. We will return to this issue later. In the meantime, we can accept Success as equivalent to our running assumption that communication succeeds only when the audience knows what the speaker intends to communicate.

From Content and Success follows

Lemma: If a speaker means a proposition p by her utterance U then her audience must entertain p if she is to understand U.

Buchanan argues that Lemma is false and lays the blame with Content. Speakers needn't intend to communicate propositions. His central counterexample to Lemma is the case of Chet and Tim, who are preparing for a party. They buy some bottled beer, which Chet leaves to chill in

the ice-filled bucket by their hot tub. Before the party starts, a parched Tim asks Chet where the beer is. Chet replies, ‘Every beer is in the bucket’. Immediately upon hearing Chet’s utterance, Tim goes to the bucket and gets himself a beer.

A few plausible stipulations. Firstly, Chet means something by his utterance. Secondly, Tim is in a position to understand Chet’s utterance. Thirdly, Chet means at least one proposition of a form determined by the sentence he uttered. Buchanan represents the meaning of the sentence Chet utters by the following template

(TEMP) [The y : Bucket(y) & $_y$] ([Every x : Beer(x) & $_x$] (x is in y))

The role of TEMP is to restrict the propositions that can be meant by the literal speaker who utters this sentence.¹⁰³ TEMP represents a proposition type, the tokens of which are all those propositions that can be derived by replacing the underscores, which function as variables, with properties. So the proposition that every beer in the apartment is in the bucket filled with ice, and the proposition that every beer from the smallest bodega in Cuba is in the bucket which belongs to Condoleezza Rice, are both of the type represented by TEMP. As with other representations of the meanings of context-sensitive sentences that we have encountered, TEMP is a function from the contextually-assigned values of the variables to a proposition. Given that Chet is speaking literally, at least one of the propositions that Chet meant by his utterance is a substitution instance of TEMP.

Chet means something by his utterance. By Lemma, there is some proposition p such that Tim must entertain p in order to understand Chet’s utterance. As Chet is speaking literally, there is some proposition p of the form specified by TEMP such that Tim must entertain p in order to understand Chet’s utterance. Buchanan argues that there is no such proposition. In fact there are a number of propositions of the right form that Tim could entertain in order to understand Chet’s utterance, but none that Tim must entertain to understand Chet’s utterance, so none that satisfies the condition specified by Lemma; none, that is, that Chet means.

¹⁰³ Buchanan (2010, 349) assumes that all candidates for what Chet meant are candidates for what he said, although his usage of the term ‘said’ is not quite the same as the usage we have established. Buchanan (2010, 345) takes there to be “a constitutive connection between what a speaker means and what she says: if S says p in uttering u , then at least one of the propositions that S means by uttering u must be P .” Given our slightly different terminology, I will try to avoid talk of ‘what is said’ in discussing Buchanan’s argument.

We have a situation very similar to that of the underdetermination arguments from Chapter 2. We assume that the speaker means something and that the audience can know what the speaker means. Given that the speaker has uttered a particular sentence and spoke literally, at least one of the propositions the speaker meant must have been of a certain form; in this case, the form specified by TEMP. Whichever proposition this is, it must satisfy Lemma, that is, the audience must entertain this proposition if they are to understand the speaker's utterance.

Buchanan identifies a number of candidates but argues that none of them satisfies Lemma. For initial simplicity, let's assume that there are only two candidates for the proposition that Chet means: the proposition that every beer in the apartment is in the bucket filled with ice (P_1) and the proposition that every beer for the party is in the bucket by the hot tub (P_2). Both of these propositions are of the form required by Lemma and are very reasonable propositions to communicate, given Tim's thirst and the knowledge that Tim and Chet hold in common about the history of the relevant beers and bucket.

The problem, however, is that neither of these propositions satisfies Lemma and so neither can be the proposition that Chet meant. Suppose that Chet means P_1 . By Lemma, if Chet means P_1 , then Tim must entertain P_1 in order to understand Chet's utterance. The consequent is false, however. Tim needn't entertain P_1 in order to understand Chet's utterance, as Tim could understand equally well by entertaining P_2 . By *modus tollens*, therefore, Chet could not have meant P_1 . Symmetrical reasoning demonstrates that Chet could not have meant P_2 . As the only propositions that Chet could have meant were P_1 and P_2 , Chet could not have meant any proposition by his utterance. Although in any realistic situation, there will be a great many more than two candidates for the proposition meant, Buchanan's contention is that no proposition will be such that the audience must entertain that proposition to understand the speaker's utterance.

Response

Buchanan (2010, 348) takes his argument to establish that there is no proposition that Chet could have meant by his utterance. As Chet clearly meant something by his utterance, we must reject Content: what a speaker means needn't be a proposition. According to Buchanan (2010, 358-359) speakers only mean to communicate *types* or *properties* of propositions. In the case of

Tim and Chet, there is no proposition of the form specified by TEMP that Chet means to communicate. What Chet means to communicate is, rather, TEMP itself, which specifies a proposition type or property. To accommodate this alternative notion of what the speaker intends to communicate, Buchanan modifies Content to allow that speakers may mean types or properties of propositions and modifies Success to allow that understanding such a speaker requires only entertaining some one or more propositions of the type the speaker meant.

This picture is not quite accurate, however. In the case of Chet and Tim, Chet will not be satisfied if Tim entertains just any proposition of the form specified by TEMP. The proposition that every beer from the smallest bodega in Cuba is in the bucket which belongs to Condoleezza Rice is of the form specified by TEMP, yet Tim clearly fails to understand Chet by entertaining this proposition. Rather, Buchanan's suggestion is that Chet means a narrower type of proposition, determined by TEMP and vague contextual restrictions that exclude the proposition that every beer from the smallest bodega in Cuba is in the bucket which belongs to Condoleezza Rice.

It is worth considering a point that Buchanan himself is less than entirely clear on. Is Buchanan's argument designed to make a merely contingent point about speakers' indifference to communicating truth conditions, or is his claim that the very limited common knowledge of sort shared by Chet and Tim renders it impossible for Chet to mean any proposition? If the former, then I see little reason to dispute Buchanan's position. Very plausibly, speakers' communicative intentions often exhibit the sort of indifference that Buchanan cites but I don't see any reason to use the language of 'meaning' and 'understanding' to explain the case, or to revise these notions to accommodate the case.¹⁰⁴ Speakers can articulate sentences for a number of different reasons. One is the communication of propositions, but this is one among many. A speaker needn't really intend to communicate anything by articulation of a sentence, even a declarative. Speakers often intend to bring about some other effect in the world. Sometimes this might be the procurement of beer. Sometimes the speaker's intention might be to direct their audience to beer. On one interpretation of the case of Chet and Tim, the sum total of Chet's intention is just to direct Tim to the beer and Chet will be satisfied so long as his utterance leads Tim to procure beer. On this interpretation, Chet's intention wasn't to

¹⁰⁴ See von Stechow and Gilles (2011) for the suggestion that epistemic modals are regularly used without the intention to communicate. No proposition is communicated, but several are "put into play".

direct Tim to the beer by communicating any information about it, even though the case, on the face of it, might appear as one of informative communication.

A stronger claim is that it is impossible for Chet to reasonably mean any proposition at all, given the very limited knowledge that Chet and Tim hold in common. This is the claim disputed by the alternative model of saying and communicating presented in this Thesis. The virtue Buchanan (2012, 359) sees in his proposal is that it does “justice to the generality and indifference characteristic of the speaker’s communicative intentions”. According to any such account, however, speakers do not intend to communicate truths and falsehoods. Proposition types, for example, even restricted by context in the way Buchanan suggests, fail to determine truth-conditions. If the speaker means only a proposition type, then there is nothing true or false that the speaker intends to communicate, only a (potentially infinite) number of different truth-conditions that are compatible with the speaker’s intention. Contrary to Buchanan (2010, 361), proposition types just don’t seem to be “the right *kind* of thing to be the object of the speaker’s communicative intentions” if we want those intentions to determine truths and falsehoods.

We can agree with Buchanan that common knowledge leads to indifference on the part of speakers, while restricting this indifference to what is said and allowing that what the speaker intends to communicate is entirely determinate. The alternative account of the case should come as no surprise. What the speaker intends to communicate is not any one of the candidates to the exclusion of the others, but each and every one of the candidates. In our simplified example, there were only two candidates for what Chet meant: the proposition that every beer in the apartment is in the bucket filled with ice (P_1) and the proposition that every beer for the party is in the bucket by the hot tub (P_2). Part of the reason that these are equally good candidates for what is said is that they are locally equivalent, given the obvious presuppositions that the beer for the party is the beer in the apartment and that the bucket filled with ice is the bucket by the hot tub. Given these presuppositions, every beer in the apartment is in the bucket filled with ice if and only if every beer for the party is in the bucket by the hot tub. If Chet communicates either of these propositions, relative to these presuppositions, then Chet communicates them both. If Chet means either, therefore, he should mean them both.

Buchanan (2010, 353-354) considers and rejects what he considers three distinct responses to the puzzle. According to the first, Chet means the conjunction of each of the candidates. According to the second, Chet means each of the candidates. According to the third, Chet means the disjunction of all of the candidates. These options are not mutually exclusive, however. According to the response proposed here, what Chet means can be represented equally well by each of the candidates for what is said, by the conjunction of all those candidates, and by the disjunction of all the candidates. As long as Chet means at least one of the candidates, he means them all, given that they are locally equivalent in the context. Buchanan's objection is that this "places the requirement for understanding Chet's utterance far too high." According to Lemma, if Chet means each of these propositions, then Tim must entertain them all in order to understand Chet's utterance. As Buchanan has already argued, however, Tim could understand Chet's utterance equally well by entertaining any one of the candidates in isolation.

Buchanan doesn't tell us much about what *entertaining* a proposition amounts to, but this objection make it clear that he takes it that Tim has to do some extra work to entertain two propositions rather than one. This notion of entertainment overly intellectualises successful communication by presenting its success condition on the model of silent speech to oneself. In the sense relevant to defining successful communication, it needn't be any more difficult for Tim to entertain a multitude of propositions than it is for Tim to entertain a single proposition. A plausible gloss on Success, and the one that we have assumed throughout this Thesis, is that the audience understands the speaker only if, for every proposition *p* that the speaker intends to communicate, they know that the speaker intended to communicate *p*. When understanding is presented on this model, it is far less clear that the speaker has to do any more work to understand a speaker who means two propositions than they do to understand a speaker who means a single proposition.

It is a virtue of the representation of local equivalence in terms of Stalnaker's framework is that it very conspicuously represents this fact. Understanding a speaker comes down to knowing the update that they intend to impose on the context set and there is no reason to think it is more difficult to grasp an update that introduces two new truths to the context set than it is to grasp an update that introduces a single new truth. There are alternative representations of context in which this is less conspicuous. Suppose, for example, that the

context is represented just by a set of presuppositions. In cases devoid of sarcasm and the like, we are told to update this set with what is said, and then with any presuppositions entailed by what is said and the member of the original set. In an unrealistically simple version of Buchanan's case, the context is represented as {<every beer in the apartment is in the bucket filled with ice if and only if every beer for the party is in the bucket by the hot tub>}. If the speaker says that every beer in the apartment is in the bucket filled with ice, we have to add this to the set, to derive {<every beer in the apartment is in the bucket filled with ice if and only if every beer for the party is in the bucket by the hot tub>, <every beer in the apartment is in the bucket filled with ice>} before finally adding the proposition entailed by these two, to derive {<every beer in the apartment is in the bucket filled with ice if and only if every beer for the party is in the bucket by the hot tub>, <every beer in the apartment is in the bucket filled with ice>, <every beer for the party is in the bucket by the hot tub>}. Such a representation of context is entirely workable but may give the false impression that the audience has to perform some additional activity of identifying the entailments of the members of the set. Using Stalnaker's notion of the context set, we can see how two propositions are added to what is presupposed by a single update operation.

Another problem that Buchanan (2012, 354) poses for these responses is that it is implausible to suppose "that speakers have meaning-intentions towards *each* of the ... candidate propositions". As Sperber and Wilson (1986, 58) note, however, while a speaker "must have a representation of the set of assumptions ... which she intends to make manifest ... it is not necessary to have a representation of each assumption in the set." While Buchanan (2012, footnote 30) presents this quote in another context, he fails to recognise the significance of the quote to this objection. Given the context, the speaker might, for example, represent the assumptions they intend to communicate by explicitly representing any one of the candidates for what is said which, in collaboration with the context set, determines everything that they mean to communicate. As Buchanan (2012, 350) notes, if asked to be more explicit as to what he meant, "Chet might volunteer *any* one (or more) of the various candidates" for what is said. Buchanan takes this to suggest that "no single such candidate, or set of candidates, perfectly capture [sic] his communicative intentions". On the contrary, I take it to suggest that *any* of the candidates would perfectly capture Chet's meaning, relative to his presuppositions.

Revenge problem 2: The underdetermination of what is presupposed

According to the view presented in this Thesis, what the speaker intends to communicate is never determined simply by what they say, but by what they say in combination with their background presuppositions. Knowing what the speaker says is neither necessary nor sufficient for knowing what they intend to communicate. It is not sufficient, because the speaker might mean different things by what they said, depending on their background presuppositions. If I presuppose that the author of *Smells and Tickle*s is the man that we are waiting to hear, then I communicate that the man we are waiting to hear is drunk by saying that the author of *Smells and Tickle*s is drunk. If I presuppose that the author of *Smells and Tickle*s is the man currently staggering up to the podium, but not that he is the man we are waiting to hear, then by saying that the author of *Smells and Tickle*s is drunk, I communicate that the man currently staggering up to the podium is drunk but not that the man we are waiting to hear is drunk. It is not necessary, because the same thing might be communicated by all candidates for what is said, relative to the speaker's background presuppositions. If you don't know whether I say that the author of *Smells and Tickle*s is drunk or that the man we are waiting to hear is drunk, but know I presuppose that the author of *Smells and Tickle*s is the man that we are waiting to hear, then you can know what I intend to communicate: that the author of *Smells and Tickle*s is drunk and that the man we are waiting to hear is drunk.

The second revenge problem suggests that audiences cannot know what speakers intend to communicate if what they intend to communicate is determined in part by their presuppositions, because audiences cannot know what speakers presuppose. Return to the case of Chet and Tim. Suppose that one of the propositions that Chet intends to communicate is that every beer in the apartment is in the bucket filled with ice. If Chet presupposes that every beer in the apartment is in the bucket filled with ice if and only if every beer for the party is in the bucket by the hot tub, then Chet also means that every beer for the party is in the bucket by the hot tub. If Tim cannot know whether or not Chet makes this presupposition, then he cannot know whether or not Chet intends to communicate the corresponding proposition.

Response

If Tim cannot know what Chet presupposes, then Tim cannot perfectly understand Chet's utterance. Note, firstly, that this problem is significantly different from the original problem of underdetermination. The original problem was that it seemed impossible for audiences to know what speakers intend to communicate in the face of *shared* presuppositions. The new problem is rather that it seems impossible for audiences to know what speakers intend to communicate in the face of *divergent* presuppositions. This position is hardly unique the view under discussion. Imperfect communication is just what we should expect in the face of divergent presuppositions. Stalnaker (1999, 85) notes, for example, that contexts in which speakers and audiences make different presuppositions "will have a kind of instability" in that they are likely to lead to failures of communication, giving everyone involved a motivation to coordinate their presuppositions.

There is another, more direct, argument to suggest that that presupposition coordination is unproblematic in these cases. According to Stalnaker (1999, 84) "A proposition is presupposed if the speaker is disposed to act as if he assumes or believes that the proposition is true, and as if he assumes or believes that his audience assumes or believes that it is true as well." Presupposition is a dispositional notion, and various authors have supported the problem of underdetermination by pointing out that we are disposed to treat various different candidates on a par. Wettstein (1981, 247), for example, points out that a speaker, asked which of the candidates they said, might reply " 'Although I meant to refer to that table' ... 'I don't think I meant to refer to it *as* the table in room 209 of Camden Hall at t_1 as opposed to, say, *as* the table at which the author of *The Persistence of Objects* is sitting at t_1 . Nor did I intend to refer to it as the table in 209 and the table at which the author ... as opposed to, say, just as the table in 209.' " Here, Wettstein draws our attention to the speaker's disposition to treat the candidates for what is said on a par, that is the disposition to act as though it is assumed that the table in room 209 of Camden Hall at t_1 is the table at which the author of *The Persistence of Objects* is sitting at t_1 . While this supports the conclusion that what is said is underdetermined, it also supports the conclusion that the speaker presupposes the identity required to render the candidates for what is said locally equivalent, whether or not the speaker has explicitly considered this identity in advance.

Buchanan and Ostertag (2005, 900-901)¹⁰⁵ make a very similar point, arguing that the speaker is equally disposed to fall back on and “unhesitatingly offer” any one of the candidates when asked to be more specific as to what they meant. While this supports the conclusion that what is said is underdetermined, however, it also supports the conclusion that the speaker presupposes the local equivalence of the candidates. Buchanan and Ostertag take this disposition to indicate that “no one response perfectly captures [the speaker’s] intention”, while I take it to indicate that any one of these responses perfectly captures the speaker’s intention, relative to what they presuppose.

We have a motivation to coordinate on presuppositions, and some evidence to suggest that we have the dispositions relevant to such coordination, but we might still doubt whether coordination is regularly achieved, given that speakers and their audiences will often enter a conversation with divergent beliefs and points of view. Setting aside our beliefs for the purpose of conversation is a difficult thing to do, one that takes practice, and one that people fail at regularly. Even if conversations tend towards a state of presuppositional coordination, we may have reason to doubt whether it is ever achieved and so whether communication ever actually succeeds, according to the view presented here.

There are several things to bear in mind here. One thing is that the speaker’s driving intention behind their utterance may be achieved, even when the information I take you to convey is not quite what you intended. The communication of propositions is rarely, if ever, an end in itself. Rather, speakers communicate propositions in order to achieve further aims. If your purpose in speaking was to get me to disapprove of Pergola, this may be achieved whether or not I grasp the precise propositions that you intended to convey, and communication may be deemed a ‘success’. As Yalcin (2014, 24) notes, “even allowing that linguistic communication does centrally involve some degree of coordination on items of content at some interesting level of abstraction, it remains open that coordination on items of content is a highly approximate, more-or-less affair, with perfect coordination on content not being especially important, and rarely or never happening.” Practical considerations affect the level of coordination we desire from communication. Speakers can have a number of different goals in making an articulation and sometimes the speaker may care about communicating

¹⁰⁵ See also, Buchanan (2010, 350).

certain information only insofar as it will further some other goal, such as leading the audience to disapprove of Pergola, or leading Tim to acquire beer.

A further point to keep in mind is that the information that a speaker intends to communicate may be determined relative to the audience's presuppositions, rather than the speaker's. One particularly clear example of this phenomenon is captured by the cases from Egan (2009). Suppose that you write a billboard that reads 'Jesus loves you'. Egan (2009, 262) offers compelling evidence that what is said varies between interpreters. So if I read the sign, it says that Jesus loves *me*. If you read the sign, it says that Jesus loves *you*. This is opposed to the view, for example, on which the sign says the same thing to everyone, perhaps that Jesus loves everyone who reads the billboard.

A natural explanation of this variation of what is said is that we take 'you' to refer to whoever we presuppose to be reading the sign. Whoever reads the sign will presuppose that they are the reader, allowing them to identify a particular proposition that the speaker intends to communicate. The sign-maker might be incapable of entertaining presuppositions about the eventual interpreter, however, given that they needn't have any acquaintance with them. Despite failing to have any concept of the eventual reader, there is a particular proposition that the speaker intends to communicate. If the reader is A, for example, then the sign-maker intends to communicate that Jesus loves A – no other proposition will do.

Arguably, the sign-maker and the reader fail to coordinate their presuppositions, the reader making a presupposition that the sign-maker is incapable of making, yet communication may be entirely successful because the proposition the speaker intends to communicate is not fully determined by the speaker's own presuppositions, but by the presuppositions of whoever reads their sign.¹⁰⁶ If the response is plausible in this case, then it should extend to any case in

¹⁰⁶ We might try to say that the speaker's presuppositions are parasitic on the audience's presuppositions, so they manage to coordinate their presuppositions. This is problematic, however, given that the sign may be read by any number of different audiences with different presuppositions. If the speaker takes on the presuppositions of each of their audiences, then they end up with the inconsistent presuppositions that the unique interpreter of the sign is A and that the unique interpreter of the sign is B. We might try to suggest that the speaker avoids inconsistency by making these different presuppositions at different times, but there needn't be any change in the speaker's brain state, depending on whether A is reading or B is reading. While the speaker's brain might not determine their presuppositions, we might expect it to be a change in brain state. There is no parallel problem for intention. The speaker intends to communicate a certain proposition to A, and another proposition to B. As there is no contradiction between these intentions, both can be held simultaneously.

which the speaker has some reason to believe that their presuppositions differ from those of their audience.

Of course, speakers might not be happy with just any presuppositions that the audience has. There might be many presuppositions that would lead the speaker to disown the content that the audience takes them to intend and to conclude that communication had failed. We might see this, for example, in the case of the German spy who reads ‘Kitchener needs you!’ on the famous signs from the First World War (in which Kitchener’s visage appears in place of the term ‘Kitchener’). The maker of the sign might generally intend to communicate, to each reader x , that Kitchener needs x , while having no intention to communicate anything to a German spy who happens to read it.

I take the arguments from Chapter 4 to show that it is possible for speakers to communicate propositions in the face of underdetermination of what is said, but whether they actually do is another question. To some extent, the answer will depend on whether speakers and their audiences are able to coordinate their presuppositions, and the extent to which they care to coordinate. I don’t think either question has an easy answer. It may well be that, as a matter of contingent fact, speakers and their audiences never successfully coordinate their presuppositions and so that audiences never perfectly understand speakers. It may well be that speakers are not generally interested in perfect coordination on propositional content. These are not, however, problematic possibilities that arise in the context of the theory presented here, but are possibilities that we should be open to from the outset.

Even if communication is regularly imperfect, by the lights of this theory, the theory itself can retain all of its importance. The role of a theory of communication is not simply to describe communication as it actually happens (and certainly not to assume that actual communication is often, or ever, perfectly successful) but to offer ways of improving what are certainly imperfect communicative practices. The theory offered here has a lot to offer here. In particular, it cautions that divergent presuppositions pose a potential barrier to perfect coordination on content. To the extent that we care about coordination on content, then, we should attempt to coordinate our presuppositions. The extent that we care about coordination on content should be expected to vary between different discourses.

Pluralism about what is said

Having presented a certain view of interpretation and communication, I want now to clearly distinguish it from another response to underdetermination arguments. According to this alternative response, we should respond to underdetermination arguments by modifying our view of what is said. According to one version of the response, speakers say sets of propositions, rather than individual propositions, although the limiting case is that in which the speaker intends to communicate a set containing a single proposition. According to another variation of the pluralist response, which may differ only terminologically, the object of the saying-relation remains the proposition, rather than a set of propositions, but the relation is one-to-many; a single speaker may at one and the same time be saying many things.

The pluralist has ready responses to the paradoxes from Chapter 3, but neither is completely satisfactory. In response to non-epistemic underdetermination, the response is clear. The pluralist can deny the premise that at most one thing is said. By adopting this response, however, the pluralist concedes a lot to the problem of underdetermination. The pluralist concedes, for example, that the hidden-indexical theory of quantifier domain restriction (according to which one thing is said in any context) is false. In my view, these underdetermination arguments offer no reason to restrict the domain of viable linguistic theory.

In response to epistemic underdetermination, the pluralist might argue that the speaker knows what is said: it is each and every one of the candidates for what is said. The first thing to say is that, from the perspective of the alternative model, it is difficult to see how pluralism about what is said could be anything but an idle wheel in the explanation. We already have a more economical solution to the problem of underdetermination. Not only is our solution more economical in that it sanctions fewer entities said, it is, as mentioned above, consistent with a broader domain of linguistic theory than the pluralist response.

There is, however, a variety of pluralism about what is said that is supported by the account offered here. Indeed, this variety of pluralism is just the account offered here, presented in an alternative terminology. Buchanan and Ostertag (2005, 889) present a defence of the hidden-indexical version of Russellianism against the problem of underdetermination according to which “no proposition can nor need be identified as the proposition expressed”

by utterances of definite descriptions. On the face of it, this looks similar to the view defended in this Thesis, but Buchanan and Ostertag are utilising a very different notion of what is said by an utterance of the form ‘The F is G’, according to which the speaker says any proposition p such that the speaker intends to communicate p and p is of the form $[\text{the } x: Fx \ \& \ Hx](Gx)$. As such, a Buchanan and Ostertag’s position is, like Buchanan (2010) which uses the same terminology, that there is no proposition of that form that the speaker intends to communicate.

In contrast, the view of this Thesis, presented in the terminology of Buchanan and Ostertag (2005) is that many things are said, as there are many propositions of the required form that the speaker intends to communicate. Looking back to the arguments from Chapter 3 and interpreting them according to this new terminology, the epistemic argument fails, not because of the failure of the Epistemic Content Constraint, but because the audience can know what was said, that is, every proposition of the required form that is communicated. The non-epistemic argument fails because the third premise fails: many things are said, according to this terminology.

Speech act pluralism

In this section, we compare the view of communication presented here with Cappelen and Lepore’s (2005) *Speech Act Pluralism*. Comparison of the two views is extremely difficult, given that Speech Act Pluralism is, according to Cappelen and Lepore (2005, 190) “not really a theory; it’s a collection of observations, one of which is that there can be no *systematic theory* of speech act content.” The view offered here, however, is a systematic theory of what is asserted by the saying/implicating of a particular proposition. The fact that only one of these views aspires to the level of a theory is a significant and fundamental difference between them, but that is not to say that they are inconsistent. The view presented here can be seen as one way of vindicating some of the predictions made by Speech Act Pluralism.

Speech Act Pluralism is summarised as follows:

No one thing is said (or asserted, or claimed, or . . .) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated. What is said (asserted, claimed, etc.) depends on a wide range of facts other than the proposition semantically

expressed. It depends on a potentially indefinite number of features of the context of utterance and of the context of those who report on (or think about) what was said by the utterance.¹⁰⁷ (2005, 4)

The two views clearly have something in common. According to both views, what is asserted by a speaker depends on more than the semantic content of the sentence they utter. There are also *prima facie* differences evident even in this brief summary, however. According to Speech Act Pluralism, what is asserted depends on “a wide range of facts other than the proposition semantically expressed”, whereas we have identified a single such fact, that is, what is presupposed. The two views may not be so far apart on this point, however, as what is presupposed may itself be seen to depend on a wide range of facts such as “*Facts about the speaker’s intentions and beliefs ... Facets about the conversational context of this particular utterance ... Other facts about the world ... Logical relations*” (2005, 193).

Cappelen and Lepore go on to say that “an utterance can assert propositions that are not even logical implications of the proposition semantically expressed” (2005, 4). Our view agrees but says nothing about what is asserted beyond the logical implications of the conjunction of what is semantically expressed with what is presupposed.¹⁰⁸ Speech Act Pluralism says much more, however, going as far as to say that “Nothing even prevents an utterance from asserting ... propositions incompatible with the proposition semantically expressed by that utterance” (2005, 4). The mechanism described in this Thesis, however, cannot yield this result. No proposition entailed by what is said and what is presupposed can be inconsistent with what is said. Again, this is not to say that the view presented here is inconsistent with Speech Act Pluralism, but it is to say at the least that Speech Act Pluralism makes predictions that we can remain neutral on.

The final sentence in the indented quote above represents a further significant difference between the two views. What a speaker asserts may depend on “a potentially indefinite number of features of the context of utterance” but it is no part of the view presented here that asserted content depends on “the context of those who report on (or think

¹⁰⁷ Not that my usage of ‘what is said’ is more in line with the usage of ‘the proposition semantically expressed’ than the usage of ‘what is said’ in this quotation and throughout Cappelen and Lepore (2005).

¹⁰⁸ In cases other than literal speech, of course, what is implicated may play the role of what is semantically expressed.

about) what was said by the utterance”. What does depend on contexts of this latter sort is what is asserted by those who report on what was said by the original utterance.

How, according to Speech Act Pluralism, does the context of a reporter affect what is asserted by a speaker? Unfortunately, we are not told. Again, we are confronted by the fact that Speech Act Pluralism is a set of assertions (or, more charitably, observations) without explanation or theoretical backing. We are, however, given some indication of the methodology that leads to this conclusion, which Cappelen and Lepore (2005, 191) describe as taking “nontheoretic beliefs and intuitions about what speaker say, assert, claim, etc. at face value” while allowing that they “can, of course, be overridden should important theoretical considerations lead us to reject specific intuitions or general considerations.” This statement of the methodology looks like nothing more than the statement that intuitions are taken as defeasible evidence, which almost anyone would accept. Perhaps more central to the methodology is the suggestion that naïve intuitions are taken at face value without “overwhelming” reasons to the contrary, the result being that naïve speech reports are privileged radically. If someone can utter a sentence of the form ‘S said/asserted/claimed/etc. that P’ without extreme impropriety, then we should take it that S said/asserted/claimed/etc. that P.

Cappelen and Lepore are correct that the term ‘said’ has uses quite different from its use in the terminology of this Thesis and can be used equivalently to ‘assert’ or ‘claim’. However, their methodology of uncritical acceptance of naïve speech reports mistakenly assumes that all uses of ‘says’ are univocal. As Lewis (1998, 41) points out, however, ‘what is said’ “is very far from univocal. It can mean the propositional content ... It can mean the exact words. I suspect that it can mean almost anything in between.” It is doubtful, therefore, that anything of interest can be gleaned by following the Methodology of Speech Act Pluralism. As we have not followed this radically permissive methodology, Speech Act Pluralism makes many predictions that we can remain silent on.

We aren’t given an explanation, but we are given an example of how, according to Cappelen and Lepore (2005, 201) “Assumptions made by people who have nothing all to do with the original context of utterance can fix what the speaker said. To take a simple illustration, suppose you uttered (1) several weeks ago:

(1) The table is covered with books.

Suppose that whatever table is under discussion currently sits comfortably in your father's office (though it did not sit there when you uttered (1)). Haven't you said with your utterance of (1) that the table in your father's office is covered with books?"¹⁰⁹ Not according to the framework presented here. Given that the table was not presupposed to sit in your father's office when you uttered (1), you did not assert that the table in your father's office was covered with books.

Granted, it might be totally unobjectionable for a reporter to report you as having 'said that the table in your father's office is covered with books'. Suppose, for example, that the reporter is relaying instructions to a third party who is looking for something to read. Call the utterer of (1) 'A', the reporter 'B', and the third party looking for reading material 'C'. B asks A where C can find some reading material. A replies with (1). By the next time B sees C, the table in question has been moved, without A's knowledge, to C's father's office. B says to C (knowing that only the location of the table has changed, not its cover of books) 'A said that the table in your father's office is covered with books.' Given such a context, however, it is clear why such a report might be acceptable, even if it was false.¹¹⁰ By slightly misrepresenting the content A asserted, B more easily represents the location of the books, which is more helpful to C. Remember also that even if the proposition semantically associated with the utterance is false, B's utterance might, according to the view of this Thesis, assert various true propositions about A's original utterance.

The most theory-like presentation of Speech Act Pluralism comes in the form of eight theses: SPAP₁ – SPAP₈ (199-204).

SPAP₁: "No one thing is said (or asserted, or claimed, or . . .) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated, etc."

¹⁰⁹ It is unclear whether Cappelen and Lepore take the speaker to have said that the table in your father's office is covered with books *at the time it sits in your father's office* or that the table in your father's office was covered with books *at the time you uttered (1)*, or both.

¹¹⁰ False, that is, if 'said' is interpreted in the terminology we established at the outset. There may of course be alternative readings of 'said' according to which it is true. I am not here defending the claim that the proposition semantically expressed was actually false. I have not defended any view of definite descriptions, although I have shown how Russell's account can avoid various objections. If definite descriptions are, in fact, directly referential expressions, then the semantic content of A's report might well be true.

On this point, this Thesis agrees, if ‘said’ is interpreted in Cappelen and Lepore’s loose terminology, so as to be distinct from what they call above ‘the proposition semantically expressed’. The proposition semantically expressed (what we have been calling ‘what is said’) is (at most) only one among many propositions asserted by an utterance.

SPAP₂: “One of the many propositions asserted by an utterance is the semantic content of that utterance (the proposition semantically expressed).”

At least in cases of literal speech, where something is said, we are again in agreement. Remember, of course, that an utterance needn’t semantically express any proposition and yet communicate many propositions.

SPAP₃: “We have been unable to figure out how to devise an algorithm that takes the proposition semantically expressed and delivers all the propositions said, asserted, etc. There might not be *any* systematic theory from which one can derive all of which is said by an utterance.”

This is in line with the point made earlier that Speech Act Pluralism is not thought of as a theory by its authors. As we are aspiring to the level of theory, however, we should be able to provide an algorithm that takes the proposition semantically expressed (and the propositions presupposed) and delivers a set of propositions asserted. We suggested such an algorithm in Chapter 4: the propositions asserted are those that are not true at every possibility in the context set prior to intersection with the proposition said (or any of the best candidates for what is said) but are true at every possibility in the context set after intersection.

SPAP₄: “[Speech Act Pluralism] is opposed to theories that are speaker (or audience) centrist, and more generally, to Original Utterance Centrism. These are theories according to which the speech act content is fixed by facts about the speaker, his audience, and their common context. On our view, that’s definitely *not* the case.”

Again, this aspect of Speech Act Pluralism is not endorsed here. All propositions communicated are fixed by facts about the speaker, the audience, and their common context, where the context is taken to include all presuppositions.

SPAP₅: “It is a corollary of SPAP₄ that speakers don’t have privileged access to all the propositions they assert or say in uttering a sentence.”

While we haven’t endorsed SPAP₄, we might endorse SPAP₅, even though presented as a consequence of SPAP₄. If endorsed, however, it is as a corollary, rather than a central component of the view. It was suggested, for example, by the response to divergent presupposition according to which the proposition a speaker intends to communicate can be determined by the audience’s presuppositions. If A reads my ‘Jesus loves you’ sign, then I intend to communicate that Jesus loves A, despite being unable to entertain this proposition if I have no concept of who A is.

SPAP₆: “A corollary of SPAP₄ and SPAP₅ is that speakers need not believe everything they say even when their saying it is sincere.”

Again, we might endorse SPAP₆ as a corollary, but it is no part of the central view. The principle looks *prima facie* plausible to me, however, as making a sincere contribution to a conversation isn’t obviously limited to expressing beliefs. In Chapter 5 we offered a unified Russellian analysis of the distinction between referential and attributive uses of definite descriptions and saw that speakers may presuppose what they believe to be false in order to communicate truths, with the result that at least one proposition the speaker asserts is believed by them to be false. Such utterances do not seem to me necessarily insincere.

SPAP₇: A further corollary is that the principle Hawthorne calls the True Belief Principle is false (Hawthorne 2003, p. 99):

The True Belief Principle (TBP). If a speaker sincerely accepts an utterance *u* and *u* has semantic value *p*, then the belief manifested by his sincerely accepting *u* is true iff *p* is true.

TBP should not be accepted for several reasons. Firstly, it labours under the mistaken assumption that there is a particular belief manifested by such acceptance. In accepting an utterance with semantic value *p* (used literally, I assume) the speaker might manifest belief in several different propositions, as the original utterance may assert much more than its semantic content. Secondly, sincere acceptance of an utterance might not even manifest belief in the

semantic value of the uttered sentence but rather manifest belief in some of the other propositions asserted. The semantic value of a sentence is often just one among many propositions communicated, and needn't occupy a privileged position in the mind of the speaker or their audience. There seems nothing obviously insincere in your accepting my utterance of 'The murderer of Smith is insane', despite the fact that you believe Smith to have been killed by a group and so that the semantic value of this sentence is false, when you know that my primary intention was to communicate that Jones is insane.¹¹¹

SPAP₈: "The very same contextual features that determine the implicatures of an utterance influence what speakers say and assert by that utterance. There is no fundamental theoretical divide between sayings and implicatures. They are both on the side of speech act content. Whatever mechanisms might generate implicatures are also all used to generate what speakers say."

We allowed from Chapter 1 that the same mechanisms may be relevant in deriving implicatures and what is said.¹¹²

The benefit of underdetermination

In closing, I want to consider a further question: why would speakers leave what they say underdetermined? Is it just laziness that leads to underdetermination of what is said? Wouldn't things be easier if speakers took care to ensure that they utter particular sentences and say particular propositions? Indeed, perhaps there is here the beginning of a further argument that underdetermination remains problematic even if communication is possible, as we require an error theory to explain why people communicate so sloppily.

On the contrary, there are some benefits to be had from leaving what you utter or say underdetermined. Suppose that a speaker utters a particular sentence, thereby saying a particular proposition that they intend to communicate. There are various sorts of evidence

¹¹¹ Why not utter instead 'Jones is insane'? I might not know their name, or there may be too many potential referents with the same name in the context. Why not use 'He' or 'they' or 'that person'? Again, there may be too many potential referents.

¹¹² At this point, Cappelen and Lepore seem to shift back to a usage of 'what is said' more in line with their earlier usage of 'semantic content'. If they do not intend to change their usage, I remain unconvinced that they are really "Against Grice", as they say. (204).

which could indicate to their audience that they have misinterpreted at some stage in the process. Broadly, the interpreter might end up with something weird: they might end up with an ungrammatical sentence, or a proposition that seems to be trivial or irrelevant. When one of these things happens, they have to go back and reinterpret; perhaps first seeking out additional evidence on which to base their interpretation; perhaps taking the failure of their first interpretive attempt as additional evidence.

By allowing for underdetermination of what is uttered and said, we allow for an additional safeguard to misinterpretation: the candidates may fail to converge on the same thing communicated. The audience can take this as evidence that they were wrong about the candidates for what was uttered or said, or as evidence that they were wrong about the speaker's presuppositions, and this evidence can allow for reinterpretation. By leaving open a number of candidates for what is uttered or said, speakers can ensure that their audiences are on the same page in terms of the presuppositions in effect. If it is clear that you have intentionally left open two different candidates for what is said, you may thereby indicate that, relative to your presuppositions, they communicate the same thing, thereby allowing me to infer something about your presuppositions. You can convey more by saying less; indeed, by saying nothing.

Conclusion:

This Thesis has attempted to dislodge the Encoding Model of communication, according to which speakers encode the propositions they want to communicate into sentences, to be retrieved by their audiences. An effect of this model has been the proliferation of underdetermination arguments and subsequent restriction on the space of viable linguistic theories in an attempt to avoid underdetermination of what is said.

On the alternative model proposed here, the point of articulating a sentence is not to pass an encoded message that, by decoding in the right way, will reveal what the speaker intended to communicate. Often, there is no unique object that the audience should associate with the sentence uttered (indeed, there may be no unique sentence uttered that the audience should associate with the sentence articulated) and yet the audience can identify every proposition that the speaker intended to communicate.

This model was employed to respond to the problem of underdetermination. The audience cannot know what was said, as there are a number of equally viable candidates for what is said. However, given that these candidates are locally equivalent, the same propositions will be communicated, whichever was said. The audience can, therefore, know what the speaker intended to communicate, even though they cannot know what the speaker said. Indeed, the audience can know what the speaker intended to communicate, even if, as argued in Chapter 3, nothing was said.

Bibliography

Bach, K. 1981: Referential/attributive. *Synthese*, 49 (2), 219 - 244.

Bach, K. 1994: Conversational Implicature. *Mind & Language*, 9(2), 124-162.

Benacerraf, P. 1965: What Numbers Could not Be. *The Philosophical Review*, 74(1), 47-73.

Blackburn, W. K. 1988: Wettstein on Definite Descriptions. *Philosophical Studies*, 53(2), 263-278.

Borg, E. 2002: Pointing at Jack, Talking about Jill: Understanding Deferred Uses of Demonstratives and Pronouns. *Mind & Language*, 17(5), 489-512.

Buchanan, R. and Ostertag, G. 2005: Has the problem of Incompleteness Rested on a Mistake? *Mind*, 114(456), 889-913.

Buchanan, R. 2010: A Puzzle about Meaning and Communication. *Noûs*, 44(2), 340-371.

Buchanan, R. 2012: Is belief a propositional attitude? *Philosophers' Imprint*, 12(1).

Cappelen, H and Lepore, E. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Cappelen, H and Lepore, E. 1997. On an Alleged Connection Between Indirect Speech and the Theory of Meaning. *Mind & Language*, 12(3/4), 278-296.

Carston, R. 1988. Implicature, explicature, and truth-theoretic semantics. R. M. Kempson (ed.) *Mental Representations: The interface between language and reality*. Cambridge: Cambridge University Press. 155-181.

Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Malden, MA: Blackwell.

Clapp, L. 2002. What Unarticulated Constituents Could Not Be. J. K. Campbell, M. O'Rourke and D. Shier (eds) *Meaning and Truth: Investigations in Philosophical Semantics*. New York: Seven Bridges Press. 231-256.

Corazza, E., Fish, W., and Gorvett, J. 2002. 'Who Is I?'. *Philosophical Studies*, 107(1), 1-21.

Donnellan, K. 1966. Reference and Definite Descriptions. *The Philosophical Review*, 75(3), 281-304.

Donnellan, K. 1968. Putting Humpty Dumpty Together Again. *The Philosophical Review*, 77(2), 203-215.

Egan, A. 2009. Billboards, bombs and shotgun weddings. *Synthese*, 166(2), 251-279.

Frege, G. 1956. The Thought: A Logical Inquiry. *Mind*, 65(259), 289-311.

Frege, G. 1979. *Logic in Mathematics* in H. Hermes, F. Kambartel, and F. Kaulbach (eds.), P. Long, R. White (trans) *Posthumous Writings*. Oxford: Blackwell, 203-250.

Fine, K. 1975. Vagueness, Truth and Logic. *Synthese*, 30(3/4), 265-300.

von Fintel, K and Gillies, A. 2011. *Might Made Right*, in A. Egan, and B. Weatherson (eds.) *Epistemic Modality*. Oxford: Oxford University Press, 108-130.

Greenough, P. 2008. Indeterminate Truth. *Midwest Studies in Philosophy*, 32(1), 213-241.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hawthorne, J. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Heim, I. 1983. File Change Semantics and the Familiarity Theory of Definites. Reprinted in P. Portner and B. Partee (eds) *Formal Semantics: The Essential Readings*. Malden, MA: Blackwell, 2002. 223-248.

Heim, I. 1983. On the Projection problem for Presuppositions. Reprinted in P. Portner and B. Partee (eds) *Formal Semantics: The Essential Readings*. Malden, MA: Blackwell, 2002. 249-260.

Kamp, H. and Reyle, U. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic Publishers.

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

Korta, K., and Perry, J. Forthcoming. Full but not saturated. The myth of mandatory primary pragmatic process. S. Conrad and K. Petrus (eds), *Meaning, Context, and Methodology*.

Kripke, S. 1977. Speaker's Reference and Semantic Reference. *Midwest Studies in Philosophy*, 2(1), 255-276.

Kripke, S. 1982. *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.

Lewis, D. 1975. Languages and Language. Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, Volume VII. Minneapolis: University of Minnesota Press, pp. 3–35.

Lewis, D. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(3), 339.

Lewis, D. 1998. Index, context, and content. *Papers in philosophical logic*. Cambridge: Cambridge University Press, 21-44.

Loar, B. 1976. The semantics of singular terms. *Philosophical Studies*, 30(6), 353-377.

Ludlow, P. 1989. Implicit Comparison Classes. *Linguistics and Philosophy*, 12(4), 519–533.

Neale, S. 1990. *Descriptions*. Cambridge, MA: MIT Press.

- Perry, J. 1998. Indexicals and demonstratives. C. Wright and B. Hale (eds) *A Companion to the Philosophy of Language*. Second Edition. Oxford: Blackwell. 586-612. First published 1997.
- Predelli, S. 1998. 'I Am Not Here Now'. *Analysis*, 58(2), 107-115.
- Predelli, S. 1998. 'Utterance, Interpretation and the Logic of Indexicals'. *Mind & Language*, 13(3), 400-414.
- Predelli, S. 2005. *Contexts: Meaning, Truth, and the Use of Language*. Oxford: Oxford University Press.
- Predelli, S. 2013. *Meaning Without Truth*. Oxford: Oxford University Press.
- Putnam, H. 1967. It Ain't Necessarily So. *The Journal of Philosophy*, 59(22), 658-671.
- Recanati, F. 1996. Domains of Discourse. *Linguistics and Philosophy*, 19(5), 445-475.
- Reimer, M. 1992. Incomplete descriptions. *Erkenntnis*, 37(3), 347-363.
- Romdenh-Romluc, K. 2002. Now the French are invading England! *Analysis*, 62(1), 34-41.
- Rostworowski, W. 2013. Roundabout Semantic Significance of the "Attributive/Referential" Distinction. *Kriterion*, 27(1), 30-40.
- Russell, B. 1905. On Denoting. *Mind*, 14(56), 479-493.
- Russell, B. 1998. *The problems of Philosophy*. Second edition. Oxford: Oxford University Press. First published 1912.
- Schiffer, S. 1981. Indexicals and the Theory of Reference. *Synthese*, 49(1), 43-100.
- Schiffer, S. 1992. Belief Ascription. *The Journal of Philosophy*, 89(10), 499-521.
- Schiffer, S. 1994. A Paradox of Meaning. *Noûs*, 28(3), 279-324.

- Schiffer, S. 1995. Descriptions, Indexicals, and Belief Reports, Some Dilemmas (But Not the Ones You Expect). *Mind*, 104(413), 107-131.
- Sidelle, A. 1991. The answering machine paradox. *Canadian Journal of Philosophy*, 21(4), 525-539.
- Soames, S. 1994. Donnellan's Referential/Attributive Distinction. *Philosophical Studies*, 73(2/3), 149-168.
- Sperber, D. and Wilson, D. 1986. *Relevance: Communication and Cognition*. Malden, MA: Blackwell.
- Stainton, R. 1995. Non-Sentential Assertions and Semantic Ellipsis. *Linguistics and Philosophy*, 18(3), 281-296.
- Stalnaker, R. 1999. *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, R. 2006. Responses. J. Thomson and A. Byrne (eds), *Content and Modality: Themes from the Philosophy of Robert Stalnaker*. Oxford: Oxford University Press. 251-294
- Stanley, J. 2000. Context and Logical Form. *Linguistics and Philosophy*, 23(4), 391-434.
- Stanley, J and Gendler Szabó, Z. 2000. On Quantifier Domain Restriction. *Mind and Language*, 15(2&3), 219-261.
- Strawson, P. F. 1950. On Referring. *Mind*, 59(235), 320-344.
- Strawson, P. F. 1954. A Reply to Mr. Sellars. *The Philosophical Review*, 63(2), 216-231.
- Travis, C. 1985. "On What is Strictly Speaking True." *Canadian Journal of Philosophy*, 15, 187-229.
- Unger, P. 1980. The Problem of the Many. *Midwest Studies in Philosophy*, 5, 411-467.
- Weatherson, B. 2002. Misleading Indexicals. *Analysis*, 63(4), 308-310.

Wettstein, H. 1981. Demonstrative Reference and Definite Descriptions. *Philosophical Studies*, 40(2), 241-257.

Quine, V. W. O. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Quine, V. W. O. 1970. On the Reasons for Indeterminacy of Translation. *The Journal of Philosophy*, 67(6), 178-183.

Yablo, S. 2006. Non-Catastrophic Presupposition Failure. J. Thomson and A. Byrne (eds), *Content and Modality: Themes from the Philosophy of Robert Stalnaker*. Oxford: Oxford University Press. 164-190.

Yalcin, S. 2014. Semantics and Metasemantics in the Context of Generative Grammar. A. Burgess and B. Sherman (eds) *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press. 18-54.