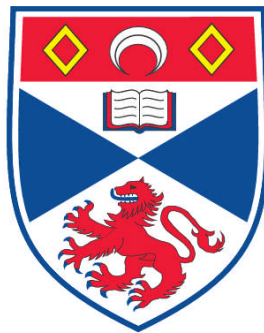


ON DESCRIBING

Anders J. Schoubye

**A Thesis Submitted for the Degree of PhD
at the
University of St. Andrews**



2011

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/2468>

This item is protected by original copyright

ON DESCRIBING

Anders J. Schoubye

*Submitted for the degree of Doctor of Philosophy
University of St Andrews, 2011*

I, Anders Johan Schoubye, hereby certify that this thesis, which is approximately 65.295 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in 2007 and as a candidate for the degree of Doctor of Philosophy in 2007; the higher study for which this is a record was carried out in the University of St Andrews between 2007 and 2011.

Date

Signature of candidate

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

Date

Signature of supervisor

In submitting this thesis to the University of St Andrews we understand that we are 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. We 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. We 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 electronic publication of this thesis: Access to Printed copy and electronic publication of thesis through the University of St Andrews.

Date

Signature of candidate

Signature of supervisor

ABSTRACT

The overarching topic of this dissertation is the semantics and pragmatics of definite descriptions and its central question is whether a sentence such as ‘the king of France is bald’ *literally asserts* the existence of a unique king (and therefore is false) or simply *presupposes* the existence of such a king (and thus fails to express a proposition).

One immediate obstacle to resolving this question is that truth value judgments seem particularly unstable with regards to these cases (viz. sentences containing non-denoting descriptions). Some of these elicit a strong intuition of falsity whereas others seem simply awkward or strange. Because of this, truth value judgments are generally considered unreliable. This is the central topic of Chapter 2 where I attempt to develop an explanation of this phenomenon. I observe that when these types of sentences are considered in the context of a discourse, a systematic pattern in judgments emerges. This pattern, I argue, should be explained in terms of certain pragmatic factors, e.g. whether a speaker’s utterance is interpreted as cooperative. I then lay out a detailed and general explanation of the phenomenon which draws importantly on recent research in the semantics and pragmatics of questions and focus. I show that the behavior of these judgments *can* be systematically explained, that truth value judgments are not as unreliable as standardly assumed, and that my proposed explanation best supports the conclusion that definite descriptions presuppose rather than assert existence.¹

If definite descriptions are assumed to literally *assert* existence, a sentence such as ‘Hans wants the ghost in his attic to be quiet’ is incorrectly predicted to be true only if Hans wants there to be a ghost in his attic. This prediction is often considered evidence against Russell’s quantifica-

¹ A version of this chapter was published in *Linguistics and Philosophy*, cf. Schoubye (2009).

tional analysis and evidence in favor of the referential analysis of Frege and Strawson. In Chapter 3, I argue that this problem is a general problem about the existence commitments of natural language determiners, i.e. not an argument in favor of a referential analysis. I show that in order to avoid these undesirable predictions, quite radical changes to the semantic framework are required. For example, it must be assumed that a sentence of the form ‘The *F* is *G*’ has the open sentence ‘*x* is *G*’ as its asserted content. I outline a uniform quantificational and presuppositional analysis of definites and indefinites which by exploiting certain features of so-called dynamic semantics unproblematically assumes that the asserted contents indeed *are* open sentences.

In view of the uniform semantic analysis defended in Chapter 3, I use Chapter 4 to resist a recent argument due to Reimer (1998) and Devitt (2004) that definite descriptions are ambiguous between *attributive* and *referential* uses. In contrast to Donnellan’s (1966) pioneering arguments in favor of this distinction, Reimer’s and Devitt’s argument rests on the assumption that definite descriptions are *regularly* used to communicate singular thoughts and that the *standard* meaning of definite descriptions *therefore* must be fundamentally indexical and directly referential. I argue that this argument relies crucially on tacit assumptions about semantic processing for which there is no empirical evidence. I then show that the argument is too general; if sound, it would be an argument for an indexical treatment of most, if not all, other determiners. I conclude by demonstrating that the view does not explain any new data and thus has no clear motivation. I conclude by providing a rough outline of how the uniform analysis presented earlier is capable of explaining referential uses.

These three chapters are preceded by a short general introduction to definite descriptions. There I explicate Russell’s highly influential theory of descriptions and I outline its motivations. I also briefly introduce four of the standard objections to Russell’s theory since these are relevant to the subsequent chapters.

The dissertation is supplemented with an appendix where I discuss a wide range of data relevant to Chapter 3. In principle, this could have been included there, but since Chapter 3 deals with a large number of theoretical issues already, I thought it best to postpone a comprehensive survey of the data to a later stage in order to maintain the focus of the chapter.

Summarizing, I have aimed at providing a detailed pragmatic explanation of a long-standing puzzle about truth value judgments and then at outlining a novel dynamic semantic analysis of definite (and indefinites) descriptions.

This analysis solves a significant problem about existence commitments — a problem that neither the standard Russellian analysis nor the Frege/Strawson analysis are equipped to handle. This proposed analysis is then defended against the claim that definite descriptions are lexically ambiguous between attributive and referential uses.

ACKNOWLEDGMENTS

This dissertation is to a great extent a product of ideas emanating from conversations with my primary supervisor, Herman Cappelen, and my greatest debt is therefore owed to him. Throughout the past four years, Herman has been the model supervisor. He has patiently listened to my often disoriented and typically handwaving ideas, read through numerous drafts, relentlessly provided crucial feedback, but also constantly pushed me to revise, improve, and clarify. In addition to this, Herman has had to cope with supervising a somewhat stubborn, impatient, and occasionally arrogant student who in one week would claim to have everything figured out and who then the next would lapse into depression because nothing made sense anymore. Without Herman's persistent support, his insightful comments, and his never failing ability to identify the interesting questions, it is doubtful that I would ever have finished this dissertation. I am thus hugely grateful to Herman for his help.

My second greatest debt in writing this dissertation is owed to Jonathan Schaffer and Josh Dever. Jonathan and Josh are undoubtedly among the smartest and most talented philosophers I have ever met and their comments and feedback have been absolutely invaluable. I cannot count the number of drafts that Jonathan has read, and the second chapter in particular owes much to his incisive and constructive comments. Jonathan's happy, always patient, and genuinely caring nature is a rare gem in this discipline and if I ever get the chance to supervise students myself, I can only hope that they will think of me as I think of Jonathan.

The third chapter was unquestionably the hardest for me to write and the chapter that it took the longest for me to finish. The bulk of this chapter was written during my stay in Austin, TX, and had it not been for hours and hours on end of me turning up at Josh's office for advice, I am convinced that I

would have given up. Josh seems to have this amazing ability to immediately grasp the depth of any given problem and moreover of the obstacles one will have to overcome in order to solve it. For several weeks, conversations with Josh were virtually the only thing that kept the chapter alive — so much so that I feel that Josh should almost have the status of co-author.

Finally, I feel I owe a special thanks to Dilip Ninan. Dilip took time out on several occasions to read and discuss my work and his comments were invariably extremely helpful. In general, Dilip always had positive and constructive things to say and this made talking to him a particularly delightful experience.

The additional people who have helped and supported my work over the past several years are far too many to list here. However, one relatively small group of people deserve a special mention — not only because of their often significant contributions to my work but also because of their friendship which during the last four years has meant a great deal to me. These are Nate Charlow, Liz Coppock, Mikkel Gerken, Gudmundur Andri Hjálmarsson, Torfinn Huvenes, Daniele Sgaravatti, and Andreas Stokke.

During my years in St Andrews, I benefited immensely from interactions with a great number of people, including Josh Armstrong, Ralf Bader, Derek Ball, Björn Brodowski, Jessica Brown, Yuri Cath, Peter Clark, Josh Clarkson, Sharon Coull, Mike De, Paul Dimmock, Dylan Dodd, Andy Egan, Patrick Greenough, Katherine Hawley, Ole Hjortland, Tom Hodgson, Jonathan Ichikawa, Antti Kauppinen, Dirk Kindermann, Julia Langkau, Federico Luzzi, Andrea Onofri, Nikolaj Jang Pedersen, Esben Nedenskov Petersen, Walter Pedriali, Graham Priest, Stephen Read, François Recanati, Stewart Shapiro, Martin Smith, Jason Stanley, Rachel Sterken, Margot Strohming, Paula Sweeney, Brian Weatherson, Crispin Wright, Elia Zardini, and Dan Zeman.

A substantial part of this dissertation was written during my visit to the University of Texas at Austin in 2009-2010. In Austin I was very fortunate to be surrounded by some amazingly smart people. I want to thank, in particular, Nicholas Asher, David I. Beaver, and Hans Kamp who not only read through drafts and provided me with extremely helpful comments and feedback, but who also were immensely supportive of my project and made me feel at home in Austin. I learned more than I could have hoped for from conversations with Nicholas, David and Hans.

I also want to thank to the rest of the philosophers and linguists in Austin that I was fortunate to interact with during my extremely enjoyable visit, namely Derek Anderson, John Bengson, Nora Berenstain, Georgui

Bronnikov, Ray Buchanan, Justin Cope, Brian Cutter, Sinan Dogramaci, Jeremy Evans, David Frank, Joey Frazee, Enrico Grube, Alex Grzankowski, Julie Hunter, Steven James, Malcolm Keating, Aidan McGlynn, Katherine Piatti, Bryan Pickel, Ian Proops, Kate Ritchie, Mark Sainsbury, David Sosa, Steve Wechsler, Malte Willer, and Briggs Wright.

Lastly, I want to thank Paul Elbourne, Kai von Fintel, and Craig Roberts whom I have corresponded with over the years and who have been very kind to take time out to answer questions and comment on my research.

I also kindly acknowledge the University of St Andrews, the Centre for the Study of Mind in Nature, and the Royal Institute of Philosophy for their financial support.

I humbly dedicate this dissertation to my mother.

Anders J. Schoubye
Oslo, April 2011

CONTENTS

I	Introduction	17
1	On Denoting.....	17
1.1	Russell's Theory of Descriptions.....	18
1.2	Motivations	19
2	Standard Objections to Russell's Analysis	25
2.1	The Mismatch between Syntax and Semantics.....	26
2.2	Presuppositions and Truth Value Gaps	27
2.3	Referential and Attributive Uses of Definite Descriptions	29
2.4	Incomplete Definite Descriptions	31
3	Concluding Remarks.....	32
II	Descriptions, Truth Value Intuitions, and Questions	35
1	Introduction	35
2	The Intuitive Pull.....	38
2.1	The Data	38
2.2	Accommodation Theories	40
2.3	Systematizing the Strawsonian Insights	44
3	SWEPPs In Context	45
4	Discourses and Questions	48
4.1	Discourse Structure	49
4.2	Consonance	51
4.3	The False Intuition Hypothesis (FIH)	53
4.4	Predictions	55
4.5	Consonance Constraints.....	56
5	Non-Denoting Definite Descriptions and Focus	59
5.1	Backgrounds and Focus	59

5.2	Structured Meanings: The Semantics and Pragmatics of Focus.....	60
5.3	Focus and Non-Denoting Definites	62
6	Extending the Proposal	65
6.1	Complex Constructions and Intensional Interpretations	65
6.2	Out-of-Context SWEPPs	66
6.3	Discourse Topics and Practical Relevance	67
7	Existential vs. Presuppositional Analyses	70
8	Concluding Remarks.....	72
III	Ghosts, Murderers, and the Semantics of Descriptions	73
1	Introduction	73
2	The Existence Problem	79
2.1	Problem and Plan	79
2.2	The Presuppositional Solution	80
2.2.1	Presuppositions and Partial Functions	81
2.2.2	Presuppositions and Referential Terms.....	84
2.2.3	Presupposition Projection in Attitude Contexts	86
2.2.4	Interim Conclusion.....	87
2.3	Indefinite Descriptions: The Failure of the Referential Analysis	88
2.4	The Vagaries of Attitude Verbs	92
3	Towards a Solution to the Existence Problem	97
3.1	Three Desiderata	97
3.2	Meaning as Update Potentials.....	100
3.2.1	Dynamic Semantics	101
3.2.2	Presuppositional Definites	103
3.2.3	Presuppositional Indefinites	105
3.2.4	Solving the Existence Problem for Conditionals.....	106
3.3	Propositional Attitudes	107
4	Conclusion	111
IV	Attributive/Referential: Against the Argument from Regularity	113
1	Attributive and Referential Descriptions	113
2	Kripke's Pragmatic Response	116
2.1	Referential Uses of Proper Names and Quantifiers	117
2.2	A Defense of Referential Descriptions.....	118

3	The Argument from Regularity	119
3.1	Regular Usage	119
3.2	Grasping Referential Uses	123
3.3	Kripke's Objections Revisited	124
3.4	Analyzing The Argument from Regularity	126
4	Evaluating the Argument from Convention	127
4.1	Communicative Success and Referential Descriptions ...	128
4.2	Generalizing the (RA) to other Determiners	130
4.3	Linguistic Meaning vs. Literal Meaning	133
4.3.1	Overgeneration I: Scalar Implicatures	134
4.3.2	Overgeneration II: Ability Modals	139
4.4	Explanatory Value	143
5	Gödelian Completions	144
5.1	Gödelian Completions and Dynamic Semantics	147
6	Concluding Remarks	150
A	Indefinites and Presuppositions	153
1	Preliminary Observations	153
1.1	Entailments, Implications and Presuppositions: Projection	155
2	Projection Tests	156
2.1	Questions	156
2.2	Conditionals	158
2.3	Modals	161
3	The 'Hey, Wait a Minute'-Test	162
4	Analysis	164
4.1	Weak Determiners and Presupposition Accommodation	164
	References	168

CHAPTER I

INTRODUCTION

1 On Denoting

With the possible exception of Frege's seminal 'Über Sinn und Bedeutung' (1892), Russell's paper 'On Denoting' (1905) is arguably the most important and influential paper in twentieth-century philosophy. Among a substantial number of philosophers of language, the essentials of the theory of descriptions put forward in 'On Denoting' is still widely considered the state of the art, and considering that 'On Denoting' was published more than a hundred years ago, this is nothing short of a truly remarkable achievement. The number of books, articles, notes, and reviews that take Russell's insights in 'On Denoting' as their starting point is plainly staggering. Russell's paper simply defined a discipline of philosophy — philosophy of language.

This dissertation is devoted primarily to a discussion of issues emanating from Russell's seminal paper. More specifically, it is devoted to subjecting some of the widely considered virtues of Russell's analysis to further scrutiny and criticism. While one cannot sensibly deny the importance of Russell's paper nor Russell's genius, one can still challenge Russellian orthodoxy. In this dissertation, I will therefore focus on a few interesting problems for the Russellian theory that even today have not been properly appreciated. Since these problems reveal various interesting complexities of natural language that seem to go beyond what one might reasonably have expected, I hope to convince the reader that these are problems that are worth discussing.

Chapter II and Chapter III (the chapters following this relatively short introduction) are devoted to a discussion of two specific aspects of Russell's analysis that among contemporary philosophers of language are typically cited as significant advantages. The aim of these chapters is to show that

when these aspects are properly examined, we find that they are in fact *disadvantages*. But we also find that the underlying issues are significantly more complex than standardly assumed and that to understand the semantics and pragmatics of definite descriptions, both our methodology and our semantic analysis must be refined.

The second aim is to show the way forward, i.e. provide solutions to the relevant problems and thereby pave the way for an adequate semantic and pragmatic analysis of definite descriptions.

In Chapter IV, I turn to a discussion of a standardly assumed challenge to Russell's theory, namely so-called referential uses of definite descriptions. In recent years, some researchers have attempted to revive this challenge, originally due to Donnellan (1966), but I argue that their main argument is flawed. That is, while I *do* argue in Chapters II and III that Russell's theory should be abandoned, I argue in Chapter IV that the existence of referential uses of descriptions is not a reason why.

Now, Russell's theory is motivated by various important epistemological, metaphysical, and semantic—pragmatic considerations, and an entire dissertation could easily be devoted to any of these. In this dissertation, my main focus is on semantic and pragmatic issues and the extent to which Russell's theory of descriptions can be said to be an empirically adequate analysis of definite descriptions as used by speakers of natural language. This means that I engage with only a tiny fraction of the issues relevant to the general topic of definite descriptions and it is therefore almost trivial to emphasize that much more would need to be said in order to conclusively determine what the best analysis of definite descriptions is.

I begin the next section by providing a short, and admittedly crude, overview of Russell's theory and of a selection of the rather astonishing number of issues that were raised by Russell's paper.

1.1 *Russell's Theory of Descriptions*

In the first paragraph of 'On Denoting', Russell writes,

By a "denoting phrase" I mean a phrase such as any one of the following: a man, some man, any man, every man, all men, the present king of England, the present king of France, the centre of mass of the Solar System at the first instant of the twentieth century, the revolution of the earth round the sun, the revolution of the sun round the earth. Thus a phrase is denoting solely in virtue of its *form*. (Russell 1905: 479)

Russell's central and controversial thesis in 'On Denoting' is that definite descriptions are not singular terms that *refer* to particular individuals. Russell argues instead that both definite and indefinite descriptions are a

species of what he calls *denoting phrases*, i.e. quantifier phrases, and he proposes to analyze these descriptions in terms of the simple unrestricted first-order quantifiers ' $\forall x$ ' and ' $\exists x$ '. Hence, according to Russell, on a proper analysis of sentences of the form 'an/some F is G ' and 'the F is G ', these are effectively shorthand for more complex quantificational constructions, viz. the constructions in (2) and (4) respectively.

- (1) An/Some F is G .
- (2) $\exists x[F(x) \wedge G(x)]$
- (3) The F is G .
- (4) $\exists x[F(x) \wedge \forall y[F(y) \rightarrow x = y] \wedge G(x)]$

Hence, on Russell's theory, the sentence in (1) is analyzed as making two distinct claims, namely A1 and A3 below, whereas the sentence in (3) is analyzed as also making the additional claim A2.

- A1 There is an individual x which is F .
- A2 For all individuals y such that y is F , if y is F , then y is identical to x .
- A3 x is G .

In short, a sentence such as (3) *asserts* that there is an individual x who instantiates the property F , that only one individual in the domain instantiates the property F , and that this individual also instantiates the property G . Let's refer to A1 as the assertion of existence, A2 as the uniqueness constraint, and A3 as the property attribution. This is, in its essentials, Russell's theory of descriptions. And despite its deceptively simple appearance, it is a remarkably powerful theory. Of the numerous discussions devoted to these three conditions above, most have revolved around the status of A2, the uniqueness condition. The majority of this dissertation is devoted to a discussion of A1, the assertion of existence (in particular Chapters II and III).

1.2 Motivations

The easiest way to demonstrate the strength of Russell's theory is by considering the range of epistemic, metaphysical and semantic problems that the theory is capable of solving. Let's begin by considering one of the metaphysical-cum-semantic problems, namely non-denoting descriptions. Consider the sentences below.

- (5) Bertrand is bald.
- (6) The president of the United States is bald.

(7) The king of France is bald.

It seems quite natural to analyze both (5) and (6) as consisting of a singular term (the term in subject position) referring to a particular individual and an attribution of a property to this individual. However, this intuitively natural analysis is problematic with respect to (7). If the description in (7) is analyzed as a referential term, it then seems that a speaker asserting (7) should be capable of expressing a thought *about* the king of France. But since it is common knowledge that France is not a monarchy, and hence has no king, it is not clear how this sentence could in any conceivable sense be used to express such a thought. After all, what would the speaker asserting this sentence be entertaining a thought about? Hence,

[...] one would suppose that “the King of France is bald” ought to be nonsense; but it is not nonsense, since it is plainly false. (Russell 1905: 484)

As Russell observes, this raises the following puzzle. Even though the subject of the sentence is a non-denoting definite description, the sentence is clearly grammatical and it seems perfectly meaningful.

One type of explanation that Russell’s theory intended to refute was the explanation put forward by Meinong (1904). Meinong proposed to distinguish between the notions of *existence* and *subsistence*, and he argued in favor of the view that *existence* is not a precondition for successful reference but rather a property with the same status as other simple properties such as *color* or *mass*. Thus on Meinong’s view, at least as Russell interprets it, definite descriptions such as ‘the king of France’ do indeed refer — only to objects that fail to instantiate the property of *existence*.¹ But Russell considered Meinong’s view to have untenable metaphysical implications.

Of the possible theories which admit such constituents the simplest is that of Meinong. This theory regards any grammatically correct denoting phrase as standing for an object. Thus “the present King of France,” “the round square,” etc., are supposed to be genuine objects. It is admitted that such objects do not subsist, but nevertheless they are supposed to be objects. This is in itself a difficult view; but the chief

¹ Whether Russell’s interpretation of Meinong’s view is correct is, I think, controversial. More recently, some researchers, e.g. Parsons (1980), have defended the Meinongian view and attempted to establish its consistency. See also Lambert (1983). Nevertheless, the standard narrative here is that ‘On Denoting’ delivered a fatal blow to Meinong’s theory by thoroughly undercutting its motivation. Interestingly, Russell himself was sympathetic to this kind of view in Russell (1903) which was published just a few years before ‘On Denoting’.

objection is that such objects, admittedly, are apt to infringe the law of contradiction. It is contended, for example, that the existent present King of France exists, and also does not exist; that the round square is round, and also not round; etc. But this is intolerable; and if any theory can be found to avoid this result, it is surely to be preferred. (Russell 1905: 482-483)

In contrast to Meinong's proposed solution to the problem of non-denoting descriptions, which relied on a rather promiscuous ontology, Russell's theory of descriptions provides an elegant solution that retains a conservative ontology (an ontology without non-existent entities) and yet explains why these sentences are meaningful.

Since Russell treats definite descriptions as complex denoting phrases, the truth conditions of (7) are captured in terms of the following existential statement.

$$(8) \quad \exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y] \wedge \text{bald}(x)]$$

When (7) is analyzed as having the truth conditions given by (8), it is predicted to be straightforwardly false, as Russell intended, and it is now easy to explain why it is meaningful; the definite description is not the *logical subject* of the sentence and this means that one need not interpret the sentence as being about the individual (potentially) denoted by the description. Rather, it is a sentence that simply makes an existential assertion, viz. the assertion that there exists a unique king of France and that this individual instantiates the property of being bald. This sentence is perfectly meaningful and straightforwardly false, since there is no king of France. One particularly elegant feature of this explanation is that the prediction relies only on the assumption that the definite description is not a singular term but instead a complex denoting phrase.

An additional related advantage of Russell's proposed analysis is that it provides an elegant and intuitive explanation of so-called negative existentials, i.e. sentences such as (9).

$$(9) \quad \text{The king of France does not exist.}$$

The intuitive problem with negative existentials is that these sentences purport to predicate the property of non-existence to the individual picked out by the expression in subject position. However, it also seems intuitively clear that for an object to instantiate the property predicated of it, it must at least exist (Meinong's theory was also devised to address this particular problem). However, this assumption cannot conceivably be satisfied when the property in question is that of non-existence, and so it seems that these sentences could not possibly make sense. Nevertheless, negative existentials,

e.g. (9), are perfectly meaningful and often intuitively true. On Russell's analysis, this is easily explained: The logical form of (9) is given by (10).

$$(10) \quad \neg \exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y]]$$

Here, the negative existential in (9) is analyzed as asserting that it is not the case that there is a unique individual x such that x instantiates the property of being the king of France — and this of course seems to capture correctly the truth conditions of (9). Again, the beauty is that this prediction falls out naturally from the assumption that definite descriptions are denoting complexes rather than referential terms. Since the definite description in (9) is a complex quantificational construction, it interacts with scope-taking operators such as negation — and by allowing the negation to scope over the existential assertion, the desired truth conditions are derived.

Similarly, Russell's analysis can also explain why a sentence such as (11) appears to have both a true and a not-true interpretation. I.e. it seems that one can use (11) to express either (11a) or (11b).

- (11) The king of France is not bald.
- a. $\neg \exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y] \wedge \text{bald}(x)]$
 - b. $\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y] \wedge \neg \text{bald}(x)]$

That is, if (11) is used to assert a true denial of existence, as in 'the king of France is not bald, because there is no king of France', the logical form of the assertion is that given by (11a). In contrast, if the speaker falsely believes that there is a king of France and intends to convey the information that the king of France is bald, the logical form of the speaker's assertion is that given by (11b) — and this is, as Russell intended, predicted to be false.

The view that (11) is structurally ambiguous fits well with our intuitions about the possible meanings of (11), but Russell's analysis of definite descriptions commits us to a more general conclusion. If definite descriptions are complex quantificational constructions, we should expect ambiguities to arise whenever these interact with various scope-taking expressions such as other quantifiers, modals, if-clauses, and quantificational adverbs.²

One type of construction where these putative scope ambiguities are apparent are sentences where a definite description is embedded in the scope of a propositional attitude verb, such (12).

- (12) Bertrand believes that the president of the United States is bald.

² This has been the subject of much debate. This general consequence of Russell's view is forcefully defended in Neale (1990: 118-164), but more recently Rothschild (2006, 2007) and Glanzberg (2007) have raised significant worries.

This sentence is ambiguous between the *de re* reading where Bertrand believes of some *particular* individual that he is bald (perhaps not knowing that this individual is the US president), and the *de dicto* reading where Bertrand believes of no particular individual (i.e. *whoever* is the president of the United States) that he is bald. These readings are however straightforwardly captured on Russell's analysis.

- (12) a. $\exists x[US\text{-}President(x) \wedge \forall y[US\text{-}President(y) \rightarrow x = y] \wedge \text{Bertrand believes that } bald(x)]$ (*de re*)
 b. Bertrand believes that $\exists x[US\text{-}President(x) \wedge \forall y[US\text{-}President(y) \rightarrow x = y] \wedge bald(x)]$ (*de dicto*)

In other words, *de dicto* and *de re* ambiguities are structural (or scope) ambiguities on Russell's analysis, and these are then resolved in the exact same way that other structural ambiguities (i.e. those arising from sentences containing multiple quantifier phrases) are standardly resolved, namely by determining the scopes of the relevant constituents. And this analysis has an even further putative advantage. Consider the sentence below.

- (13) Bertrand believes that the king of France is bald.

Suppose Bertrand falsely believes that France is a monarchy and the king of France is bald. If so, it seems that there is a straightforwardly true reading of (13) and this is easily explained using Russell's analysis. That is, the analysis allows us to capture not only Bertrand's belief that the king of France is bald, but also the prerequisite belief in the existence of a king of France. This is done by analyzing the definite description in (13) as taking scope under the attitude verb as in (14).

- (14) Bertrand believes that $\exists x[king\text{-of-France}(x) \wedge \forall y[king\text{-of-France}(y) \rightarrow x = y] \wedge bald(x)]$ (*de dicto*)

Russell's analysis thus elegantly explains why certain attitude attributions are intuitively judged true even though they contain an apparent reference to a non-existing entity. And on the face of it, this is something that any non-quantificational analysis of definite descriptions should have considerable problems explaining.

The explanatory virtues of Russell's theory discussed thus far have been mostly of a semantic or metaphysical nature, but there are reasons to believe that Russell's primary motivations were epistemological.³ For example, one of Russell's many key insights was that speakers can grasp the meaning

³ Some of these reasons can be found in Russell (1917).

of sentences of the form ‘the F is G ’ without knowing, i.e. without being acquainted with, the object that the description ‘the F ’ denotes. Moreover, as already mentioned, Russell noted that speakers in general seem perfectly capable of grasping the meaning of sentences of the form ‘the F is G ’ even when they are not (and could not possibly be) acquainted with its denotation (for example, when the description fails to denote). At the same time, Russell maintained that understanding the meaning of a sentence containing a genuinely singular term, i.e. a referential expression, required acquaintance, he was led to the conclusion that descriptions such as ‘the F ’ are not referential expressions and that sentences of the form ‘the F is G ’ express *object-independent* propositions. With the theory of descriptions, Russell had a theory which straightforwardly captured these epistemological considerations.

But because Russell insisted that object-dependent thoughts strictly require the existence of actual objects — and hence that seemingly object-dependent sentences such as ‘Pegasus flies’ or ‘Vulcan is a planet’ cannot give rise to object-dependent thoughts — Russell was eventually led to the now widely considered *extreme* view that only sense data can serve as objects of singular thoughts. And since even proper names can fail to refer, Russell concluded that proper names were simply covert descriptions. This assumption was also necessary if the theory of descriptions was to achieve one of its stated goals, namely to provide a solution to Frege’s Puzzle.⁴

Much more could be said about the motivations behind Russell’s theory, but this should suffice for my purposes. Summarizing, Russell’s analysis is widely considered to have the following significant explanatory advantages.

- B1** The theory provides a simple and elegant way of assigning classical truth values to sentences containing non-denoting descriptions such as (7) and (9) without committing us to intuitively implausible metaphysical assumptions (à la Meinong) or infringing the law of excluded middle.

⁴ Frege’s puzzle is the problem of explaining the difference in cognitive significance between statements of the form $a = a$ and $a = b$, cf. Frege (1892) and when the assumption that proper names are simply covert definite descriptions is made, this puzzle has a straightforward solution. Yet due to the pioneering work of Kripke (1980), this aspect of Russell’s theory is now almost unanimously considered incorrect, and for that reason I will not discuss this aspect of Russell’s theory here. However, this is not to imply that absolutely no one endorses Russell’s descriptivist thesis about proper names. For some recent attempts to defend this view, cf. Geurts (1997), Elbourne (2005), and Hawthorne and Manley (forthcoming).

- B2 The theory straightforwardly explains how negative existentials such as (9) can be true.
- B3 The theory yields a simple explanation for the ambiguities that arise when definite descriptions interact with other scope-taking expressions, e.g. modals, if-clauses, quantifiers, quantificational adverbs, and various logical operators. In particular, it provides an explanation of *de dicto* and *de re* ambiguities.
- B4 The theory explains why certain attitude attributions containing non-denoting descriptions are nevertheless intuitively true.
- B5 The theory explains how speakers can grasp the meaning of the *F* is *G* without being acquainted with the individual denoted by 'the *F*'.

This list hardly exhausts the aspects of Russell's theory which are generally considered virtues, but these are among the more significant ones. Nevertheless, several of these have been seriously questioned. One of the most famous objections to Russell's theory due to Strawson (1950) was essentially an argument against B1 (and this objection is briefly outlined in the next section and extensively discussed in Chapter II.) As for B3, it was more or less conclusively shown by Fodor (1970) and Bäuerle (1983) that Russell's analysis of *de dicto* and *de re* ambiguities in terms of scope is insufficiently powerful, i.e. simple scope permutations are simply not going to be enough to capture the relevant readings.⁵ B4 is also still widely considered a virtue of Russell's theory and while Heim (1991), Kripke (2005), Elbourne (2005, 2010) have done much to challenge this assumption, their objections seem to have been either ignored or outright dismissed. Chapter III is devoted to showing that this is a mistake and that the scope of the problems arising from these cases have been seriously underestimated.

2 Standard Objections to Russell's Analysis

Over the years, numerous objections of varying strength have been raised against Russell's theory. The following is a brief outline of the most influential objections as these are relevant to the remaining chapters.

⁵ For discussion and some recent advanced approaches to the analysis of *de dicto* and *de re* ambiguities, cf. Musan (1995), Percus (2000), Maier (2006, 2009), von Stechow (2007), Keshet (2008), Romoli and Sudo (2009).

2.1 The Mismatch between Syntax and Semantics

Russell's analysis of descriptions is *syncategorematic*, i.e. it captures the meaning of definite and indefinite descriptions in terms of the unrestricted unary first-order quantifiers '∃' and '∀' and assigns the descriptions no meaning in isolation. Rather, definite and indefinite descriptions have meaning only when occurring in fully well-formed grammatical sentences. Several critics have therefore pointed out that there is a blatant (and implausible) lack of connection between surface form and logical form on Russell's analysis, cf. (3) and (4). These critics also often argue that Russell's syncategorematic analysis appears to violate the intuitively plausible principle that the meaning of a sentence is a product of the meaning of its parts and their syntactic combination, viz. the principle of compositionality.

While this criticism of Russell's theory is serious, it is today widely agreed that there is a straightforward response. During roughly the past four decades, research on natural language quantifiers has led most researchers in philosophy of language and natural language semantics to believe that natural language quantifiers are so-called *generalized* quantifiers — or more specifically *restricted* quantifiers.⁶

There is a vast amount of research on the properties of generalized quantifiers and I am unable to survey this research here. In short, one of the main arguments in favor of analyzing various natural language determiners

⁶ So-called generalized quantifier theory owes much to the pioneering Mostowski (1957), and extensions specific to natural language quantifiers were chiefly developed by Barwise and Cooper (1981), Higginbotham and May (1981), Keenan and Stavi (1986). Roughly speaking, a quantifier is a restricted quantifier when it denotes a relation between two sets (a so-called type (1,1) quantifier) and it obeys the principle of conservativity. The principle of conservativity states that for each $A, B \subseteq D$ (where D is the domain of the model), a quantifier Q is conservative if only if:

$$\text{CONSERVATIVITY: } Q(A, B) \leftrightarrow Q(A, A \cap B)$$

When a quantifier Q obeys conservativity, its quantificational nature is restricted in the following sense: to determine whether a sentence of the form $Q(A, B)$ is true, one need only inspect the set A , i.e. one need not quantify over the entire universe of individuals. Natural language determiners have in general been shown to obey this principle. But, that quantifiers of natural language are restricted is also supported by syntactic evidence, i.e. it seems that a quantificational determiner (D) combines syntactically with a common noun phrase (CNP or just NP) to form one syntactic and semantic unit, a determiner phrase (DP). This unit then combines with a verb phrase (VP) to form a sentence. As a result, proponents of Russell's analysis today more or less uniformly assume that definite and indefinite descriptions should be analyzed in terms of restricted, rather than unrestricted, quantification.

as restricted quantifiers is that numerous determiners of natural language are simply not amenable to semantic analysis in terms of the unary quantifiers ‘ \exists ’ and ‘ \forall ’ — for example, determiners such as ‘most’, ‘more than half’, and ‘few’.⁷ But these determiners *are* amenable to an analysis as restricted quantifiers where their semantic values are treated as relations between sets.

Hence, the consensus today among proponents of Russell’s theory is that Russell’s theory of descriptions should be subsumed under generalized quantifier theory so that definite and indefinite descriptions are both treated as genuine syntactic and semantic units. In order to indicate this, their logical forms are typically represented as in (15) and (16) respectively, and their truth conditions are then stated in terms of set theoretic relations.

- (15) $[\text{An } x: Fx](Gx)$ is true iff $|F \cap G| \geq 1$
 (16) $[\text{The } x: Fx](Gx)$ is true iff $|F \cap G| \geq 1$ and $|F| = 1$

When definite descriptions are assigned these truth conditions, the resulting semantic analysis is almost identical to Russell’s. For example, as on Russell’s original analysis, the conditions A1-A3 outlined above are satisfied; ‘the F is G ’ asserts that there is a *unique* F , viz. the cardinality of F equals one, and that the individual, who is F , is G . And again, as witnessed by the truth conditions for (15), the uniqueness constraint is simply an extra condition that is easily discarded. The close relation between definites and indefinites is thus preserved. In conclusion, with restricted quantifiers we have an analysis that appears to capture every essential aspect of Russell’s theory yet maintains a neat correspondence between syntax and semantics and allows for easy integration into a compositional semantic theory.⁸ Throughout this dissertation, I stick to the standard unary quantifiers of first-order logic in my formalizations, but this is a completely innocuous simplification. Everything I argue could be argued just the same if definite and indefinite descriptions were treated as restricted quantifiers.

2.2 Presuppositions and Truth Value Gaps

One of the most famous objections to Russell’s theory was put forward by Strawson (1950). While Strawson seemed in principle opposed to the idea

⁷ For discussion see e.g. Wiggins (1980), Neale (1990), Ludlow (1995).

⁸ This is oversimplifying a bit as the integration of quantifiers into a compositional semantic theory remains a rather complex and controversial issue. Useful discussions of the various problems are available in Heim and Kratzer (1998: 178-200), Dever (2005), and Glanzberg (2006). However, the relevant complexity here is thoroughly general and so does not pertain exclusively to Russell’s theory.

of analyzing natural language in terms of an artificial formal language such as first-order logic, his main argument against Russell in 'On Referring' is that the predictions of Russell's theory — in particular as these regard non-denoting descriptions — are inconsistent with our immediate truth value judgments.

Now suppose some one were in fact to say to you with a perfectly serious air: "The king of France is wise". Would you say, "That's untrue"? I think it's quite certain that you wouldn't. But suppose he went on to *ask* you whether you thought that what he had just said was true, or was false; whether you agreed or disagreed with what he had just said. I think you would be inclined, with some hesitation, to say that you didn't do either; that the question of whether his statement was true or false simply *didn't arise*, because there was no such person as the king of France. You might, if he were obviously serious (had a dazed astray-in-the-centuries look), say something like: "I'm afraid you must be under a misapprehension. France is not a monarchy. There is no king of France." (Strawson 1950: 330)

Strawson observes that we are not immediately inclined to judge that an out of the blue assertion of 'the king of France is wise' is plainly false, and hence the prediction of Russell's theory that these sentences are straightforwardly false looks empirically inadequate. Strawson is thus effectively challenging A1 above, the assumption that sentences of the form 'the *F* is *G*' somehow *assert* the existence of a unique *F*. If these sentences did in fact assert existence, it should be perfectly felicitous to simply respond 'that's not true' to such assertions. However, as Strawson points out, this is somewhat questionable.

Instead of an assertion of existence, Strawson maintained that when speakers use a definite description, they are instead *presupposing* that the description has a unique denotation. Strawson continues,

[...] to use an expression such as "The king of France" at the beginning of a sentence was, in some sense of "imply", to imply that there was a king of France. When a man uses such an expression, he does not *assert*, nor does what he says *entail*, a uniquely existential proposition. But one of the conventional functions of the definite article is to act as a *signal* that a unique reference is being made — a signal, not a disguised assertion. When we begin a sentence with "the such-and-such" the use of "the" shows, but does not state, that we are, or intend to be, referring to one particular individual of the species "such-and-such". *Which* particular individual is a matter to be determined from context, time, place and any other features of the situation of utterance. Now, whenever a man uses any expression, the presumption is that he thinks he is using it correctly: so when he uses the expression, "the such-and-such", in a uniquely referring way, the presumption is that he thinks

both that there is *some* individual of that species and that the context of use will sufficiently determine which one he has in mind. To use the word “the” in this way is then to imply (in the relevant sense of “imply”) that the existential conditions described by Russell are fulfilled. But to use “the” in this way is not to *state* that those conditions are fulfilled. (Strawson 1950: 331-332)

Strawson thus maintained that definite descriptions are expressions that speakers *use* to *refer* and that such uses require that the speaker is *presupposing* that he is referring to something, viz. that there is an individual who satisfies the descriptive component. Thus, on Strawson’s view, a sentence containing a non-denoting description simply fails to express a proposition, and such sentences are therefore neither true nor false, i.e. the question of truth and falsity simply *fails to arise* for such sentences.

Since Frege (1892) also maintained that uses of ‘the *F*’ presuppose a reference and that ‘the *F* is *G*’ fails to express a proposition when the presupposition is not satisfied, the presuppositional view of definite descriptions is now standardly referred to as the Frege/Strawson analysis. Proponents of this analysis thus in general hold that the putative advantage of Russell’s view stated in B1 above is actually an empirically questionable consequence rather than a virtue. I discuss the Frege/Strawson analysis in more detail in Chapter III.

Strawson’s objection to Russell’s theory initiated a debate about the relative value of truth value judgments. While I think a majority agrees with Strawson’s intuitive judgment about cases such as (7), the problem is that there are numerous sentences containing non-denoting descriptions that *do* elicit robust judgments of falsity which thus conflict with Strawson’s general view. Strawson was, or at least became, quite aware of this problem and it led him to revise his position, cf. Strawson (1964).

A few people continue to maintain that the empirical data actually favors the Russellian analysis and they have dismissed Strawson’s objections on those grounds, e.g. Neale (1990). But, most people seem to have concluded from this debate that truth value intuitions are simply too unstable or unreliable to shed any significant light on the semantics of definite descriptions, and moreover that other types of data are required to settle this dispute. This question is the central topic of Chapter II.

2.3 Referential and Attributive Uses of Definite Descriptions

In the seminal ‘Reference and Definite Descriptions’ (1966), Donnellan argued that there are two distinct uses of definite descriptions, namely *attributive* and *referential* uses. On the standard attributive use, the speaker

uses the description to assert a general proposition about an individual whom the speaker quite possibly is not acquainted with. But on the referential use, the speaker uses the description to communicate a singular proposition, i.e. the speaker has a particular individual in mind and the description then only serves as a descriptive guide to the individual in question.

As evidence for the referential uses, Donnellan relied on cases involving misdescriptions, i.e. cases where the description in actual fact fails to denote but where communication nevertheless succeeds. For example, it is easy to imagine cases where both speaker and interlocutor are under the misapprehension that some particular individual *a* is the unique *F*, and where the speaker uses a description of the form 'the *F*' in an effort to *refer* to *a*. Since communication appears to proceed smoothly in such cases and since it seems that the speaker can often be interpreted as having asserted something *true* about *a* in such cases, Donnellan concluded that these cases are semantically significant. That is, Donnellan concluded that sentences of the form 'the *F* is *G*' can be true even when there is no unique *F*.

This presents an obvious difficulty for the standard Russellian analysis, because any case involving a misdescription is, on the Russellian analysis, predicted to be false. Donnellan thus concluded that definite descriptions in English are ambiguous between referential and attributive uses.⁹

Donnellan's objection has received an overwhelming amount of attention and it is discussed in numerous papers.¹⁰ Recently, researchers such as Devitt (2004, 2007) and Reimer (1998) have offered a new argument in favor of a semantic distinction between referential and attributive uses, sometimes referred to as *the argument from convention* or *the argument from regularity*. This argument is the topic of Chapter IV, where I attempt to show that it is flawed and that it does not provide grounds for thinking that there is a genuine semantic distinction between referential and attributive uses.

⁹ Donnellan's view is explicated in more detail in Chapter IV.

¹⁰ The most famous response is probably Kripke's (1977) 'Speaker's Reference and Semantic Reference', but for other discussions of Donnellan's distinction, cf. Barwise and Perry (1983), Devitt (1981b), Devitt and Sterelny (1997), Wettstein (1981, 1983), Wilson (1978), Schiffer (1998, 2005), Sainsbury (1979), Searle (1979), Bach (1981, 1987, 2004, 2007), Evans (1982), Neale (1990), Ludlow and Neale (1991) Salmon (1982, 1991), Soames (1986, 1994). Donnellan's paper actually gave rise to a similar discussion about referential uses of *indefinites*, mostly as a result of a now famous paper by Fodor and Sag (1982). Fodor and Sag's paper also initiated a comprehensive discussion of so-called specific indefinites which has received a great deal of attention, especially in linguistics. See e.g. Ludlow and Neale (1991), Ruys (1992), Abusch (1994), Reinhart (1997), Kratzer (1998), Winter (2001), Schwarzschild (2002).

2.4 Incomplete Definite Descriptions

The final objection to Russell's theory of descriptions that I discuss here is the problem of incomplete definite descriptions. This problem was also raised by Strawson but it has subsequently been stressed to a greater degree by other researchers.¹¹ Strawson writes,

Now it is obviously quite false that the phrase "the table" in the sentence "the table is covered with books", used normally, will "only have an application in the event of there being one table and no more". It is indeed tautologically true that, in such a use, the phrase will have an application only in the event of there being one table and no more *which is being referred to*, and that it will be understood as being used to refer to. To use the sentence is not to assert, but it is (in the special sense discussed) to imply, that there is only one thing which is *both* of the kind specified (*i.e.* a table) *and being referred to* by the speaker. (Strawson 1950: 332-333)

The problem here concerns Russell's uniqueness constraint (A2 above), viz. the constraint that 'the *F* is *G*' asserts that there is a *unique* individual who instantiates the property *G*. As Strawson points out, it seems that descriptions of the form 'the *F*' can unproblematically be used in communication even when it is common ground that there is more than one single individual/object that instantiates the relevant property. In other words, speakers can succeed in asserting something true even when the uniqueness constraint is clearly not satisfied.

Some researchers, e.g. Strawson (1950), Donnellan (1966), Devitt and Sterelny (1999), took these types of problems to bolster the case for the existence of referential descriptions, but see in particular Neale (1990) for a response. Other researchers have argued that uniqueness as a semantic constraint should simply be abandoned, e.g. Szabó (2000), Ludlow and Segal (2004), and that uniqueness is merely a pragmatic implication of definite descriptions. Others still have argued that incomplete definite descriptions are nothing but instances of a much more general phenomenon that arises with basically every other determiner in natural language, namely the phenomenon of quantifier domain restriction. As Neale writes,

[...] incompleteness affects quantifiers more generally, not just definite descriptions. And to that extent, appeals to contextual coordinates and ellipsed material are independently required by any adequate theory of natural language quantification. (Neale 1990: 101)

¹¹ See e.g. Wettstein (1981), Salmon (1982), Soames (1986), Neale (1990).

I.e. on standard uses of other natural language quantifiers, we also encounter this type of apparent incompleteness. For example, it is easy to imagine the (a)-sentences below used to communicate what is literally expressed by the (b)-sentences.

- (17) a. Every student passed the exam.
b. Every student *in my class* passed the exam.
- (18) a. Only three students passed the exam.
b. Only three students *in my class* passed the exam.
- (19) a. More than half of the students passed the exam.
b. More than half of the students *in my class* passed the exam.

The topic of quantifier domain restriction has in recent years been given a great deal of attention, for example in Stanley (2002, 2000) and more extensively in von Stechow (1994).

3 Concluding Remarks

This concludes my introduction to Russell's theory of descriptions and the overview of the standard objections raised to Russell's theory. The objections discussed above are of course far from exhaustive, but the objections mentioned are surely among the most influential. And since several of the chapters to follow discuss aspects of these objections, I felt it was important to at least provide a crude overview. The topics discussed in the remainder are as follows.

Chapter II — Here I discuss the problem raised by Strawson concerning truth value intuitions and non-denoting descriptions. I present a model which explains this variations in intuitions and I argue that when the data is properly analyzed, it favors a presuppositional analysis.

Chapter III — Here I discuss the putative virtue of Russell's theory explicated in B4 above. I argue that the problems raised by Heim, Kripke, and Elbourne's are serious, but show that their counterexamples do not clearly provide evidence for the Frege/Strawson analysis (as Heim and Elbourne maintain). I then propose a dynamic semantic analysis of definites and argue that this analysis is the way forward if the relevant problems are to be solved.

Chapter IV — Here I discuss an argument the *argument from regularity* mentioned above. I show that this argument severely overgenerates

and as such provides no convincing reasons for positing a semantic ambiguity. I then outline in rough details how the dynamic analysis can utilize various resources to explain referential uses.

Appendix A — Here I survey relevant data concerning indefinite descriptions and presuppositions relevant to the analysis proposed in Chapter III.

CHAPTER II

DESCRIPTIONS, TRUTH VALUE INTUITIONS, AND QUESTIONS

1 Introduction

In contemporary philosophy of language and linguistics, two mutually incompatible semantic analyses of definites are predominant. These are of course Russell's (1905) famous quantificational analysis and the presuppositional analysis as defended by Strawson (1950, 1952, 1964).¹ The central point of disagreement between these analyses is whether definites should be analyzed as asserting or merely presupposing the existence of an individual satisfying the relevant description. In the profusion of spirited defenses of these analyses, appeals to truth value intuitions are fairly common. For proponents of a Russellian analysis, the intuition of falsity elicited by a sentence such as (20) is often considered an argument in favor of an existential analysis — an analysis on which sentences containing non-denoting descriptions are assigned a classical truth-value. In contrast, for proponents of a Strawsonian analysis, the fact that no clear truth value intuition is elicited by sentences such as (21) is considered an argument in favor of a presuppositional analysis — an analysis on which sentences containing non-denoting descriptions are undefined, viz. neither true nor false.

(20) The king of France shot my cat last night.

Neale (1990)

¹ With the important precursors to the presuppositional analysis appearing in Frege (1892, 1918).

(21) The king of France is bald.

Strawson (1964)

The immediate problem is that these variations in truth value intuitions are pervasive and as a result many researchers today maintain that (a) truth value intuitions are too unstable to provide decisive support for either analysis, (b) that providing an explanation of these pervasive variations is essentially intractable, and (c) that appeals to truth value intuitions merit no serious attention. In other words, truth value intuitions have acquired a bad reputation.²

But simply disregarding these putative semantic intuitions is not obviously feasible. Both the existential and the presuppositional analyses provide uniform treatments of no-reference cases, and as a result a wide variety of cases are left unexplained on both analyses. This is unsatisfactory for several reasons, but one particular reason is that an important factor in assessing the adequacy of a semantic analysis is its predictive capabilities. Yet, given that both the existential and the presuppositional analyses give uniform treatments of no-reference cases, the result is numerous incorrect predictions.

One way to resolve the problem of apparently conflicting intuitions is to take the existence of these fragile, highly context-sensitive, and frequently unstable intuitions as constituting a direct argument against both aforementioned analyses. This option amounts to the claim that only a non-uniform analysis of definites, i.e. an analysis which is capable of assigning both classical and non-classical truth values to sentences containing non-denoting definites, will be adequate, see e.g. Lasersohn (1993). However, another option, available to both proponents of Russell's theory and proponents of Strawson's theory, is to combine the analyses with a systematic explanation of the problematic cases, e.g. an explanation that justifies disregarding certain parts of the data. Such an approach would be successful if it was effectively demonstrated that the problematic parts of the putative data are semantically irrelevant.

² Occasionally, the problem is categorically dismissed. For instance, in his book *Descriptions*, Neale (1990) writes.

[...] all sorts of factors may conspire to deter the native speaker from *saying* that a given utterance is true or false, but that is hardly enough to show that the utterance lacks a truth-value. (Neale 1990: 28)

While Neale is right that the existence of cases which elicit no robust intuitions is insufficient for concluding that a given sentence lacks a truth value, Neale neglects to consider which factors might deter the native speaker from saying that a given utterance is *neither true nor false*.

While investigating the behavior of truth value intuitions is in itself instructive, this chapter is an attempt to explore the second option. In particular, it is an attempt to explain why truth value intuitions exhibit such extensive variations and in turn evaluate to what extent a semantic analysis of definites can and should rely on such intuitions. The aim is to show that in order to make sense of these recalcitrant truth value intuitions, an alternative perspective on the putative data, so far neglected in the debate, is required. The upshot of my analysis is that the fairly commonplace intuitions of falsity associated with sentences containing non-denoting definites quite often result from factors which are not semantically relevant. This yields, *ceteris paribus*, reasons to favor a presuppositional analysis of definites.

The hypothesis here is that intuitions of falsity elicited from sentences suffering from existential presupposition failure are a result of evaluating the content of the sentences in relation to a contextually salient question. In particular, I show that there is an important relation between robust intuitions of falsity and the satisfaction of certain conversational conditions. I argue that the existence of this relation puts into question whether these intuitions track the semantic status of the sentences in question. The proposal here is essentially an attempt to revive and systematize important insights made by Strawson (1964) while avoiding certain problematic consequences of his proposed solution.

I begin this chapter by briefly illustrating the complexity of the phenomenon to be explained. The data shows quite vividly that there is no simple or obvious systematic pattern which provides an easy explanation of the variance in intuitions. Next, I briefly discuss a couple of recent attempts to explain the behavior of these truth value intuitions and argue that these proposals are problematic in various respects. In light of these proposals, I contend that to gain a better understanding of the behavior of truth value intuitions, we must first consider their behavior in context, namely their sensitivity to discourse topics and to salient questions. Adopting Roberts' (1998, 2004) framework for discourses, I show how intuitions are systematically sensitive to context, and that there is an important connection between intuitions of falsity and, what I label, conversational consonance. Subsequently, I study what happens when non-denoting definites interact with prosodic focusing. The effects of prosodic focus are perplexing and a neglected phenomenon in recent discussions of the problem. However, I show that the present proposal provides a straightforward explanation of these effects. Finally, I evaluate the explanatory adequacy of the proposal.

2 The Intuitive Pull

First some useful terminology: I use the abbreviation SWEPPF for *Sentences With Existential Presupposition Failure* but this is interchangeable with e.g. *sentences containing non-denoting definite descriptions* or similar. I use the term *robust truth value intuitions* (or sometimes just *robust intuitions*), to refer to intuitions of falsity or intuitions of truth, however I focus almost exclusively on cases which elicits robust intuitions of falsity. I use Strawson's term *squeamish* to refer to SWEPPFs that elicit no robust intuitions, i.e. *squeamish SWEPPFs*. These are SWEPPFs where one feels an immediate inability to judge the sentence either true or false — i.e. the SWEPPF elicits a feeling of squeamishness.

2.1 The Data

The paradigm case of a SWEPPF which elicits squeamishness is Strawson's famous example (21).

- (21) The king of France is bald. Strawson (1964)

An immediate inability to evaluate this sentence for truth or falsity is a common reaction — the reaction that is normally reported in the literature. Confronted with (21), there is an intuitive sense that the speaker has made a mistake and that one cannot simply dismiss the sentence as false. In other words, it appears to elicit no robust intuitions. However, a small tweak of the content in (21) seems to make intuitions change.

- (22) The king of France is a bald nazi.
(23) The king of France is a skinhead.

These SWEPPFs differ from (21) in that they seem to elicit a much stronger sense of falsity and there is thus a marked contrast between (21) and (22)-(23).

Let's consider another contrast pair. The SWEPPFs in (24) and (25) are standard examples of SWEPPFs which elicit robust intuitions of falsity. Assume that the demonstratives below pick out contextually salient objects.

- (24) The king of France is sitting in that chair. Lasersohn (1993)
(25) The king of France owns that pen. von Fintel (2004)

While these SWEPPFs elicit robust intuitions of falsity, a change in intuition is again facilitated by a seemingly insignificant change of the content, e.g. (26-27).

- (26) The king of France is sitting in a chair.
(27) The king of France owns a pen. von Fintel (2004)

If these sentences do not elicit straightforward squeamishness, I think the pull towards squeamishness is certainly non-negligable. If so, there is again a distinct contrast between (24-25) and (26-27).

From the point of view of semantic analysis, the above SWEPPs seem to suffer from the same defect, namely that the definite description fails to denote. One would therefore *prima facie* expect our intuitions to be uniform, yet they are clearly not. There are numerous contrast cases such as the above and below I provide a select sample of relevant cases. Since the data is somewhat fine-grained, contrast cases are particularly useful for demonstrating the peculiar variations in intuitions. What is important at this point is to consider the intuitive pull of the sentences. These minimal pairs show how sentences which are structurally quite similar nevertheless pull intuitions in different directions.³

SWEPPs (SENTENCES WITH EXISTENTIAL PRESUPPOSITION FAILURE)

- (21) # The king of France is bald.
(22) ^F The king of France is a bald nazi.
- (26) # The king of France is sitting in a chair.
(24) ^F The king of France is sitting in that chair.
- (27) # The king of France owns a pen.
(25) ^F The king of France owns that pen.
- (28) # The king of France is on a state visit.
(29) ^F The king of France is on a state visit to Australia this week.
- (30) # The king of France heard about Goldbach's conjecture.
(31) ^F The king of France proved Goldbach's conjecture.
- (32) # The king of France read Anna Karenina.
(33) ^F The king of France wrote Anna Karenina.
- (34) # The king of France always enjoys a croissant in the morning.

³ In this chapter, I follow von Fintel's (2004) convention of using '#' to indicate feelings of squeamishness and a superscripted 'F' to indicate intuitions of falsity. Several of these examples are adapted from the literature on truth value intuitions, i.e. Lasnik (1993), von Fintel (2004), and Yablo (2006, 2009).

- (35) ^F The king of France always enjoys some human flesh in the morning.
- (36) # The king of France walked his dog last night.
- (37) ^F The king of France ate his dog last night.
- (38) # The king of France is BALD.⁴
- (39) ^F The king of FRANCE is bald.
- (40) # The king of France is on a STATE visit.
- (41) ^F The king of FRANCE is on a state visit.
- (42) # The king of France is bald.
- (43) ^F Only the king of France is bald.

Briefly surveying the above contrast pairs makes perfectly clear the complexity of the problem. Let's consider some of the apparently intuition-changing causes. In (21) the change in intuitions appears to be caused by the addition of the noun 'nazi'. However in (24-26) and (25-27) the intuition-changing factor is the difference between the indefinite determiner and the demonstrative. In (30-31) and (32-33) intuitions change due to a mere substitution of verbs, whereas in (38-39) and (40-41) intonational stress forces an intuition of falsity. In (43), the focus-sensitive particle 'only' has a similar effect. In light of these cases, it seems that hopes of finding some simple structural explanation are quickly dashed. While it is possible to formulate somewhat intuitive but ad hoc explanations for isolated cases, providing a uniform explanation is a real challenge. Yet, if our semantic analysis is to capture the data and refrain from making incorrect predictions, a uniform explanation is clearly preferable, if not required.

2.2 *Accommodation Theories*

In recent years, a couple of interesting proposals to explain the behavior of truth value intuitions and their significance to semantics have emerged, namely Lasersohn (1993), von Stechow (2004), Yablo (2006, 2009).⁵ While

⁴ The capitals here indicate a pitch accent.

⁵ There are important differences between these three analyses, but the objections I discuss here are sufficiently general to apply to all three proposals. Two papers that also deserve mentioning are Glanzberg (2005) and Geurts (2007). Glanzberg's paper is a proposal to characterize the conditions under which a sentence fails to express a proposition and it touches briefly on the topic of truth value intuitions and presupposition failure.

these theories are different in various respects, they converge on a general idea. The common contention is that truth value intuitions in no-reference cases are a result of a mechanism of temporary accommodation of the false presupposition. When evaluating SWEPFs, evaluators (i.e. individuals consulting their intuitions) employ a temporary accommodation of the presupposition to permit an evaluation of the content of the SWEPF. The hypothesis is that if the result of accommodating the presupposition is a proposition which the speaker is in a position to falsify by deriving a contradiction, the sentence elicits an intuition of falsity. Conversely, if the speaker is incapable of falsifying the proposition, e.g. not in a position to derive a contradiction, the sentence is predicted to elicit no robust truth value intuition. On Lasersohn's (1993) and von Fintel's (2004) proposals, this basic idea is formally explicated in terms of belief revision procedures. The false presupposition is added to the evaluator's body of beliefs and the sentence under evaluation is tested against this body of beliefs.⁶ If the evaluator is in a position to derive an inconsistency after (a) accommodating the false presupposition and (b) adding the propositional content of the relevant sentence, the SWEPF is judged false. If the evaluator is not in a position to derive an inconsistency, the SWEPF is predicted to elicit squeamishness.⁷

Glanzberg's general contention is that truth value intuitions are too unstable to provide solid data for theorizing and he suggests instead to rely on a range of discourse-based tests, so-called *repair tests*. This line of reasoning is similar to the proposal I present here, but since Glanzberg's paper is not an attempt to explain the peculiar behavior of truth value intuitions, it engages with a very limited number of examples. Geurts (2007) defends a Strawsonian analysis of strong determiners and contends that Strawson's view of presuppositions should be assimilated to the binding theory of presuppositions, cf. van der Sandt (1992). In this chapter I also defend a broadly Strawsonian analysis, but I do not discuss Geurts' proposal since it fails to engage with, and cannot explain, a number of the cases introduced in e.g. von Fintel (2004). Moreover, since Geurts' proposed analysis is quite similar to Strawson's, I believe it falls prey to a number of the problems that afflicts Strawson's original proposal. These problems are discussed in von Fintel's paper. So, while I am sympathetic to the overarching idea in Geurts' paper, I believe the proposal presented here provides a more satisfactory analysis.

⁶ Providing a detailed account of the relevant belief revision procedure faces a number of challenges. For instance, bodies of beliefs must be closed under some suitable consequence relation, but adequately characterizing that consequence relation is obviously no trivial task. Nevertheless, to attain correct predictions, an adequate characterization is required. The most comprehensive attempt to provide an adequate characterization of the belief revision procedure is given in von Fintel (2004). I will ignore problems related to this issue for the remainder of this chapter.

⁷ The crucial differences between these proposals lie in the explications of these revision procedures. However, the problems related to these approaches are sufficiently general,

Consequently, these proposals share the basic assumption that truth value intuitions must be explained in terms of a verification/falsification procedure.

While these proposals are capable of explaining considerable parts of the data, I believe their general strategy is problematic and I will briefly discuss what I consider the two primary problems. First, there is a class of cases for which these proposals yield systematically incorrect predictions and these incorrect predictions stem solely from the limitations of the strategy, namely the strategy where predictions rely on the derivability of inconsistencies. The existence of this class puts into question whether a verification/falsification strategy is equipped to explain how truth value intuitions are in general generated. Let's begin by considering a number of the SWEPPs listed above. Using the verification/falsification procedure, the accommodation theories make correct predictions for the paradigm cases (21) and (25). Adding to our stock of beliefs that there is a king of France and that he is bald, gives rise to no inconsistency and hence is predicted to elicit squeamishness. Conversely, adding to our stock of beliefs that there is a king of France and that he is sitting in a contextually salient chair *does* give rise to an inconsistency; either the relevant chair is empty or it is occupied by someone who we know is not the king of France. And so, doing simple derivations on our body of beliefs, we quickly derive a contradiction. This is the desired result, since (25) elicits an intuition of falsity.⁸

However, when we start considering the additional cases, these proposals run into problems since a number of the listed contrast cases cannot be explained. For instance, the stark contrast between (21) and (23-22) is undetected by the verification/falsification analysis; while there is no way of falsifying that the king of France is bald (i.e. deriving a contradiction), there is also no way of falsifying that he is a bald nazi. Analogous results obtain for the contrast pairs (34-35) and (36-37). To explain the contrast between e.g. (21) and (22), the accommodation theories must place heavy demands on bodies of beliefs. They must assume that bodies of beliefs in general contain a proposition which is either directly inconsistent with (22) or entails a proposition that is inconsistent with (22). To capture the contrast, they must therefore either attribute to speakers a general belief such as 'no

so that we can ignore these differences here. The proposals also diverge importantly in their assessment of the semantic significance of these intuitions, but I will also ignore these differences here.

⁸ It should be noted that von Fintel's (2004) revision procedure is more sophisticated than e.g. Lasnik's (1993) which allows him to capture more data. But, the cases adduced below are problematic for both Lasnik, von Fintel, and Yablo.

kings are nazis' or the more odious belief 'the king of France is not a nazi'. Now, I assume that it is not feasible to ascribe the latter belief to speakers in general and as regards the former, viz. the general belief that no kings are nazis, the question arises whether it is reasonable to assume that speakers in general have such a belief. This is, to say the least, unclear. However, simply assuming that they do, only because (22) elicits an intuition of falsity, deprives the proposals of explanatory power. In that case we are attributing beliefs to people simply to attain correct predictions, but if the proposals are to *explain* intuitions, we need independent justification for the assumption that speakers in general have such beliefs. It is rather difficult to see what would justify that assumption.

Also, since these analyses rely on verification/falsification, they rely importantly on facts – facts that are in general known. But this means that if the facts required to falsify a sentence are not known (or on von Fintel's account, something that could not in principle be known), the sentence is predicted to elicit a sense of squeamishness. As a result, any sentence which ascribes e.g. a non-factive mental state (for example a non-factive propositional attitude) to a non-denoting definite is predicted to elicit squeamishness. There are of course no facts about what, say, the king of France's mental states are, and as a result we have no beliefs about the mental states of the king of France (presumably because he does not exist). Consequently, any sentence of that form should elicit a sense of squeamishness. Nevertheless, constructing examples of that form that elicit robust intuitions of falsity is an easy task. Consider (44-46).

- (44) ^F The king of France hates your mother.
- (45) ^F The king of France is planning a nuclear attack on the US.
- (46) ^F The king of France wants to steal your car.

These sentences seem to me to elicit robust intuitions of falsity, but because of the inherent limitations of the verification/falsification strategy, the accommodation theories cannot predict this. These results constitute a real challenge to the claim that we are somehow relying on a tacit verification/falsification procedure when generating these intuitions. It is just not clear that verification/falsification is at issue here.

This brings us to what I consider the other primary problem with the accommodation theories. On these analyses, there is no explanation provided for the fragility of truth value intuitions. Many SWEPFs elicit no robust intuitions when considered in isolation, but suddenly elicit quite robust intuitions of falsity when considered in context. In other words, truth value intuitions are delicate and easily affected by contextual factors. Any analysis attempting to explain the behavior of these delicate intuitions that

nevertheless fails to explain their significant sensitivity to context is bound to be afflicted with predictive shortcomings. And since the accommodation theories have no resources to explain such variations, there are numerous cases of SWEPPs whose behavior in context remains but a mystery. Given a SWEPP and a relevant body of beliefs, the accommodation theories make a firm prediction, and as long as the context engenders no change in the body of beliefs, this prediction remains the same. But, as I will demonstrate in the next section, context can engender a change in truth value intuitions without altering in any substantial way the relevant body of beliefs.⁹ The accommodation theories are at a loss to explain this kind of variation.

I submit that if we are to have any hope of sorting out this mess, we must focus our attention elsewhere, viz. away from SWEPPs evaluated in isolation. If we want to gain a better understanding of truth value intuitions in regards to SWEPPs, we must examine their behavior in context.

2.3 *Systematizing the Strawsonian Insights*

My proposal is very much an elaboration of ideas put forward in Strawson's seminal (1964) paper. My analysis is in a sense an attempt to systematize Strawson's main insights while avoiding the unfortunate consequences of his proposed explanation. Strawson originally cast his explanation of varying truth value intuitions in terms of a distinction between topic and focus, and he conjectured that intuitions of falsity are elicited only when non-denoting definites are not in topic position. He therefore proposed to analyze definites in focus position as non-presuppositional, since this would allow these sentences to be assigned classical truth values in accordance with our purported intuitions.

But Strawson's proposal is problematic on several counts. First, it is convincingly argued in von Fintel (2004) that definites are presuppositional as both topics and focus.¹⁰ Secondly, on Strawson's proposed method for identifying sentence topics, his proposal makes a number of incorrect predictions. Many sentences elicit strong intuitions of falsity even though the non-denoting definite is in, what Strawson considered, topic position.¹¹

⁹ Relevant cases here include (CI) and (CII).

¹⁰ Cf. the 'Hey, Wait a Minute!'-test, (von Fintel 2004: 322-325). This test is also described in Section 3 of Appendix A.

¹¹ Whether the topic/focus interpretation of Strawson's proposal is correct is somewhat controversial. This is the standard interpretation of Strawson in the literature, but some authors, e.g. Anne Bezuidenhout (p.c.) and Kadmon (2001) have suggested alternative

The present analysis is designed to avoid these consequences while retaining Strawson's basic idea that discourse topics, or rather questions under discussion, play an essential role in generating intuitions.

3 SWEPPs In Context

The accommodation theories afford little attention to the influence of context on truth value intuitions. The tacit assumption seems to be that contextual information plays no significant role since conflicting intuitions are elicited even when no explicit contextual information is supplied — as witnessed in the contrast pairs listed earlier. It is therefore presupposed that it is sufficient to provide an explanation of isolated SWEPPs, or in Strawson's words, SWEPPs evaluated as 'asserted abruptly and out context' (Strawson 1964: 68). Nevertheless, explaining how minor variations in sentence content and sentence form are capable of affecting intuitions in such a remarkable way has proved exceedingly complicated, so in this section I want to explore and discuss the behavior of truth value intuitions in context. The hope is that this will provide important insights into their curious behavior.

An interesting characteristic of squeamish SWEPPs, that is SWEPPs which are liable to elicit a sense of squeamishness when asserted abruptly and out of context, is the ease with which these sentences are transformed into SWEPPs which elicit an immediate sense of falsity. In particular it seems that only minimal contextual framing is required to transform a sense of squeamishness into a robust intuition of falsity. To illustrate this, consider the cases below.

- (21) # The king of France is bald.
- (27) # The king of France owns a pen.
- (34) # The king of France always enjoys a croissant in the morning.

Without any contextual information, these sentences are liable to elicit a feeling of squeamishness, i.e. an inability to evaluate them for truth or falsity. However, with a minimal amount of contextual information, the sentences are quickly transformed into robust SWEPPs. Here is a simple case. Suppose you are compiling a list of every monarch in Europe who

interpretations. Kadmon (2001: 402-403) contends that Strawson has in mind questions under discussion, rather than topic/focus, and on that interpretation my proposal is very much an elaboration of Strawson's suggestions.

owns a pen. Suppose someone utters (27).¹² In this context, when evaluating the sentence in relation to a salient question, the sentence seems to elicit an intuition of falsity. Asserted abruptly and out of context, that sentence elicits squeamishness, but this minor tweak in the context seems to change intuitions. In other words, we do not need much from context to get falsity judgments.

To drive the point home, I present two cases where a suitably enriched context appears to facilitate falsity intuitions. And since these enriched contexts do not engender any relevant changes to our beliefs, the accommodation theories discussed in the previous section are unable to explain the variations in judgments.

Context I The famous hair-dresser, Pierre, is organizing a fashion show, where the hair-models are all European royalty. Incidentally, minimalism is the new black in hair-style fashion, so Pierre is trying to determine whether he can convince some royalty to go bald for the show. As he is deliberating, he turns to you and his other assistants and says “I can’t think of any royal who would agree to go bald just for the sake of the fashion show, and I can’t think of any who is already bald. You guys get on this immediately. Whoever finds a solution gets to go to the fashion show” Another assistant standing next to you exclaims: “The king of France is bald.” (CI)

I think there is a natural inclination here to judge that the assistant’s utterance of (21) is straightforwardly false. Now, recall that (21) is the paradigm example of a sentence that elicits no robust intuitions when asserted abruptly and out context. Here is another example.

¹² Strawson uses the question: ‘What examples, if any, are there of famous contemporary figures who are bald?’. These cases are noted in the literature, but proposals to explain them are fairly infrequent. However, see von Fintel (2004: 337).

Context II	<p>A salesman from the company Froggy Croissants is attempting to persuade you that your high-end bakery should start selling croissants from Froggy Croissants. You're sceptical and say: "The problem is this. Our customers are primarily wealthy snobs with poor taste who only buy products which are trendy or hip. So, unless Paris Hilton or some other celebrity is eating croissants on a regular basis, we will not be able to sell them." The salesman retorts: "I take your point, but you see, our product is extremely popular in France. For instance, the king of France always enjoys a croissant in the morning."</p> <p style="text-align: right;">(CII)</p>
-------------------	--

When (34) is embedded in the above context, it seems to me that it is natural to again judge it straightforwardly false. And again, observe that this sentence elicits no robust intuitions when evaluated without explicit contextual information.

The principal question is why intuitions change in the above contexts. It is easily recognized that not *any* addition of contextual information is sufficient for achieving an intuition-changing effect, and that the context must be primed in a suitable way. For instance, suppose that the salesman in the latter context retorted by uttering *I take your point, but you see, our product is extremely popular in France. For instance, the king of France is bald*. While this would be a very odd thing for the salesman to say, indeed a clearly inappropriate assertion, I think there is no sense in which a robust judgment of falsity automatically results. Yet, it takes little imagination to construct utterance contexts which are suitable for changing a sense of squeamishness into a robust intuition, so the question remains.

In addition to the contextual framing above, there are other curious ways of altering intuitions. For instance, prosodic focus is often noted to have effects similar to those observed with contextualized SWEPPs. Consider the sentences in (47-49) — capitals indicate intonational stress.

- (47) ^F The king of FRANCE is bald.
- (48) ^F The king of FRANCE owns a pen.
- (49) ^F The king of FRANCE always enjoys a croissant in the morning.

Here I think there is a rather striking contrast between the sentences above and their unfocused counterparts. In particular, when the non-denoting definite is focused, indicated above by intonational stress (or a pitch accent) on the proper name 'France', these sentences pull immediate intuitions towards falsity. This is surprising, because it is standardly assumed that absent certain focus-sensitive particles (e.g. 'only' and 'even'), focus has no

truth conditional effects. Here truth value intuitions nevertheless appear to be affected by focus.¹³

Similar perplexing effects are sometimes created by certain syntactic environments, e.g. clefts.

(50) ^F It is the king of France who is bald.

Again, I think there is a strong pull towards falsity here rather than squeamishness. Notice that the existential presupposition triggered by the cleft should not affect our intuitions here since, let's assume, it is common ground that *someone* is bold.¹⁴

The remainder of this chapter is devoted to sketching an alternative explanation of the relation between truth value intuitions and SWEPFs and providing a tentative procedure for predicting when SWEPFs elicit intuitions of falsity (or at least pull our intuitions in the direction of falsity). This procedure is then argued to also provide an explanation of the curious behavior of these intuitions – in particular, why they are so easily changed by context.

4 Discourses and Questions

SWEPFs typically elicit one of two contrastive reactions. One either judges that the SWEPF is false or one feels a sense of squeamishness, viz. an inability to judge whether the SWEPF is true or false. As emphasized by von Stechow (2004), the unifying characteristic of both these reactions is that the informational content of the SWEPF is *rejected*. For instance, if the SWEPF is asserted in the context of a discourse, then regardless of the interlocutors' particular reaction to the SWEPF, the informational content

¹³ For the accommodation theories, which purport to explain intuitions of falsity in terms of inconsistencies, this data is particularly problematic. While focus is assumed to have the semantic effect that a set of contextually constrained alternatives is computed, it is unclear how such a set of alternatives is to be used to derive an inconsistency. There is therefore no obvious way for the accommodation theorists to explain the effects of prosodic focusing on truth value intuitions. I am grateful to Nate Charlow (p.c.) for making this point more clear to me.

¹⁴ One might think that the culprit here is an exclusivity implicature which clefts are often argued to convey. However, if in the context it is clear that the identity of a single individual is sought, and thus a potential exclusivity implicature is satisfied, (50) still elicits an intuition of falsity. I thank an anonymous reviewer at L&P for pointing this out.

of the SWEPPF is rejected from the common ground.¹⁵ The question is how to explain this prominent difference in reactions, and I suggest that an answer can be found by considering certain facts about discourse structure, discourse goals, and questions.

The central contention here is that a very large class of falsity intuitions result from evaluating the content of a SWEPPF in relation to a contextually salient question. The hypothesis is that when certain contextual/conversational conditions are satisfied, SWEPPFs elicit intuitions of falsity. In particular, if a SWEPPF is evaluated as what I label a *consonant* response to a contextually salient question and the relevant question has true answers, the SWEPPF is judged false. The purpose of the present section is to provide a systematic and detailed explication of these conditions. I therefore adopt a theory of information structure in discourses based on Roberts (1998, 2004). This is a convenient framework to examine how intuitions behave in certain types of discourses, namely discourses where information exchange is the primary objective.

4.1 Discourse Structure

Following Roberts (1998, 2004), let's assume that the goal of a discourse is to address and answer certain questions under discussion (QUDs).¹⁶ Roberts writes,

I assume that the primary goal of discourse is communal inquiry — the intention to discover with other interlocutors “the way things are”, to share information about our world. [...] The linguistic counterpart of an inquiry is a question. Thus, we might take questions to be formal objects that reflect interlocutors' intentions in conducting discourse. (Roberts 2004: 208).

Let's make the additional assumption that in order to count as a discourse participant, one must have an antecedent interest, or immediately acquire an interest, in addressing or answering the QUDs — or as Roberts puts

¹⁵ I am deliberately ignoring cases involving presupposition accommodation, since these are irrelevant here. I am only concerned with cases where the interlocutors know that the presupposition is false and resist accommodation.

¹⁶ Roberts' model relies importantly on the notions of a *common ground* and a *context set* as developed in Stalnaker (1970, 1974, 1978, 1998, 2002). On Stalnaker's theory, adding a proposition to the common ground is represented by the set theoretic operation of intersection, namely intersecting the set of worlds denoted by the proposition and the set of worlds already in the context set.

it, to count as a discourse participant is to accept a *proffered* question. Finally, let's assume that when a speaker engages in a discourse, the speaker is presupposing that someone in the discourse (e.g. the speaker herself) is capable of relevantly addressing or answering the QUD. Adopting this model of discourses allows us to examine the relation between reactions to SWEPPs and QUDs. In particular, it helps illustrate that the extent to which the asserted SWEPP addresses a QUD has a significant influence on our reactions to the SWEPP.

On this picture, discourses contain at least one QUD and this question is accepted by the discourse participants. The goal of the discourse is to answer the QUD, and speakers may employ different strategies to achieve that goal. One strategy is to attempt to answer the QUD, the primary question, by addressing and answering a range of subquestions. A subquestion is defined as a question entailed by the primary question, i.e. the question in (51) entails the subquestions s_1, s_2, s_3 etc.

(51) What did Mary eat?

$$(51) \left\{ \begin{array}{ll} s_1 & \text{Did Mary eat clams?} \\ s_2 & \text{Did Mary eat braised short ribs?} \\ s_3 & \text{Did Mary eat veal shanks?} \\ \vdots & \vdots \\ s_n & \text{Did Mary eat a tomato salad?} \end{array} \right.$$

This entailment relation for questions is defined as follows: (51) entails $s_1, s_2, s_3 \dots s_n$ when a complete and exhaustive answer to (51) entails an answer to $s_1, s_2, s_3 \dots s_n$. Conversely, answering a subquestion is to provide a partial answer to the primary question and answering every subquestion is therefore to provide a complete answer to the primary question. Now, on this rough picture of discourse structure, Roberts defines a notion of relevant conversational moves. A conversational move m , Roberts says, is relevant to the QUD if and only if m introduces a partial answer to the primary question or is part of a strategy to answer the primary question. Partial answers to a QUD can be given both directly or indirectly, i.e. one can relevantly address a QUD by asserting a sentence which (contextually) entails or implicates an answer to the QUD. Moreover, a conversational move m is part of a strategy to answer the primary QUD if e.g. it proposes a subquestion to the primary QUD or serves to restrict the domain of inquiry. For example, one can propose a subquestion but restrict the domain so as to exclude certain answers as irrelevant with regards to the QUD.

4.2 Consonance

There are roughly two predominant semantic analyses of questions, namely proposition-based analyses on which questions are analyzed as sets of propositions where these propositions constitute possible answers, see e.g. Hamblin (1973) and Groenendijk and Stokhof (1984), and *functional* (or *categorical*) analyses on which questions are analyzed as functions from individuals to truth values, see e.g. Hausser (1983), Scha (1983), Krifka (2001, 2004).

Now, since Roberts defines relevant assertions in terms of (direct or indirect) possible answers, relevant conversational moves do not include assertions that contain non-denoting descriptions (or non-denoting determiner phrases (DPs) in general) as constituents. Since a non-denoting DP cannot be mapped to an individual in a model, such an assertion cannot be a member of an answer set. Consequently, on both the proposition-based and the functional analyses of questions, assertions of sentences with non-denoting constituents do not constitute possible answers. In this section I therefore intend to characterize an alternative set of conversational moves which are intuitively interpreted as *cooperative* and as *addressing* a QUD even if they fail to provide an answer to the QUD. So, let's say that a conversational move *m* is *consonant* only if (a) the speaker intends to answer a QUD, (b) the speaker could reasonably be interpreted by the interlocutors as intending to answer the QUD, and (c) if the presuppositions of the relevant assertion were accommodated, the assertion would provide a (direct or indirect) answer to the QUD. In other words, for a conversational move to be consonant, it is strictly speaking not required that it provides an answer to the QUD, viz. a proposition/individual contained in the relevant answer set. Rather, what is required is that the conversational move could reasonably be interpreted by the interlocutors as an attempt to answer the QUD and that if its presuppositions were accommodated, it would either entail or implicate a proposition/individual contained in the answer set. On this characterization, the set of conversational moves consisting of assertions of true and false answers to a QUD is only a proper subset of the set of consonant conversational moves. The set of consonant conversational moves also contains speech acts which do not provide genuine answers to the QUD, yet are interpreted as such. The set of consonant conversational moves is a proper subset of the set of appropriate conversational moves. Not every appropriate conversational move is consonant, since not every appropriate conversational move can reasonably be interpreted as intending to answer a QUD. Here is a crude illustration.

(52) Who ate the braised short ribs?

a ₁	Mary.	+relevant / +consonant / +appropriate
a ₂	Mary is a vegetarian.	+relevant / +consonant / +appropriate
a ₃	I do not know.	-relevant / -consonant / +appropriate
a ₄	Mary jumped the fence.	-relevant / -consonant / -appropriate

The response in (52-a₁) is a genuine answer to the QUD and thus a consonant conversational move even if Mary did not eat the braised short ribs. While the response in (52-a₂) is not a direct answer to the QUD, it implicates an answer to a subquestion and is therefore also a consonant conversational move. Conversely, a response such as (52-a₃) is an inconsonant conversational move. It is inconsonant because a speaker asserting (52-a₃) is not attempting to answer the QUD. Indeed, the speaker asserts that she's not in a position to answer the question. Nevertheless, this is a perfectly appropriate response to the question insofar as the speaker fails to know the answer. Finally, the response in (52-a₄) is also an inconsonant conversational move. Here the speaker cannot reasonably be interpreted as attempting to answer the QUD.¹⁷

In the representation in (52) it seems that the set of consonant and relevant conversational moves coincide, so why is the notion of consonance needed? Consonance is designed to capture an intuitive aspect of conversational pertinence that extends beyond the set of conversational moves which provide (directly or indirectly) answers to a QUD. Where a *relevant* conversational move must ultimately provide an answer (directly or indirectly), a consonant conversational move need not do this. The idea is that intuitions of falsity with regards to SWEPFs depend importantly on the relation between a speech act, its content, and a given QUD, and it is this important relation which consonance is designed to capture. To characterize the set of conversational moves that give rise to falsity intuitions, a class of conversational moves extending beyond a set of genuine answers is required. As emphasized already, if *wh*-questions are defined in terms of sets of propositions, or as functions from individuals to truth values, purported responses which contain a non-denoting DP fail to constitute

¹⁷ On this characterization of sets of conversational moves (CM), the following relations obtain: $CM_1 \subset CM_2 \subset CM_3$, where $A = \{p : p \text{ is a true or a false answer to a QUD}\}$, $CM_1 = \{m : m \text{ is an assertion of } p \text{ where } p \in A\}$, $CM_2 = \{m : m \text{ is consonant}\}$, and $CM_3 = \{m : m \text{ is appropriate/felicitous}\}$. In other words, the set of appropriate/felicitous conversational moves is going to contain both true answers, false answers, merely purported answers, and claims of ignorance. In this characterization, answers are assumed to be propositions, but it should be obvious that it can easily be assimilated to a functional analysis of questions where e.g. single constituent questions denote sets of individuals.

a genuine answer. For example, if the relevant set of answers is a set of propositions, it will not contain any elements which have non-denoting DPs as their constituents. Similarly, if the relevant set of answers is a set of individuals, it will not contain non-existing individuals. Nevertheless, there is a clear intuitive sense that certain responses to questions, where the asserted response contains a non-denoting or non-referring expression, *can* be interpreted as a purported answer even though it is not a genuine answer in the technical semantic sense. For example, suppose we are discussing who proved Goldbach's conjecture. If I sincerely assert that the mathematician who proved Goldbach's conjecture is a woman, there is a clear sense in which my assertion is acceptable, cooperative, and addressing the QUD. Of course, no one has proved Goldbach's conjecture, so the real problem is that the discourse is defective. But despite defectiveness, my response is easily interpreted as an intended answer to the QUD. In other words, the deciding factors for obeying consonance is whether the speaker can reasonably be interpreted as answering the QUD and whether the assertion could answer the QUD were its presuppositions accommodated. Consequently, as the notions are used here, (52-a₅) is a consonant conversational move in response to (52). It is not a relevant conversational move, because it fails to provide a genuine answer (directly or indirectly) to the QUD.

- (52) a₅ The mathematician who proved Goldbach's conjecture ate
the braised short ribs. -relevant / +consonant / +appropriate

4.3 The False Intuition Hypothesis (FIH)

SWEPPs asserted in discourses are never incorporated into the common ground as partial or complete answers to a relevant QUD. The question is why SWEPPs nevertheless elicit contrasting judgments. I contend that when the content of a SWEPP is interpreted as relating to a QUD, in particular when it is interpreted as a consonant conversational move, the SWEPP elicits an intuition of falsity rather than squeamishness. The important point to emphasize is that a SWEPP is evaluated in relation to a QUD *only* when it is interpreted as a consonant conversational move. However, falsity judgments also depend on the relevant discourse. In particular, if a discourse is defective, i.e. if the QUD has no true answers, a consonant assertion of a SWEPP is not guaranteed to elicit a falsity judgment. The hypothesis is therefore that when an assertion of a SWEPP satisfies the conditions below, the SWEPP elicits an intuition of falsity.

- (a) The SWEPP is interpreted as a *consonant conversational* move as regards the QUD or a subquestion to the QUD.

- (b) The relevant QUD has true answers.¹⁸

But what would explain this hypothesized generalization? Why do intuitions change when these conversational conditions are satisfied? Here is a tentative explanation: A non-defective discourse contains a QUD which has true answers and the only relevant assertions in such a discourse are assertions which (directly or indirectly) provide a partial or complete answer to the QUD. Every other assertion is irrelevant. However, when we evaluate an assertion in a given discourse, the only important issue is whether the assertion succeeds in providing a *true* answer to the QUD, i.e. only true answers further the goal of the discourse. But in the cases under consideration here, namely assertions of SWEPPs, we are aware that the definite fails to denote, and so it is immediately clear that the assertion could not succeed in answering the QUD. As a result the assertion is not furthering the goal of the discourse. However, if it is consonant, it is then interpreted as cooperative and as an attempt to provide a true answer to the QUD. For this reason, it is grouped among consonant assertions which do not further the goal of the discourse, namely false answers to the QUD. In other words, it seems plausible that consonant assertions of SWEPPs have the same conversational status as consonant assertions of straightforwardly false non-SWEPPs. From the point of view of the discourse participants, these are simply not furthering the goal of the discourse, which is to resolve the QUD. Consequently, the SWEPPs are grouped with other consonant non-SWEPPs (viz. other false answers) and since conversational moves that further the goal of the discourse correlate with true answers, the hypothesis is that consonant conversational moves which do not further the goal of the discourse are naturally interpreted as correlating with false answers.

Before considering the predictions of the present conjecture, it is important to emphasize its scope. A SWEPP *may* give rise to a falsity judgment even when the conditions of the FIH are seemingly unsatisfied. In other words, squeamishness is harder to predict. However, it is plausible that a sense of squeamishness is elicited only when at least one of the following conditions are satisfied: (a) the SWEPP is interpreted as an *inconsonant* conversational move, (b) the speaker is unable to determine which of several salient questions the assertion is purporting to answer, or (c) the discourse is defective.

¹⁸ Probably, what is required here is actually a strictly speaking weaker condition, namely that the QUD is not *known* to have *no* true answers. However, I rely on this slightly simpler condition for the remainder of the chapter.

4.4 Predictions

The FIH provides a straightforward explanation of the changes in truth value intuitions observed with the contextualized SWEPPs from section 2, namely (CI) and (CII). Why do we judge that the assistant's assertion of (21) is false in the context described in (CI)? The simple explanation is this: (CI) is a discourse context where the QUD is arguably 'who would agree to go bald for the fashion show'. On the most easily accessible interpretation, the assistant's response is understood as a purported answer to the QUD. In other words, a consonant conversational move, but a false answer to the question. On this explication of (CI), the SWEPP asserted by the assistant satisfies both conditions of the FIH and is therefore predicted to be false. This is the correct prediction, since (21) seems to elicit an immediate intuition of falsity when asserted in the context of (CI).

An analogous explanation applies to (CII) where the initial QUD is arguably 'should you buy croissants from Froggy Croissants?' it is subsequently determined in the discourse that the answer to that question is conditional on another question, namely whether a famous celebrity also buys croissants from Froggy Croissants. This then generates a new QUD, roughly 'which if any celebrity buys Froggy Croissants?' In response to this QUD, the salesman's assertion of (34) is again interpreted as a consonant conversational move, but a not true answer to the QUD. Both conditions of the FIH are thereby satisfied and the SWEPP is predicted to elicit falsity. Again, this is the desired prediction.

Let's consider a couple of additional examples. The SWEPPs in (53a) and (54a) seem to elicit a sense of squeamishness when evaluated as asserted abruptly and out of context. However, when these SWEPPs are interpreted as consonant responses to a QUD and the discourse is non-defective, it seems to me that these sentences then elicit intuitions of falsity.

- (53) Who is on a state visit?
 - a. ^F The king of France is on a state visit.
- (54) Is anyone reading Russian literature these days?
 - a. ^F The king of France read Anna Karenina.

(53a) is a consonant conversational move because it is easily interpreted as a purported answer to (53) and were its presuppositions satisfied, it would constitute a genuine answer to the question. It is judged false, but this is predicted by the FIH, because both conditions are satisfied. The response in (54a) is also very easily interpreted as implicating an answer to the question in (54) and were its presuppositions satisfied it would. This is sufficient for it

to count as a consonant conversational move. Again, the SWEPF is judged false and this is predicted by the FIH — both conditions are satisfied.

While these cases have a relatively artificial feel, it is simple to construct these examples and they appear to confirm the FIH. I conclude this section by emphasizing that what I am suggesting here is a general criterion for determining when and why SWEPFs are judged false. They are judged false when the relevant conditions are satisfied and they are judged false *because* when the relevant conditions are satisfied, the defect exhibited by the SWEPF is negligible. When a response is consonant, the speaker is not immediately interpreted as uncooperative (perhaps just confused) and this is the reason that the defect is negligible. In subsequent sections, I demonstrate that a number of *prima facie* perplexing cases can be explained along the lines suggested here, viz. if a SWEPF satisfies the conditions of the FIH, it is judged false.

4.5 *Consonance Constraints*

One interesting class of sentences that could seem to constitute *prima facie* problems for the FIH, are SWEPFs where the non-denoting description contains a (complex) demonstrative. Uses of such descriptions, even in the context of a QUD, can elicit squeamishness rather than judgments of falsity. For instance, suppose we are engaged in a discourse, where the QUD is (55a). If a speaker asserts (55b) as an answer to this question, but demonstrates, say, an empty chair, it seems to me that this sentence fails to elicit a robust intuition of falsity.

- (55) a. Who's is bald?
 b. # The woman sitting in that chair is bald.

However, we should not accept that (55b) is a consonant conversational move. Consonance requires that it is reasonable for the interlocutors to interpret the speaker as intending to provide an answer to the QUD. However, in the discursive exchange above, this is clearly not reasonable. Since the speaker is demonstrating an empty chair, it is not feasible to ascribe to her the belief that there is woman sitting in the relevant empty chair. Moreover, there is no way of interpreting the speaker as simply *not knowing* that the chair is empty. In other words, it is not possible to interpret the speaker as if she simply assumes that the demonstrative refers while maintaining the assumption that she's cooperative. As a result, there is no way of evaluating the speaker's claim. The speaker's assertion of (55b) is therefore an inconsonant conversational move and the conditions of the FIH are not

satisfied. In other words, consonance places demands on communicative cooperation. Now consider (56).

- (56) a. Who's is bald?
b. ?? The naked and invisible donkey standing in the center of the room with an iPhone in its mouth is bald.

My immediate judgment of (56b) is that it is perplexing and odd rather than false. But again, I think this judgment should be explained in terms of inconsonance. The answer in (56b) is just so bizarre that it obscures intuitions. Why? Because it is very hard to attribute to the speaker the belief that there is an individual satisfying the definite's restrictor 'naked and invisible donkey standing in the center of the room with an iPhone in its mouth'. And hence, it is quite difficult to evaluate the speaker's claim. The response thus fails to count as a consonant conversational move, because the speaker asserting (56b) cannot reasonably be interpreted as cooperative and attempting to provide an answer to the QUD.

The last provisional problem for the FIH discussed here concerns cases where an answer is implicated by reporting. Consider the following question-answer pair.

- (57) Who ate the braised short ribs?
a. Mary thinks that Barney ate the braised short ribs.

Insofar as the response in (57a) is consonant, it is consonant because it succeeds in implicating an answer to the QUD, namely that Barney ate the braised short ribs. Compare this to (40b) which strikes me mostly as odd rather than straightforwardly false.

- (57) b. ?? The king of France thinks that Barney ate the braised short ribs.

If the response in (57a) succeeds in implicating an answer to the QUD, one might be inclined to think that (40b) should also succeed in implicating an answer, namely that the king of France ate the braised short ribs. And since (40b) fails to elicit an immediate intuition of falsity, this seems to constitute a counter-example to the FIH.

The crucial question here is when and why a response such as (57a) is consonant, viz. succeeds in implicating an answer to the QUD. It seems plausible that a speaker asserting (40a) is tacitly assuming that Mary's thoughts on the issue are relevant to resolving the QUD and that the interlocutors must accept (or come to accept) this assumption. In other words, one might think that the response in (40a) is consonant only if it is common ground,

or *could* become common ground, that Mary's opinion on the issue can help settle the question. In contrast, it is not common ground (and it could not become common ground) that the king of France's thoughts on the issue could help settle the question. And for this reason (40b) sounds odd. I also think, in support of this explanation, that for most discourses revolving around a QUD such as (57), a response such as (40c) elicits roughly the same intuitive reactions as (40b).

(57) c. ?? Noam Chomsky thinks that Barney ate the braised short ribs.

In other words, if a speaker proffers the opinion of some individual in response to a QUD, this individual's opinion is relevant only on the supposition that she/he is likely to know the answer to the QUD. Consequently, one might think that each response above is interpreted as an answer to a subquestion, namely what the opinion of *x* is where it can be established in the context that *x* is an individual whose opinion is relevant to resolving the QUD.¹⁹ One might also think that the common ground determines which subquestions of this kind are relevant in a discourse, but that a speaker can introduce additional subquestions into the discourse as long as the discourse participants do not object. So if a speaker asserts (40b) in response to (57), the speaker is interpreted as addressing an irrelevant subquestion. i.e. since there is no king of France, it cannot be established in the context that the opinion of the king of France is relevant to resolving the QUD.

This explanation salvages the FIH because with respect to the primary QUD, the response in (40b) is not interpreted as providing either a partial or complete answer to the primary QUD, but rather as providing an answer to a question which is irrelevant and moreover has no true answers.

In the next section, I discuss a particular advantage of the present proposal, namely that it provides a simple and intuitive explanation of the focus cases briefly introduced earlier. No proposal for explaining the curious behavior of truth value intuitions has, to my knowledge, provided a substantial and systematic explanation of these cases.

¹⁹ In support of an explanation in terms of subquestions, it seems that a natural prosodic contour for the responses above involve so-called B-accented subject phrases which is the standard indicator for *contrastive topics*, cf. Roberts (2009). If these subject phrases function as contrastive topics, it is natural to interpret them as intended answers to additional subquestions. This relation between contrastive topics and subquestions is extensively discussed in Kadmon (2001, 2009). I discuss the effects of prosodic focusing more generally in the subsequent section. I thank an anonymous reviewer at L&P for suggesting the explanation involving contrastive topics.

5 Non-Denoting Definite Descriptions and Focus

Focus is standardly assumed to have no truth conditional effects save for interactions with focus-sensitive operators. Nevertheless, when the definite descriptions in (47-49) are focused on account of intonational stress, this *does* seem to impact immediate judgments.

- (47) ^F The king of FRANCE is bald.
- (48) ^F The king of FRANCE owns a pen.
- (49) ^F The king of FRANCE always enjoys a croissant in the morning.

While these sentences normally elicit no robust intuitions, focusing appears to pull immediate intuitions towards falsity. These effects are quite curious and existing analyses of truth value intuitions are at a loss to explain them. The crucial question is how focusing affects intuitions and moreover why. However, these effects can be readily explained by the FIH, but demonstrating this requires a minor digression. In the subsequent paragraphs I therefore provide a brief explication of the notion of *prosodic focus* and present an analysis of its pragmatic and semantic properties.

5.1 Backgrounds and Focus

In a variety of languages (incl. English) pitch accents are used to realize *prosodic focus*. Intuitively, speakers use prosodic focusing in order to contrast the focused material with something already salient in the discourse. For instance, the use of prosodic focus is particularly prevalent when speakers attempt to emend incorrect information in a discourse. This is typically labeled ‘contrastive focus’ and illustrated in (58).

- (58) a. Mary ate clams.
- b. No, JOHN ate clams.

On standard analyses, it is assumed that prosodic focus divides the content of an expression into two parts, the background and the focus. The background is construed as information which is already given in the context, whereas the focused material is construed as new information. Alternatively, one might roughly say that focus-marked material is information intended to be added to the common ground, whereas the non-focused material (the backgrounded material) is assumed to be antecedently established. This division of the content can be neatly represented using function/argument pairs, where the backgrounded information is represented by a function and the new information is represented by an argument (cf. the (c)-cases below).

- (59) a. MARY ate clams. \Rightarrow

- b. $\langle \text{Background, Focus} \rangle \Rightarrow$
 - c. $\langle \lambda x.x \text{ ate clams, Mary} \rangle$
- (60)
- a. Mary ATE clams. \Rightarrow
 - b. $\langle \text{Background, Focus} \rangle \Rightarrow$
 - c. $\langle \lambda R.Mary R \text{ clams, eat} \rangle$
- (61)
- a. Mary ate CLAMS. \Rightarrow
 - b. $\langle \text{Background, Focus} \rangle \Rightarrow$
 - c. $\langle \lambda x.Mary \text{ ate } x, \text{ clams} \rangle$

Although focus is realized by pitch accenting single words or syllables, it is generally accepted that focus often extends to larger syntactic constituents (i.e. DPs, VPs etc.). Even though the pitch accents in (47-49) are placed on the proper name ‘France’, the focus-marking can extend to the entire DP. The problem of understanding the relation between prosody and focus-marking — sometimes referred to as *focus projection* — has spawned a lively discussion which I will refrain from engaging in here. However, it should be noted that focal scope places important restrictions on use. So, in order to distinguish syntactic–semantic focus-marking from mere indicators of prosodic contours, I indicate focal scope using square brackets with a subscripted *f*, and use capitals to indicate stress/pitch accenting. So, (47-49) should be represented as below.

- (47) $^F [\text{The king of FRANCE}]_f \text{ is bald.}$
- (48) $^F [\text{The king of FRANCE}]_f \text{ owns a pen.}$
- (49) $^F [\text{The king of FRANCE}]_f \text{ always enjoys a croissant in the morning.}$

5.2 *Structured Meanings: The Semantics and Pragmatics of Focus*

It is generally assumed that there is an important relation between the semantic and pragmatic effects of focus and the semantics of questions. The currently predominant semantic analyses of focus, namely *Alternative Semantics* and *Structured Meanings*, are therefore extensions of the currently predominant analyses of questions. In particular, *Alternative Semantics* is a proposition-based analysis extending the proposition-based semantic analysis of questions, whereas *Structured Meanings* is a functional analysis, extending the functional semantic analysis of questions. In the exposition here, I use a *Structured Meanings* analysis, but I want to emphasize that an analysis using *Alternative Semantics* makes identical predictions.²⁰

²⁰ *Structured Meanings* is discussed in von Stechow (1989, 1991) and Krifka (1991, 2001). *Alternative Semantics* was introduced by Mats Rooth in (1985) and further discussed in

On the functional analysis of questions, as previously mentioned, *wh*-questions denote functions from individuals to truth values which are obtained by lambda-abstracting over the *wh*-phrases (or rather their denotations). This analysis therefore assigns a very intuitive meaning to questions, namely a function which takes answers as its arguments. A question such as (62a) is therefore analyzed as (62b).

- (62) a. What did Mary eat?
 b. $\lambda x. \text{Mary ate } x$ ← λ -abstracting on the *wh*-phrase

The Structured Meanings account of focus extends the basic idea in the Structured Meanings semantics for questions. Focus-induced interpretations of sentences are obtained by lambda-abstracting on the focus-marked constituents. This produces the focus-induced *structured meaning* of the sentence, namely an ordered sequence whose members are a property (a function) and an individual (an argument). This analysis of focus is represented in (59) — repeated below.

- (59) a. $[\text{MARY}]_f$ ate clams.
 b. $\langle \lambda x. x \text{ ate clams}, \text{Mary} \rangle$ ← λ -abstracting on the focus

Given this analysis, recovering semantic content, viz. truth conditional content, is straightforward, but it also provides a useful resource for explaining various effects. For instance, semantic effects observed when focused elements interact with focus-sensitive particles and pragmatic effects such as question-answer congruence.

A common observation in the study of focus is that focus restricts the set of contexts in which sentences can be felicitously asserted. The paradigm example of this effect is question-answer pairs. The general observation is that focus-marked sentences asserted in response to *wh*-questions are felicitous only if the focus-marked constituent is correlated with the *wh*-phrase of the particular question.²¹ For example, consider the question (63) and each response in (63a-63c). Observe that responses where the focus-marked constituent is not correlated with the *wh*-phrase are infelicitous.

- (63) What did Mary eat?

Rooth (1992, 1996). These theories diverge in important respects, but the differences only emerge at a fairly detailed level of analysis. As regards the cases under consideration here, the predictions of these two theories are identical. Illuminating discussions are available in Kadmon (2001) and Beaver and Clark (2008).

²¹ Cf. Kadmon (2001), Krifka (2001), Rooth (1996) for discussion.

- a. Mary ate [CLAMS]_f.
- b. * Mary [ATE]_f clams.
- c. * [MARY]_f ate clams

Hence, we can think of the focus induced constraints as constraints on what questions a sentence can felicitously answer.

The Structured Meanings analysis provides a straightforward explanation of these constraints. Since focus-induced interpretations are derived by lambda-abstraction on the focus-marked constituents, this procedure provides a method for determining which questions the focus-marked sentence is most naturally understood as purporting to answer. It can only felicitously answer questions with a denotation that corresponds to the function obtained from abstraction on the focus, cf. (62b).

5.3 *Focus and Non-Denoting Definites*

The close correlation between questions and focus provides a simple and intuitive explanation of the changes in truth value intuitions noted in the beginning — cf. (47-49). The explanation is that focus-marking fixes the questions that the sentence can felicitously answer. With no focus-marked constituents, (47-49) could serve as felicitous answers to different questions, but the result of focusing the DPs in (47-49) is that these sentences are now most naturally interpreted as providing an answer to a specific question.

- (64) a. [The king of FRANCE]_f is bald.
 b. $\langle \lambda x.x \text{ is bald, the king of France} \rangle$ ← λ -abstracting on the focus
 c. $\lambda x.x \text{ is bald}$ ← derived function
 d. Who is bald? ← the relevant question
- (65) a. [The king of FRANCE]_f owns a pen.
 b. $\langle \lambda x.x \text{ owns a pen, the king of France} \rangle$ ← λ -abstracting on the focus
 c. $\lambda x.x \text{ owns a pen}$ ← derived function
 d. Who owns a pen? ← the relevant question
- (66) a. [The king of FRANCE]_f always enjoys a croissant in the morning.
 b. $\langle \lambda x.x \text{ always enjoys a croissant in the morning, the king of France} \rangle$ ← λ -abstracting on the focus
 c. $\lambda x.x \text{ always enjoys a croissant in the morning}$ ← derived function
 d. Who always enjoys a croissant in the morning? ← the relevant question

λ -abstraction on the focus yields an ordered sequence, a function and an argument. The derived function is the denotation of a question, namely the denotation of a *wh*-question where we have λ -abstracted on the *wh*-phrase. Hence, we have obtained a simple method for determining which question these sentences are most naturally interpreted as answering.

(47-49) elicit intuitions of falsity and given the analysis above this is also predicted by the FIH. The (a) condition is satisfied, because (47-49) are most easily interpreted as responses to particular questions and therefore as consonant conversational moves. The (b) condition is satisfied, because the questions against which each of these sentences are evaluated have true answers (cf. the questions above). The prediction of the hypothesis is therefore that these sentences should elicit immediate intuitions of falsity and I believe the data verifies this prediction.

The hypothesis that there is a systematic connection between robust intuitions of falsity and consonant conversational moves, i.e. putative answers to QUDs, is thus supported by the cases involving focus. But in order to evaluate in more detail the predictions of the present hypothesis, it is instructive to consider what happens if we change focus-structure in (47-49). Interestingly, it seems that the result of focusing constituents outside the non-denoting DP is again a change in intuitions. In particular, I find that there is a clear contrast between (47-49) and (67-69).

- (67) # The king of France is [BALD]_f
- (68) # The king of France owns [a PEN]_f
- (69) # The king of France always enjoys [a CROISSANT]_f in the morning.

These cases seem to me to elicit a sense of squeamishness rather than falsity. Out of context assertions of (67)-(69) are distinctly odd rather than obviously false. But why is there this marked difference in intuitions when the focus-marking is changed from the non-denoting subject DP to a constituent in the VP? Let's again analyze these constructions.

- (70) a. The king of France is [BALD]_f.
 b. $\langle \lambda R. \text{The king of France is } R, \text{ bald} \rangle$ \leftarrow λ -abstracting on the focus
 c. $\lambda R. \text{The king of France is } R$ \leftarrow derived function
 d. The king of France is what? \leftarrow the relevant question
 (what is the king of France like?)
- (71) a. The king of France owns [a PEN]_f.
 b. $\langle \lambda P. \text{The king of France owns } P, \text{ a pen} \rangle$ \leftarrow λ -abstracting on the focus
 c. $\lambda P. \text{The king of France owns } P$ \leftarrow derived function

- d. The king of France owns what? ← the relevant question
(what does the king of France own?)
- (72) a. The king of France always
enjoys [a CROISSANT]_f in the morning.
- b. $\langle \lambda P. \text{The king of France always} \leftarrow \lambda\text{-abstracting on the focus}$
enjoys P in the morning, a croissant}
- c. $\lambda P. \text{The king of France always enjoys P} \leftarrow \text{derived function}$
in the morning.
- d. The king of France always enjoys ← the relevant question
what in the morning?
(what does the king of France always enjoy in the morning?)

Given this analysis, it is unsurprising that these sentences elicit squeamishness when constituents outside the non-denoting DP are focus-marked, because this produces a different question-answer structure. The sentences in (67-69) are most naturally interpreted as answering different questions than the sentences in (47-49), namely questions which lack true answers; absent heavy duty contextual assumptions, there are no true answers to a question about the properties of a non-existing entity and a discourse revolving around a QUD that has no true answers is a defective discourse. As a result, these sentences are most naturally interpreted as consonant answers in a defective discourse. This result is consistent with the FIH. The sentences (70-72) satisfy the (a) condition, because despite the fact that the relevant questions have no true answers these sentences are naturally interpreted as consonant conversational moves. However, since the relevant questions have no true answers, (70-72) fail to satisfy the (b) condition. Again, this supports the hypothesis that falsity intuitions are a result of evaluating sentences with respect to QUDs – namely QUDs which have true answers (or are not known to have no true answers).²²

²² One potential problem for the account sketched here is multiple focus constructions, where one of the focused constituents is a non-denoting DP. While I am not confident that the data from such cases is entirely stable, it seems that these elicit falsity judgments less frequently. A similar problem arises for discourse contexts with multiple constituent questions, viz. questions which contain two *wh*-phrases such as ‘who ate what?’. Multiple constituent questions and multiple focus constructions are already a point of controversy in the literature on the semantics and pragmatics of focus, cf. e.g. Krifka (1991, 2001), so how to resolve these potential problems is a delicate issue. An in-depth discussion would require at least a serious study of the relevant data and plausibly a more detailed and more complex story about prosody, focus-distribution, and question-answer congruence. The aim here is to provide a fairly programmatic account of the relation between truth value intuitions and context, so I eschew a discussion of these potential problems here.

6 Extending the Proposal

6.1 Complex Constructions and Intensional Interpretations

The problem of conflicting truth value intuitions, as ordinarily conceived, concerns a puzzling and apparently non-systematic distribution of truth value judgments. In short, the problem is explaining why some SWEPFs elicit judgments of falsity while others elicit a sense of squeamishness. However, if complex/intensional cases are taken into consideration, this distribution appears to break down. In particular, it seems that complex/intensional SWEPFs often give rise to robust judgments of truth. Consider for example (73).

- (73) Who does Mary fear?
a. Mary fears the king of France.

If Mary believes that there is a king of France and also fears the king of France, it seems that (73a) is true. It also seems that an adequate semantic analysis should be able to predict this. However, the truth of (73a) clearly depends on an intensional interpretation of the sentence. (73a) cannot be judged true unless it is assumed that Mary believes that there is a French king. In contrast, suppose it is common ground that Mary knows that there is no king of France. In this case, i.e. where (73a) is interpreted extensionally, I think there is a natural inclination to judge (73a) false. And this is what the FIH predicts.

In other words, explaining truth value intuitions as regards complex/intensional cases is a quite different project than explaining truth value intuitions as regards non-complex/extensional cases. First, the distribution of intuitions is, for various reasons, importantly different in the complex cases. Secondly, it is widely agreed that a semantic analysis of complex constructions requires a more general discussion of the semantics of epistemic and root modals, attitude verbs, if-clauses, quantificational adverbs, and the full range of logical connectives. This is obviously a rather daunting task. I therefore think that it is quite reasonable to assume that complex cases, and in particular intensional cases, should be set aside. There are good reasons to believe that intuitions about such cases must be explained along quite different lines and involve a range of considerations which are not relevant in non-complex/extensional cases.²³ This, I take it, is also the

²³ For example, a presuppositional analysis should not predict truth value gaps for every complex sentence containing a non-denoting definite, i.e. when embedded in modal constructions, conjunctions, conditionals etc. it is standardly assumed that an adequate

reason that complex/intensional cases are rarely discussed in the literature on the present problem. So, while (73a) is a putative counter-example to the FIH (i.e. both conditions of the FIH are satisfied, yet one reading of the sentence fails to elicit a sense of falsity), it is a counter-example only on an intensional reading. However the FIH should be understood as a hypothesis which aims to explain a particular distribution of intuitions about non-complex/extensional cases — not complex/intensional cases.

Now, having set complex/intensional cases aside, one might nevertheless worry whether my proposal is extendable to SWEPPs evaluated as asserted abruptly and out of context. The *prima facie* problem is that an explanation of such non-contextualized SWEPPs cannot obviously rely on facts about the discourse context or QUDs. In the next sections, I therefore provide a rough sketch of an explanation of how non-contextualized cases can be assimilated to the present proposal.²⁴

6.2 Out-of-Context SWEPPs

The notion of evaluating a sentence as *asserted abruptly and out of context* is frequently used in the literature, but everyone of course agrees that assertions do not occur *out of context*. Any assertion, real or imagined, occurs in context and against the background of several contextual assumptions. This much is uncontroversial. But for that very reason, it is hard to believe that we are capable, in any substantial sense, of evaluating SWEPPs as asserted abruptly and out of context. The question then is what process is invoked and against what contextual assumptions we are evaluating the sentences.

A plausible answer is that we evaluate the sentence by trying to determine which question an assertion of the sentence is intended to address. If so, an explanation of the distribution of intuitions as regards non-contextualized

account of presupposition projection is what is required to avoid this consequence and the current literature contains numerous attempts to provide such an account. In contrast, while a standard Russellian analysis of definites is capable of computing classical truth values for definites embedded in complex constructions, this requires a range of not unproblematic assumptions about scopal ambiguities, see e.g. Fodor (1970), Bäuerle (1983). Related problems with an existential analysis and modal verbs is discussed in detail in Chapter IV but see also Heim (1991), Elbourne (2005, 2010).

²⁴ I should emphasize that the aforementioned accommodation theories are not obviously inconsistent with my proposal. While I have argued that these theories suffer from certain inadequacies, one could, I believe, combine my analysis with e.g. von Stechow's (2004) analysis of non-contextualized cases. From the point of view of theoretical parsimony, a uniform account is obviously desirable, but it is perfectly conceivable that different explanations are required in order to account for the full range of data.

cases is that if it is determined that the assertion is intended to address a question which is known to have no true answer, it elicits a sense of squeamishness. Similarly, if it cannot be determined which question the assertion is intended to address or the question cannot be answered, this also gives rise to a sense of squeamishness. In contrast, one might think that SWEPFs elicit an intuition of falsity when it is determined and interpreted as an intended answer to a question which is not known to have no true answers. The conjecture is thus that these robust intuitions also result from evaluating the SWEPF against a particular question — a question which the evaluator for various reasons interprets the sentence as addressing. What requires explaining is how relevant questions are identified and how we select among them.

6.3 Discourse Topics and Practical Relevance

Successful discourse comprehension depends crucially on determining what information an assertion is intended to convey — e.g. what question an assertion is intended to address. One particularly important interpretative task is therefore determining the topic of discussion. In simple subject-predicate constructions, one might think that absent contextual information there is a natural inclination to interpret the grammatical subject of a sentence as the *topic* where topic is understood as the topic of the discourse or a question under discussion. While topics in the sense of particular syntactic constituents, i.e. grammatical subjects, and topics in the sense of QUDs should not be conflated, it is not implausible that these *are* conflated in discourse interpretation when no information about the discourse context is supplied.²⁵ This would help explain a specific pattern of judgments that was observed by Strawson, namely that when a non-denoting DP occurs in a sentence as the grammatical subject, the sentence is liable to elicit a sense of squeamishness. But when the same non-denoting DP occurs as grammatical object (or as a prepositional object), the sentence is liable to elicit an intuition of falsity. Indeed, it is actually somewhat difficult to find examples of sentences with non-denoting DPs in object position that elicit a sense of squeamishness. This contrast is captured in Strawson's famous examples (74-75) and these furnished the grounds for Strawson's proposed analysis.

- (74) # The king of France visited the Exhibition. (Strawson 1964: 67-68)
 (75) ^F The Exhibition was visited by the king of France.

²⁵ See Roberts (2009) for a discussion of the difference between these notions of topics.

If Strawson's observation is correct, it provides an explanation of a wide range of non-contextualized SWEPPs that is consistent with the FIH: If the grammatical subjects are interpreted as topical, i.e. as the topics of discussion, these DPs are naturally interpreted to be part of the question that is being addressed. In other words, the non-denoting DPs are treated as *given* or as *backgrounded*. This means that a sentence such as (74) is naturally interpreted as addressing a question about the king of France whereas (75) is naturally interpreted as addressing a question about the exhibition. This provides a straightforward explanation of the distribution of intuitions which is consistent, indeed predicted, by the FIH.

But this explanation is, unfortunately, insufficiently general. We have already encountered a variety of SWEPPs with non-denoting DP subjects which nevertheless elicit intuitions of falsity. This is one of the reasons that Strawson's proposed analysis fails. However, Strawson's comments on the observed contrast are quite illuminating.

We might, for example, have felt a shade more squeamish if we had written 'The king of France visited the Exhibition yesterday' instead of 'The Exhibition was visited yesterday by the king of France'. We feel very squeamish indeed about 'The king of France is bald' presented abruptly, out of context, just because we do not naturally and immediately think of a context in which interest is centred, say, on the question *What bald notables are there?* rather than on the question *What is the king of France like?* or *Is the king of France bald?* Of course, to either of these two questions the statement would not be just an incorrect answer. *These* questions have no correct answer. (Strawson 1964: 68-69)

Here Strawson raises an interesting question, namely how interests are centered in interpretation. While there might be a natural inclination to interpret grammatical subjects as topical, i.e. as the topics under discussion, it is possible that this natural inclination is sometimes overridden. For example, one might think that that is one of the roles of prosodic focusing. But another factor which could plausibly impact how interests are centered on specific questions is *practical relevance*. It seems quite likely that evaluators are liable to interpret sentences as addressing a specific question if that question is somehow pertinent to their interests. For example, consider the SWEPPs below which seem to elicit intuitions of falsity even when evaluated without any information about the discourse context.

- (76) ^F The king of France hates your mother.
 (77) ^F The king of France is planning to bomb your city.

These sentences can serve as felicitous answers to a wide range of questions, but absent contextual information, these would be most naturally interpreted as responses to one of the two questions in (78) and (79) respectively.

- (78) a. Who hates your mother?
b. Who does the king of France hate?
- (79) a. Who is planning to bomb your city?
b. What is the king of France planning to bomb?

The hypothesis here is that speakers engaged in discourse interpretation select among immediately available questions with a default to interpreting grammatical subjects as topical. If nothing prompts selecting a particular question, the grammatical subject is assumed to be the topic of discussion. However, if one of the available questions is practically relevant to the evaluator and is assumed to have true answers, the evaluator's interest is centered on that question and the default interpretation of the grammatical subject as topical can thereby be overridden. If so, the sentence is interpreted as addressing that question, rather than the other available question. For example, if (76) triggers the questions in (78), the question (78a) has a clear practical relevance to most people, viz. most people have an interest in the answer to that question. In contrast, the question in (78b) has no practical relevance to most people. If this is right, evaluators should normally select the (a) questions rather than the (b) questions which are known to have no true answers. And if that is the case, then (76) and (77) should both elicit intuitions of falsity—which they do.

This line of reasoning should also explain why an out-of-context assertion of a SWEPP such as (21) elicits squeamishness. Both questions triggered by a sentence such as (21) are practically irrelevant, and as a result evaluators default to an interpretation on which the non-denoting DP is topical. This yields a prediction consistent with the FIH. However, for this idea to work, practical relevance must be construed quite broadly. For example, if a putative answer to a question would be controversial or out of the ordinary, this should suffice for it to be practically relevant.

A good analogy here is presupposition accommodation. It is fairly well established that speakers will refuse to accommodate a presupposition if the presupposed content is controversial or out of the ordinary (see e.g. Section 4.1 of Appendix A). I believe that something quite similar is going on in the cases under consideration here. If an assertion is interpreted as providing an answer to a question where the answer would be controversial or out of the ordinary, speakers are liable to center their interest on that particular question. This would for example explain the contrast between (21) and (23–22) and also why there is a contrast in intuitions between the two sentences below.

- (80) a. # The king of France lives in castle.
b. ^F The king of France lives in a spaceship.

I acknowledge that this explanation amounts, at best, to a very rough sketch of a general strategy for explaining non-contextualized cases. To carry any real weight, the notion of practical relevance would have to be explicated in much more detail and even then a number of additional psychological and pragmatic factors must be taken into consideration. However, it would be no surprise if explaining and determining how interest is centered in discourse interpretation turns out to be exceedingly complex. It is beyond any reasonable doubt that informational contents play a very important role in the generation of truth value intuitions, so a complete story about interest-centering must in all likelihood include an account of the relation between informational contents and speaker's interests, desires, intentions and so on. In other words, the really complicated step in analyzing these recalcitrant non-contextualized cases is to explain why, and predict when, evaluators select particular questions. The very rough explanation above is merely a crude example intended to convince you that psychological factors could play a potentially crucial role. If nothing else, one might think that the extremely complex behavior of non-contextualized SWEPFs is a good reason not to rely on such cases in our semantic theorizing.

In conclusion, there is a methodological lesson to be learned here: when doing semantics, we should focus on contextualized cases, since these are more stable and more easily explained. The reliability of non-contextualized cases relies on having a satisfactory theory of interest-centering available.

7 Existential vs. Presuppositional Analyses

I stated in the introduction that the present proposal provides *prima facie* support for a presuppositional analysis of definite descriptions. The reason is that if my proposed analysis is correct, i.e. if intuitions of falsity are the result of a particular relation of conversational pertinence obtaining between speech acts, their contents, and QUDs, this seems to support the conclusion that falsity intuitions are a result of factors which look *prima facie* irrelevant to semantics. In particular, the fact that a sentence elicits an intuition of falsity when its expressed content relates in a suitable way to a QUD seems to reveal nothing significant about the semantic status of that sentence, e.g. its truth value. This puts some pressure on an analysis which maintains that these sentences should be assigned classical truth values. In particular, proponents of such an analysis must explain why these intuitions really are semantically relevant despite the fact that they seem to track various pragmatic factors.

In response, proponents of a Russellian analysis might be tempted to adopt the present proposal and argue that squeamishness occurs *because*

the contents of the relevant sentences fail to relate to a QUD—the failure of conversational pertinence simply *is* the reason that some SWEFPs elicit squeamishness. In other words, the Russellian might attempt to adopt the present proposal only with the order of explanation reversed.

But it should then be noted that such a response is not obviously feasible. My analysis supplies only a sufficient condition for determining when SWEFPs elicit falsity judgments, not a necessary condition and if the explanation is to work in the opposite direction, i.e. if we are to explain the squeamishness intuitions in general, a necessary condition is needed. Of course, this is essentially a shortcoming of my proposal since determining both necessary and sufficient conditions is desirable for independent reasons. Nevertheless, it is important to emphasize that such a necessary condition must be established if the Russellian is to employ the present proposal as an explanation of the variance in truth value intuitions.

Yet even if a necessary condition is determined, there remains a couple of issues for the Russellian to address. For example, if falsity judgments track the semantic status of these sentences, why is there such an intimate relation between conversational pertinence and truth value intuitions? On the present proposal, robust intuitions of falsity are inherently related to the thoroughly pragmatic notion of consonance, but we should not expect such pragmatic factors to influence intuitions which *do* track the semantic status of the sentences. If these are genuine *semantic* intuitions, why are they so closely correlated with the satisfaction of various *pragmatic* conditions?

Secondly, it seems that if conversational pertinence plays as crucial a role, as I have suggested here, we should expect squeamishness to occur much more often. But if we are engaged in a discussion about, say, who won the Super Bowl, and I assert (81), this sentence does not elicit a sense of squeamishness as regards its truth value.

(81) Barack Obama is a Russian astronaut.

Even though this assertion is inconsonant, indeed inappropriate, I still think that we do not hesitate to judge it false. However, if the Russellian is right that the simple lack of conversational pertinence is what causes squeamishness intuitions, squeamishness should occur for a whole range of sentences when these are inconsonantly asserted in discourse. It seems somewhat questionable whether there is such a general effect.

A presuppositional analysis of definites typically analyzes definite determiners as partial functions, which are defined only for a subset of the domain of $\langle e, t \rangle$ -functions. If squeamishness tracks semantic defectiveness, it seems quite appropriate to capture this in a semantic analysis in terms of an undefined function. I therefore take it that the presuppositional

analysis of definite descriptions has a *prima facie* empirical advantage over an existential analysis.

8 Concluding Remarks

The problem of conflicting truth value intuitions is often dismissed as intractable, unimportant, and as a result uninteresting, see for example (Geurts 1999: 29-30). In this chapter I have attempted to demonstrate that this problem *is* important, that it is *not* intractable, and that it *is* interesting. I have focused exclusively on the definite determiner, and it would of course be desirable to extend the analysis to the full range of putatively presuppositional DPs and perhaps to presuppositional expressions in general. Since incorporating other determiners into the present analysis would almost surely introduce a number of additional complexities, I have focused on providing a concrete and immediately plausible proposal for understanding the relation between truth value intuitions and non-denoting definites. I have argued that falsity intuitions are the result of a particular interpretational mechanism. Robust intuitions arise when SWEPPs are evaluated in relation to QUDs and judging a SWEPP false is therefore equivalent to (tacitly) recognizing that the SWEPP is an intended, but not true, answer to the relevant QUD. In other words, judgments of falsity track whether a SWEPP is a member of a set of true answers. The behavior of contextualized SWEPPs seem to conform to this explanation and I have gestured towards a similar explanation of non-contextualized SWEPPs as well.

Given the available data, I have advocated a methodological change, but assuming that the data from non-contextualized cases is relevant, I have argued that the remaining key problem is determining how, in the non-contextualized cases, speakers select particular questions. So, in conclusion, much work remains, but hopefully the proposal here can serve as a guiding explanatory strategy.

CHAPTER III

GHOSTS, MURDERERS, AND THE SEMANTICS OF DESCRIPTIONS

1 Introduction

Suppose Bertrand falsely believes that there is a king of France and that the king of France is bald. If so, a speaker can truly and felicitously report Bertrand's belief by asserting (82).

(82) Bertrand believes that the king of France is bald.

The gospel among proponents of Russell's (1905) analysis of definite descriptions is that Russell's analysis provides a straightforward way of assigning adequate truth conditions to sentences such as (82), viz. sentences with non-denoting descriptions embedded in the scope of a propositional attitude verb. This assignment of truth conditions relies on two important assumptions. First, the assumption that definite descriptions *assert* the existence of a unique individual satisfying the descriptive content, i.e. 'the F is G ' is analyzed as an existentially quantified formula, namely the formula below.¹

$$\exists x[F(x) \wedge \forall y[F(y) \rightarrow x = y] \wedge G(x)]$$

¹ I assume that propositions are the objects of assertion, but in order to avoid excessively convoluted language I will often talk about what e.g. *the description asserts* and about *the asserted content of the description*. This is to be understood only as shorthand for what the description standardly contributes to the asserted content of the sentence in which it occurs. Similarly, I sometimes talk about particular expressions triggering presuppositions. This is to be understood in the following sense: When these expressions occur in a sentence S , S expresses a proposition only if some other proposition (the

Second, the assumption that *de dicto*—*de re* ambiguities, which arise when modal and quantificational expressions interact, are *scope* ambiguities; when a modal takes scope over a quantificational expression, the resulting interpretation is *de dicto* and when the embedding order is reversed, the resulting interpretation is *de re*. Hence, in the context where Bertrand falsely believes that there is a bald king of France, the true reading of (82) is captured by analyzing the definite description as embedded in the scope of the propositional attitude verb, cf. (83). This yields an interpretation where the existence of a unique French king is part of Bertrand's belief.²

$$(83) \quad \text{BEL}_b[\exists x[\text{king-of-France}(x) \wedge \forall y[\text{king-of-France}(y) \rightarrow x = y] \\ \wedge \text{bald}(x)]]$$

The Russellian analysis thus predicts that the sentence in (82) is true if and only if Bertrand believes the following conjunction: there is a (unique) king of France and he (the king of France) is bald. That is, relying on a minimal number of fairly uncontentious assumptions, the Russellian analysis succeeds in explaining why some sentences containing non-denoting descriptions have straightforwardly true interpretations.³

But against this frequently emphasized virtue of Russell's theory, Heim (1991), Kripke (2005), and Elbourne (2005, 2010) argue that Russell's analysis makes correct predictions only when the propositional attitudes in question are doxastic. That is, when the definite description is embedded under a *non-doxastic* attitude verb, the Russellian predictions are intuitively incorrect. Consider the example below.

$$(84) \quad \text{Hans wants the ghost in his attic to be quiet tonight.} \quad (\text{Elbourne 2010})$$

(84) has both a narrow scope (*de dicto*) and a wide scope (*de re*) interpretation, yet neither interpretation appears to provide the correct truth conditions.

$$(85) \quad \text{DES}_h[\exists x[\text{ghost-in-attic}(x) \wedge \forall y[\text{ghost-in-attic}(y) \rightarrow x = y] \\ \wedge \text{quiet}(x)]]$$

presupposition) is common ground. I discuss the notions of presuppositions and common grounds in detail later.

² I abbreviate '*a* believes/wants/hopes that ϕ ' as ' $\text{BEL}_a\phi$ ', ' $\text{DES}_a\phi$ ', ' $\text{HOPE}_a\phi$ ' respectively. I leave the semantics of these attitude verbs unspecified for now and focus instead on the semantics of the complement clauses. I will assume throughout that propositional attitude verbs are modals, in particular quantifiers over possible worlds. These assumptions are discussed in detail later.

³ See e.g. Neale (1990: 27) and Ludlow (2007: 3.2)

On the narrow scope analysis above, the Russellian analysis predicts that (84) is true if and only if Hans desires the truth of the following conjunction: there is a (unique) ghost in Hans' attic and the ghost is quiet. However, this is intuitively inconsistent with the most natural interpretation of (84). The truth of (84) is compatible with Hans having no desire that there be a ghost in his attic.

But now consider the wide scope interpretation.

$$(86) \quad \exists x[\text{ghost-in-attic}(x) \wedge \forall y[\text{ghost-in-attic}(y) \rightarrow x = y] \\ \wedge \text{DES}_h[\text{quiet}(x)]]$$

Here the Russellian analysis predicts that (84) is true if and only if there *is* a ghost in Hans' attic and he wants it to be quiet. The problem is that the truth of (84) does not appear to saddle the speaker with such a spurious ontological commitment. The speaker can truly report Hans' desires using (84) even if there, as a matter of fact, are no ghosts.

Now, the meaning of the definite is analyzed above in terms of the unary quantifiers of first-order logic and it is well known that these are inadequate for a whole range of determiners of natural language. And since Russell's original analysis of 'the *F* is *G*' is syncategorematic (i.e. no constituent in Russell's proposed 'logical form' corresponds to either 'the' or 'the *F*'), most proponents of Russell's analysis now maintain that the analysis should be cast in terms of generalized (or rather restricted) quantifiers.⁴ Neale (1990) thus adopts the notation in (87) to represent the 'logical form' of 'the *F* is *G*' and its (Russellian) truth conditions are then stated as in e.g. (88). The definite determiner is now analyzed as a relation between two sets. Below, the horizontal lines ' | ' indicate set cardinality.

$$(87) \quad [\text{The } x: F(x)](G(x))$$

$$(88) \quad [\text{The } x: F(x)](G(x)) \text{ is true if and only if } |F| = 1 \wedge |F \cap G| \geq 1$$

What is important to emphasize here is that the notation in (87) serves only to indicate that the definite article is analyzed as a restricted quantifier. This notational convention is supposed to better reflect that 'the *F*' and 'the' are now treated as proper syntactic and semantic units.⁵ In other words,

⁴ For the now standard analysis of generalized quantifiers in natural language, see e.g. Barwise and Cooper (1981) and Keenan and Stavi (1986).

⁵ Proper syntactic units in the sense that the logical form now reflects the syntactic structure of a determiner combining with a sister node to form a determiner phrase (DP) which is then combined with a verb phrase (VP) to form a sentence (S). And proper

[...] this is not to propose an alternative to Russell's theory; it is just to find a more congenial method of stating it. [the $x: Fx](Gx)$ is *definitionally equivalent* to $(\exists x)((\forall y)(Fy \equiv y = x) \ \& \ Gx)$. (Neale 1990: 45).

That is, if definite descriptions are analyzed using restricted quantifiers, then when the truth conditions of the complement clauses of (82) and (84) are computed, the Russellian analysis continues to make the right prediction for (82), cf. (89) below, but it also continues to make incorrect predictions for sentences such as (84), cf. (90) below.

$$(89) \quad \text{BEL}_b[|\text{King-of-France}| = 1 \wedge |\text{King-of-France} \cap \text{Bald}| \geq 1]$$

$$(90) \quad \text{DES}_h[|\text{Ghost-in-Attic}| = 1 \wedge |\text{Ghost-in-Attic} \cap \text{Quiet-Tonight}| \geq 1]$$

That is, (89) is predicted to be true if and only if Bertrand believes (a) that there is exactly one king of France and (b) that the set of bald French kings is equal to, or greater than, one. This appears to correctly capture the truth conditions of (82) just as (83) did before. In contrast, (90) is predicted to be true if and only if Hans desires (a) that there is exactly one ghost in his attic and (b) that the set of quiet ghosts in his attic is equal to, or greater than, one — which is equivalent to the result with (85) above. Again, this is not the correct prediction since (84) is perfectly compatible with Hans not having this desire.

Hence, if one favors the restricted quantifier analysis of definite descriptions (which one arguably should), one still faces the same dilemma, namely that the analysis yields correct predictions for doxastic attitude verbs but incorrect predictions for non-doxastic attitude verbs. For various presentational purposes, I use the standard unary quantifiers in my quasi-formal representations for the remainder of this chapter. I should however emphasize that the discussion to follow could have used restricted quantifiers and restricted quantifier notation instead with the same result.

The problem outlined above now prompts the following question: Why is there this divergence in the predictions of the Russellian analysis as regards doxastic and non-doxastic attitude verbs? The answer is the Russellian anal-

semantic units in the sense that both the determiner itself and the DP can now be assigned independent meanings, cf. below.

- i. $\llbracket \text{The } F \rrbracket = \lambda G_{\langle e,t \rangle} \cdot |F| = 1 \wedge |F \cap G| \geq 1$
- ii. $\llbracket \text{The} \rrbracket = \lambda F_{\langle e,t \rangle} \cdot \lambda G_{\langle e,t \rangle} \cdot |F| = 1 \wedge |F \cap G| \geq 1$

ysis' failure to distinguish between the meanings, i.e. the truth conditions, of sentences such as (91) and (92).⁶

- (91) The ghost is quiet.
- (92) There is a unique ghost and it is quiet.

Since definite descriptions are assumed to *assert* the existence of a unique individual satisfying the description, the Russellian analysis treats (91) and (92) as having the same meaning, i.e. as having the same truth conditions. And this appears to be unproblematic in standard extensional cases since (91) and (92) are true in precisely the same contexts, viz. contexts that contain exactly one (perhaps maximally salient) ghost which is quiet. Moreover, in intensional contexts of the doxastic variety, the putative difference between (91) and (92) again seems inessential. Since it is a precondition on the truth of (82) that Bertrand *believes* that there is a unique king of France, then if *c* is a context where (82) is true, *c* is also a context where (93) is true.

- (82) Bertrand believes that the king of France is bald.
- (93) Bertrand believes that there is a unique king of France and that he is bald.

However, it is not intuitively clear that sentences such as (82) express existential beliefs on the part of the subject. It might seem more natural to understand (82) as *presupposing* something about Bertrand's existential beliefs rather than as asserting it. But when the report in (82) is de dicto, this presupposition is a precondition on the truth of (82), and so building the presupposition into the asserted content of (82) helps output the correct truth conditions.

The problem arises when definite descriptions are embedded in the scope of attitude verbs that do not express beliefs, i.e. non-doxastic attitude verbs. While a sentence such as (84) also appears to presuppose an existential belief on the part of the subject, for example that Hans *believes* that there is a ghost

⁶ An anonymous referee suggested that if one assumes that the logical form of 'the *F* is *G*' is $[\text{The } x: F(x)](G(x))$, then one is not committed to assuming that (91) expresses the same proposition as (92). But since this logical form is only a convenient piece of notation which, according to e.g. Neale (1990), is *definitionally equivalent* to the logical form given by Russell, and since what Russell called 'logical form' simply is what Russell took to be the truth conditions, it is not entirely clear to me how spell out the details of this idea. When that is said, even if one was not committed to the claim that e.g. (91) and (92) express the same proposition, the problem that I am describing here is a problem about truth conditions. I.e. the problem for the Russellian analysis is its failure to distinguish the truth conditions of (91) from the truth conditions of (92).

in his attic, such an existential belief cannot be captured by assuming that the definite description asserts existence. Instead, when the attitude in question is non-doxastic, e.g. a desire, a hope, or a fear, the relevant sentences are predicted to express that the subject of the attitude verb *wants*, *hopes*, or *fears* that there *exist* a particular individual satisfying the description. And this is the cause of the incorrect predictions.

To further illustrate the problem, if a speaker intended to communicate that Hans really *wants there to be* a ghost in his attic and for it to be quiet, it seems implausible that the speaker would succeed by simply asserting (84). Instead, the speaker would have to use a sentence such as (94).

- (84) Hans wants the ghost in his attic to be quiet tonight.
- (94) Hans wants there to be a unique ghost in his attic and for it to be quiet.

This problem concerning assertions of existence is however not exclusive to cases where a definite description is embedded under a propositional attitude verb. The problem also arises with definite descriptions in conditionals. To illustrate, consider (95) and (96) below.

- (95) If the ghost in my attic is quiet tonight, I might finish my thesis.
- (96) If there is a unique ghost in my attic and it is quiet tonight, I might finish my thesis.

If definite descriptions assert existence, as the Russellian analysis maintains, and hence are shorthand for complex quantificational statements, one should be able to effectively paraphrase (95) with (96). But this is intuitively incorrect. If a speaker asserted (95), it would be natural to assume that she is *presupposing* (rather than asserting) that there is a ghost in her attic. That is, it would *not* be possible to interpret the speaker as simply contemplating a hypothetical scenario where there happens to be a ghost in her attic. And, as Elbourne (2010) observes, this intuition is further supported by the felicity of utterances of sentences such as (97) and the infelicity of utterance of sentences such as (98).

- (97) # I'm not sure there are any ghosts in my attic. If the ghost in my attic is quiet tonight, I might finish my thesis.
- (98) I'm not sure there are any ghosts in my attic. If there is a unique ghost in my attic and it is quiet tonight, I might finish my thesis.

Now, one might be inclined to defend the Russellian analysis by arguing that the definite description in (95) is naturally read as taking wide scope. This would straightforwardly explain why (97) sounds contradictory. But

there are at least two significant problems with this explanation. First, given the Russellian analysis, one would think that both wide and narrow scope readings of (95) should be available, but this is not consistent with the data.⁷ There just does not seem to be any narrow scope reading of the definite description in (95), i.e. (96) is not a possible interpretation of (95). The Russellian would therefore have to maintain that definite descriptions take *obligatory* wide scope in conditionals.⁸ Second, it is widely accepted that antecedents of conditionals are so-called *scope islands*, i.e. syntactic constituents from which syntactic movement is prohibited, see e.g. May (1985) and Szabolcsi (2001). Making the assumption that definite descriptions take obligatory wide scope would therefore violate this well-established syntactic principle. In sum, it seems that there is a significant problem here for the standard Russellian analysis of definite descriptions.⁹

2 The Existence Problem

2.1 Problem and Plan

What the right semantic analysis of definite descriptions is and what existential commitments these descriptions incur are issues that have traditionally divided philosophers of language and linguists into roughly two groups — those siding with Russell and those siding with Frege and Strawson. While there are a number of considerations relevant to answering these questions, the aim of the present chapter is to illuminate them by answering the more specific question below.

THE EXISTENCE PROBLEM

What existence commitments do definite descriptions incur and how

⁷ An anonymous referee notes that one needs not assume that a reading is available simply because it can be generated by the syntax. For example, one could supplement the syntax with filtering mechanisms which would then eliminate the unavailable readings. While I agree with the referee's point here, I do think that it leaves an explanatory gap. The narrow scope readings are available with most other, if not all, quantificational expressions, so it seems to me that the proponents of Russell's analysis owes an explanation for why these readings are not available with definite descriptions. Moreover, if the syntax were supplemented with relevant constraints to filter out the undesirable readings, one should think that such constraints would need to have independent syntactic motivation in order to be explanatory. And it is not obvious what would motivate such constraints.

⁸ And for an in-depth discussion (and refutation) of this suggestion, cf. Elbourne (2010)

⁹ I discuss a range of responses on behalf of the Russellian analysis in subsequent sections.

must these be analyzed if the apparently incorrect predictions manifested by sentences such as (84) and (95) are to be avoided?

The chapter is divided into two sections. In the first, I discuss whether the problematic predictions can be avoided by adopting a presuppositional analysis of definite descriptions, in particular the referential Frege/Strawson analysis (to be explicated in the following sections). Several researchers, cf. Heim (1991) and Elbourne (2005, 2010), maintain that the Frege/Strawson analysis has a significant advantage over the Russellian analysis, but I will demonstrate that with regards to the existence problem, the Frege/Strawson analysis has serious limitations. In particular, I will argue that the existence problem, as it is standardly framed, is a much more general problem than normally assumed. Next, elaborating on Elbourne (2005, 2010), I show that the existence problem is not a result of complexities introduced by propositional attitude verbs and that attempts to defend the Russellian analysis using arguments of that nature fail. I conclude by arguing that both the orthodox analyses of definite descriptions, namely the Russellian analysis and the Frege/Strawson analysis cannot solve the existence problem, i.e. cannot generally output correct truth conditions for sentences such as (84) and other related sentences.

In the second section, I outline a range of conditions that any semantic analysis must satisfy if the existence problem is to be avoided. I sketch a dynamic semantic analysis of descriptions that satisfies these conditions. I show that my proposed analysis provides a solution to the problem as it arises in conditionals and I explain what would be needed to generalize the solution to propositional attitude verbs. I also show that my analysis is more generally adequate than both of the orthodox analyses.

2.2 *The Presuppositional Solution*

As demonstrated in the introduction, the existence problem is a product of a semantic analysis that fails to distinguish the meaning of (91) from the meaning of (92). One might therefore think that an analysis which does distinguish would provide a simple solution. One such analysis is a *presuppositional* analysis.

On the standard view, presuppositions are constraints on language use. For example, if an expression ϕ triggers the presupposition that χ , it is felicitous for a speaker to use ϕ in a context c only if χ is antecedently established in c , viz. presupposed. On a presuppositional analysis of definite descriptions, existence of a unique individual is therefore assumed to be a presupposition, viz. a constraint on its use rather than an asserted component of its meaning. This means that a sentence of the form ‘the F is G ’ is not analyzed as making

the existentially quantified claim that there is a unique *F*, but rather as a sentence whose use is licensed only if the existence of a unique individual who is *F* is already established in the discourse context — only if it is *common ground* that there is a unique *F*.¹⁰

So, how is a presuppositional analysis supposed to provide a solution to the existence problem? The intuitive thought is that if definite descriptions are analyzed as merely *presupposing* the existence of a unique *F*, the existence of this unique *F* will not be part of the content of an attitude verb in whose scope the description is embedded — and hence the existence problem should fail to arise. This is, crudely speaking, the view of Heim and Elbourne. However, as it turns out, matters are not quite this simple. This section is devoted to demonstrating why.

2.2.1 Presuppositions and Partial Functions

There are several ways to formally explicate the existence commitments that definite descriptions incur. Let's begin by clearly distinguishing three such ways.

- (O1) Analyze definite descriptions as *asserting* existence. For example, assume that a definite description contributes an existentially quantified expression to the truth conditions of the sentence in which it occurs. A speaker who asserts that the *F* is *G* thus effectively asserts that a (unique) *F* exists.
- (O2) Analyze definite descriptions as *presupposing* existence, i.e. assume that the asserted content contains no existential quantifier and instead that an existential claim is antecedently established in the common ground. This is what licenses the use of the definite description. So, a speaker who asserts a sentence such as 'the *F* is *G*' does not thereby assert that a unique *F* exists — the speaker presupposes that this is common ground.¹¹

¹⁰ This conception of presuppositions is due, in its essentials, to the pioneering work of Stalnaker (1970, 1974, 1978).

¹¹ Notice that existence is logically entailed even on the analysis in (O2). A sentence ϕ containing a presuppositional expression ψ can be true only if the presuppositions of ψ are satisfied and hence if logical consequence is defined in terms of truth preservation, it follows that if ϕ contains ψ and ϕ is true, ϕ logically entails the presuppositions of ψ , or ϕ could not be true.

- (O3) Analyze definite descriptions as both presupposing *and* asserting existence.

The Russellian analysis captures the existence commitments of definite descriptions as described in (O1). Definite descriptions are analyzed as contributing to the truth conditions an existentially quantified formula and sentences containing definite descriptions are thus analyzed as *asserting* existence. That is, existence is part of the *asserted* content. This is, as we saw earlier, the reason that the Russellian analysis yields correct truth conditions for (82) but also the reason that it runs into the existence problem.

Two different types of semantic content can be distinguished on a presuppositional analysis, namely the *presupposed* content and the *asserted* content. The challenge for a putative presuppositional solution to the existence problem is to provide a formal analysis of these contents which avoids the existence problem and which in general succeeds in outputting correct truth conditions for sentences containing definite descriptions.

Now, since presuppositions are characterized as constraints on use — in the sense that they impose a contextual requirement that the presupposition must be common ground — existence presuppositions are standardly captured using *partial functions*.

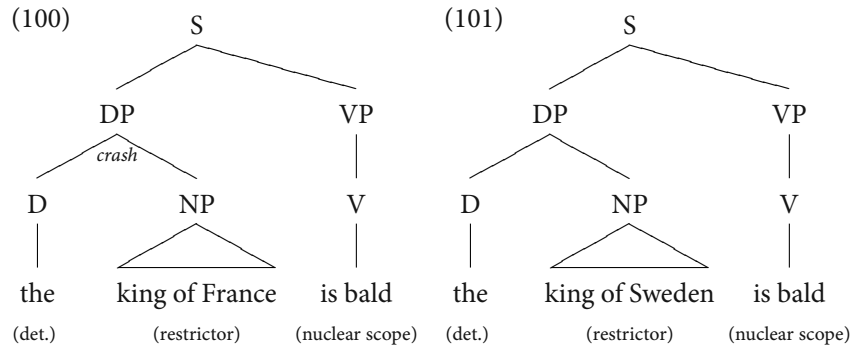
Let's assume that the input for a semantic computation is an LF (in the syntactic sense used in linguistics) and that semantic values are computed on the basis of pairwise functional application of lexical constituents. The definite determiner is then treated as a function that is defined only if its input argument F is a singleton set — and otherwise undefined. This captures the above constraint on use since a sentence containing a definite description is predicted to express a proposition (which, let's assume, is a function from worlds to truth-values) only if F has exactly one member. F is also often called *the restrictor*. In less technical terms, 'the F is G ' can be assigned a classical truth value only if there is a unique F . And so, treating the definite determiner as a quantificational determiner, viz. typing it as a function from two predicate sets to a truth value, we can represent this presuppositional requirement formally as in (99).¹²

$$(99) \quad \llbracket \text{the} \rrbracket = \lambda P: \exists! x P(x). \lambda Q. \exists x [P(x) \wedge \forall y [P(y) \rightarrow x = y] \wedge Q(x)] \langle \text{et}, \langle \text{et}, t \rangle \rangle$$

The expression $\exists! x P(x)$ is to be understood as a restriction on the initial λ -function. The domain of this function is the set of predicates of type $\langle \text{e}, t \rangle$

¹² On this analysis, the definite determiner is of type (O3), i.e. 'the' both presupposes *and* asserts existence. I begin by considering this lexical entry in order to highlight that a *presuppositional* analysis in itself is insufficient for solving the existence problem.

(i.e. a set of sets of individuals) but it is defined only for arguments (i.e. sets of individuals) that have exactly one member. This condition could thus also have been stated as a cardinality constraint, e.g. $\lambda P: |P|=1$. The presuppositional nature of a definite description is captured in the sense that unless the restrictor set contains precisely one individual, the semantic computation of the sentence crashes. As a crude illustration, consider (100) and (101).



The restrictor set in (100) is empty and this means that the determiner is undefined for its restrictor argument. The computation of D+NP therefore fails and as a result the sentence fails to express a proposition, viz. a truth value cannot be computed for the sentence as a whole. In contrast, the restrictor set in (101) contains exactly one individual which means that the computation of D+NP succeeds and that the sentence as a whole can be assigned a truth value. What the lexical entry in (99) achieves is thus to impose the requirement on computations of definite descriptions that only if a unique individual satisfying the restrictor of the description exists will the sentence as a whole express a proposition. In other words, for a sentence such as ‘the *F* is *G*’ to express a proposition, it is a *precondition* that a unique *F* exists.

However, we can now see why adopting a presuppositional analysis *in and of itself* is not sufficient for solving the existence problem. This becomes particularly clear when we consider the truth conditions generated by a computation of (101).

$$\begin{aligned}
 & \left[\lambda P: \exists! x P(x). \left[\lambda Q. \exists x [P(x) \wedge \forall y [P(y) \rightarrow x = y] \wedge Q(x)] \right] (\lambda z. \text{bald}(z)) \right] \\
 & (\lambda v. \text{KoS}(v)) \\
 \Rightarrow & \left[\lambda Q. \exists x [\text{KoS}(x) \wedge \forall y [\text{KoS}(y) \rightarrow x = y] \wedge Q(x)] \right] (\lambda z. \text{bald}(z)) \\
 \Rightarrow & \exists x [\text{KoS}(x) \wedge \forall y [\text{KoS}(y) \rightarrow x = y] \wedge \text{bald}(x)] \quad (101')
 \end{aligned}$$

The truth conditions in (101') are identical to the truth conditions predicted by a standard Russellian analysis. In other words, if (101') is embedded under an attitude verb or in the antecedent of a conditional, we will make the same problematic predictions that Russell's analysis was previously shown to make.

It is now clear that the lexical entry in (99) — and more generally any analysis on which existence ends up being *asserted* — will invariably retain the prediction that for e.g. (84) to be true, Hans must desire that there is a ghost in his attic and this was of course one of the predictions we were aiming to avoid. The upshot is this: Using partial functions to ensure that the existence of a unique *F* is a *constraint* on use is not, in itself, sufficient for solving the existence problem. The analysis must also ensure that existence does not end up being *asserted*.

In conclusion, if the solution to the existence problem must rely on an analysis which distinguishes the meaning of an existential-there sentence such as (92) from the meaning of a sentence with a definite description in subject position such as (91), and if this is best achieved by a presuppositional analysis, then only a *purely* presuppositional analysis in the sense of (O2) above can work. If existence presuppositions are captured in terms of partial functions, viz. in terms of partial interpretability, the challenge is how to formally characterize the *asserted* content without using expressions that assert existence.

2.2.2 *Presuppositions and Referential Terms*

On the so-called Frege/Strawson analysis of definite descriptions, definite descriptions are not only assumed to trigger existence presuppositions, i.e. assumed to be constrained in their use to contexts where it is common ground that there is a unique individual satisfying the restrictor, but also assumed to be *referential*. Expressions of the form 'the *F*' are thus assumed to *refer* to the unique individual who is *F*. This is typically captured formally using a lexical entry such as the following.¹³

¹³ This (or notational variants) is standard in contemporary semantics, see e.g. Elbourne (2005), von Stechow (2007), Heim and Kratzer (1998), Schwarz (2009). In some cases the lexical entries are embedded in an intensional situation semantics and thus amended to contain λ -abstractions over situation variables. This makes no difference to the points pursued here. It should also be emphasized that a referential treatment in this sense is compatible with an analysis of the definite determiner where it is typed (semantically) as a quantificational determiner, viz. a relation between two predicate sets, $\llbracket \text{the} \rrbracket = \lambda P: \exists! x P(x). \lambda Q. Q(xP(x))$. While there are both pros and cons to adopting a referential analysis of definite descriptions and this in general is a point of controversy,

$$(102) \quad \llbracket \text{the} \rrbracket = \lambda P: \exists! x P(x). \iota x P(x) \quad \langle \text{et}, e \rangle$$

Here the definite determiner is treated as a function from properties to individuals but again the function is partial: it is defined only for arguments (sets of individuals) whose cardinality equals one. The complex ι -expression is a referential term which refers to the unique member of the restrictor set F . Hence, a computation of the semantic value of the determiner and its argument can succeed only if the argument (the relevant set of individuals) has exactly one member. If the function is defined for its argument, the resulting ι -expression is effectively a constant which refers to this individual. Given this lexical entry, a computation of (101) now yields the truth conditions stated in (101'').

$$\begin{aligned} & [\lambda y. \text{bald}(y)]([\lambda P: \exists! x P(x). \iota x P(x)](\lambda z. \text{KoS}(z))) \\ \Rightarrow & [\lambda y. \text{bald}(y)](\iota x \text{KoS}(x)) \\ \Rightarrow & \text{bald}(\iota x \text{KoS}(x)) \end{aligned} \quad (101'')$$

What is important to notice is that (101'') does not *assert* that a unique individual who is a king of Sweden exists *and* that this individual is bald. Instead, (101'') asserts *of* the individual who is the unique king of Sweden that this individual is bald. This means that when this formula is embedded under an attitude verb, the existence problem should simply fail to arise. Similarly, in cases where a definite description is embedded in the antecedent of a conditional, e.g. (95), the prediction is that the conditional as a whole presupposes the existence of a French king. If this presupposition is satisfied, the antecedent of the conditional will simply contain a referential term referring to this individual (whoever it may be) and the semantic value of the conditional can then be unproblematically computed (according to whatever semantic analysis of conditionals one might favor). In sum, presupposing existence and treating the asserted content of the description as a referential term allows one to effectively mimic the result of wide-scoping without relying on any dubious syntactic assumptions. Moreover, one can now straightforwardly explain the contrast noted earlier between (97) and (98) (repeated below).

- (97) #I'm not sure there are any ghosts in my attic. If the ghost in my attic is quiet tonight, I might finish my thesis.
- (98) I'm not sure there are any ghosts in my attic. If there is a unique ghost in my attic and it is quiet tonight, I might finish my thesis.

I neglect discussing any of these issues here. I focus exclusively on showing why Heim's and Elbourne's favored analysis does not get to the heart of the existence problem.

An assertion of (95) is felicitous only if it is common ground that there is a unique ghost in the speaker's attic. Hence, if the speaker declares herself agnostic as to whether this *is* common ground, it is infelicitous for her to assert the sentence. This is the reason that (97) sounds contradictory. On the other hand, asserting (98) is acceptable simply because this sentence contains no presuppositional expressions and therefore have no restrictions imposed on its use.

Since the Frege/Strawson analysis appears to provide a simple and intuitive resolution of the existence problem, Heim (1991) and Elbourne (2005, 2010) both take the existence problem as providing evidence that the Frege/Strawson analysis is the correct analysis of definite descriptions. However, one more important assumption is needed if this solution is to work for *non-denoting* descriptions too.

2.2.3 *Presupposition Projection in Attitude Contexts*

On the Frege/Strawson analysis, the existence problem now fails to arise in cases where the definite description in fact refers, but the original problem case remains problematic.

(84) Hans wants the ghost in his attic to be quiet tonight.

There are no ghosts, so the computation of the complement clause of (84) crashes. The problem is that the Frege/Strawson analysis predicts that an utterance of a sentence such as (84) presupposes that there is a unique ghost and this prediction seems intuitively incorrect. It seems more natural to think that (84) only presupposes that Hans *believes* that there is a ghost. In fact, it seems that when a presuppositional expression occurs inside the scope of an attitude verb, the constraint that it *generally* imposes on its use is different from the constraint imposed in standard non-intensional contexts.

This behavior of presuppositional expressions was first emphasized by Karttunen (1974) who concluded that when presuppositions are triggered in the scope of a non-factive attitude verb, the presuppositions *project* into belief contexts. That is, while an assertion of (103) triggers the presupposition that Sue smoked — and hence requires that this is common ground — this is not so for (104).

(103) Sue stopped smoking.

(104) Bertrand $\left\{ \begin{array}{l} \text{hopes} \\ \text{believes} \\ \text{wishes} \end{array} \right\}$ that Sue stopped smoking.

(104) only presupposes that Bertrand *believes* that Sue smoked. To see this, we need only observe that it is infelicitous for a speaker to continue an utterance of (103) by uttering (105a-105b) whereas these continuations are perfectly felicitous for (104).

- (105) a. # However, I'm not sure she ever smoked.
b. # But she never smoked.

The presupposition triggered by (104) is not that Sue smoked but instead that the subject of the attitude, Bertrand, *believes* that she smoked. The discourse participants must mutually accept this *doxastic* claim in order for an utterance of (104) to be felicitous.¹⁴

When Karttunen's generalization is accepted, (84) is no longer predicted to presuppose that there is a unique ghost in Hans' attic. Rather, it is predicted to presuppose that Hans *believes* that there is. And since the definite description in (84) is analyzed as a referential term (which simply *refers* to the individual that Hans believes exists), the analysis is now in a position to predict that (84) has a true interpretation without incurring any unfortunate assertions of existence.¹⁵

2.2.4 Interim Conclusion

The Frege/Strawson analysis might appear to completely resolve the existence problem, but I now intend to demonstrate that sentences such as (84) and (95) are simple instances of a thoroughly general problem concerning the existence commitments incurred by natural language determiners. While the Frege/Strawson analysis *is* in a position to make correct predictions for cases involving definite descriptions (when combined with the aforementioned auxiliary assumptions), this result relies on exploiting a

¹⁴ This generalization is further supported by the infelicity of utterances such as (i) and (ii) below.

- (i) # Bertrand hopes that Sue stops smoking but believes she never smoked.
(ii) # Bertrand wants Sue to stop smoking but he's not sure she ever smoked.

¹⁵ Strictly speaking, this prediction requires both a suitable semantics for propositional attitude verbs and an account of the behavior of presuppositions in attitude contexts, see e.g. Heim (1992) for a thorough discussion. I discuss some of these problems briefly in later sections. However, if for example one accepts Karttunen's generalization and adopts Heim's (1992) ordering semantics for desire-verbs, this should suffice for making correct predictions in cases such as (84).

particular feature of the meaning of definite descriptions — a feature that cannot be exploited in cases involving various other determiners. When this is acknowledged, it becomes clear that the Frege/Strawson analysis fails to get to the heart of the existence problem. It also becomes clear that much more radical changes to our semantic system are needed, if a proper semantic analysis of the existence commitments incurred by various natural language determiners is to be devised.

2.3 Indefinite Descriptions: The Failure of the Referential Analysis

To demonstrate that the Frege/Strawson analysis fails to get to the heart of the existence problem, we need only consider a simple scenario.

Context I Suppose that three known murderers are on trial for their crimes and suppose that Bertrand has a desire that at least one of the three murderers be convicted. Now, let's assume that Bertrand has no specific individual in mind; if either of the three murderers are convicted, Bertrand's desire is satisfied. For example, suppose that Bertrand has a bet that at least one murderer will be convicted. But, let's also assume that Bertrand is a normal and rational individual who prefers that murders were never committed. Bertrand therefore has no desire that there be a murderer. His desire that a murderer be convicted is conditional on the antecedent belief that the individuals on trial are in fact murderers. Now suppose, on the night of the trial, I report this by uttering (106). (CI)

(106) Bertrand wants a murderer to be convicted tonight.

In the scenario described in (CI) my utterance of (106) is clearly both felicitous and true. But a familiar problem now resurfaces: On the standard (Russellian) analysis of indefinite descriptions, i.e. 'an *F*' or 'some *F*', these descriptions are analyzed as complex quantificational constructions. Since the indefinite description interacts with a modal (the propositional attitude verb), (106) has both a narrow scope (*de dicto*) interpretation and a wide scope (*de re*) interpretation.

(106) a. $DES_b[\exists x[murderer(x) \wedge convicted(x)]]$

On the narrow scope analysis above, the standard Russellian analysis predicts that (106) is true only if Bertrand desires the truth of the following conjunction: There is a murderer and he is convicted. However, this is directly inconsistent with the scenario described in (CI) — Bertrand has no desire that there is a murderer.

But the wide scope reading in (106b) makes an equally problematic prediction.

(106) b. $\exists x[\text{murderer}(x) \wedge \text{DES}_b[\text{convicted}(x)]]$

The prediction here is that (106) is true only if there *is* a murderer and Bertrand wants that murderer to be convicted. This is problematic for two different reasons: First, the truth of (106) does not in fact require that there is a murderer. The scenario could have been described so that the existence of murderers was nothing but a figment of Bertrand's imagination and yet (106) could be true. Second, Bertrand is now predicted to entertain a *de re* desire about a particular individual, but this is inconsistent with the scenario described in (CI). Bertrand has no desire pertaining to any specific murderer, because he does not care which murderer is convicted.¹⁶

The analogy between (106) and (84) is striking. In both cases, permutations of the scope taking expressions yield incorrect truth conditions and in both cases there is a strong intuition that existence is not *asserted* and hence not part of the attitudinal content. Instead it seems that in both cases existence is *presupposed*. In other words, (106), (84) and also (95), are instances of the same (existence) problem. For example, notice that the problem with indefinites, as manifested by (106), can also arise in conditionals. Consider the following scenario.

Context II Mary wants an apple pie for her birthday tomorrow and so she asks Bertrand if he would help bake one. While Bertrand is happy to help, Mary suddenly realizes that the stores are closed and that she is unsure whether they have the required ingredients. Now suppose Bertrand utters (107). (CII)

(107) I know we have flour, sugar, cinnamon, and nutmeg. If some apples in the pantry are ripe, we should be alright.

Here it seems incorrect to analyze the indefinite description in the conditional in (107) as *asserting* the existence of apples in the pantry; Bertrand

¹⁶ As far as I know, Hawthorne and Manley (forthcoming) were the first to argue that the existence problem generalizes to other determiners. However, demonstrating that such a generalization is *in fact* problematic requires a fairly meticulous contextual setup and a thorough diagnosis of the problem which Hawthorne and Manley do not provide. Moreover, they seem to assume that given that the problem generalizes, it can simply be dismissed. Yet, if anything, the generalization shows that this problem is much worse than previously assumed.

is not making the hypothetical claim that if *there are some apples in the pantry* and they are ripe, they can then bake the pie. Bertrand seems to be presupposing that there are apples in the pantry. To see this, observe the contrast in judgments about the following sentences.

- (108) # I know we have flour, sugar, cinnamon, and nutmeg, but I'm not sure we have any apples. If some apples in the pantry are ripe, we should be alright.
- (109) I know we have flour, sugar, cinnamon, and nutmeg, but I'm not sure we have any apples. If there are some apples in the pantry and they are ripe, we should be alright.¹⁷

(CII) shows that the existence problem can also arise with *indefinite* descriptions when these are embedded in conditionals and even if the consequences of treating the description in (107) as asserting existence leads to somewhat less problematic predictions than in the case of (95), it seems clear that this is not the correct analysis.

When we acknowledge that the existence problem can arise with indefinite descriptions, it is easy to see that it could also arise with many other determiner phrases, e.g. 'three *Fs*', 'many *Fs*', 'several *Fs*', 'most *Fs*' etc. In other words, the existence problem raises a *general* problem for more or less every determiner phrase — not just definite descriptions. And because of the general nature of the existence problem, it can now easily be demonstrated that the Frege/Strawson solution fails to get to the core of the existence problem.

On the Frege/Strawson analysis, 'the *F*' is analyzed as referential rather than quantificational and this is a feasible assumption with respect to definite descriptions simply because these descriptions pick out a *unique* individual. One can thus more or less unproblematically assume that a definite description *refers* to the unique individual who happens to satisfy the description. The advantage is that when the meaning of a definite description is treated as a referential term, one can embed the description inside the scope of an attitude verb or a conditional without incurring any existence problems.

However, this solution to the existence problem is simply exploiting the uniqueness of definite descriptions and cannot be extended to cases such as (106) or (107) for this reason. Even if we assume that the indefinite description in (106) presupposes, rather than *asserts*, the existence of a murderer (as immediate intuitions might suggest), we cannot analyze the

¹⁷ These cases are variants of cases discussed in von Fintel (1998) who presents a number of cases suggesting that weak determiners do, on occasion, trigger presuppositions.

asserted content of the description as a referential term for the following reasons: There are several murderers in the context, i.e. there is no *unique* murderer in the context, so if the indefinite ‘a murderer’ is analyzed as a referential term, which individual is this referential term supposed to pick out? Moreover, if the indefinite description is analyzed as a referential term, we would predict that Bertrand has a desire about a specific murderer which is directly inconsistent with the scenario described in CI. In CI, Bertrand’s desire is satisfied regardless of which murderer is convicted and in order to capture this, we must allow the indefinite description to range over multiple individuals, i.e. we must treat the indefinite description as a genuine quantificational expression. Any referential analysis of indefinite descriptions is thus by its very definition incapable of capturing that Bertrand’s desire is general rather than specific.

In conclusion, the strategy underlying the Frege/Strawson solution to the existence problem, i.e. characterizing the asserted content using referential terms, is incapable of providing a *general* solution to the existence problem. This strategy fails because the existence problem generalizes to cases where there is no uniqueness to exploit and where adequate truth conditions require quantificational force.¹⁸ In sum, the existence problem is a *general* problem about the existence commitments of natural language determiners and it seems reasonable to assume that it has a *general* solution.

Given the above observations, it seems clear that a general and uniform solution to the existence problem, viz. a uniform analysis of existence commitments which is adequate for definites but also indefinites, is going to require significant changes to our semantic system. Summarizing what we have observed so far helps illustrate what an adequate semantic analysis of e.g. definite and indefinite descriptions must look like.

- The analysis must be capable of distinguishing existential-there sentences from sentences with descriptions or determiner phrases in subject position. This means that the *asserted* contents cannot invariably be characterized using expressions which assert existence, i.e.

¹⁸ I should note here that the generalization problem also rules out an analysis on which the asserted content contains just a universal quantifier ranging over every individual satisfying the restrictor predicate. As regards definite descriptions, the partial function guarantees that the predicate set is singleton and as a result, universal quantification over that set in the asserted content would range only over a single individual without incurring existential commitments. However, for e.g. indefinite descriptions, this strategy fails miserably, since on such an analysis an utterance of (106) would be predicted to be equivalent to asserting that Bertrand wants every murderer to be convicted.

existential quantifiers. As shown in (84) and (106), when existence is asserted, the problem in question arises.

- The asserted contents cannot invariably be characterized using referential terms. If the asserted contents are characterized as such, this yields incorrect predictions when the descriptions or determiner phrases in question do not imply uniqueness, cf. (106) and (107).
- Even if indefinite determiners are analyzed as presuppositional, these must be analyzed as ranging over multiple individuals. If they are not, this yields inadequate predictions, cf. (106) and (107).

This raises the following question: *Is there any viable way of characterizing what these descriptions contribute to the asserted content that avoids the existence problem?* Or in other words, is there a viable analysis which can plausibly satisfy each of the conditions stated above? I think that there is and in the second half of this chapter, I outline such an analysis. However, before proceeding to the positive part of this chapter, I want to first discuss some possible responses on behalf of the proponents of a Russellian analysis.

2.4 *The Vagaries of Attitude Verbs*

Given the rather serious obstacles to solving the existence problem explicated above, a standard defensive maneuver on behalf of Russell is to push the problem back to the attitude verbs. A proponent of Russell's theory could for example attempt to defuse the problem by stressing one or more of the points below.

Response 1	The putative problem posed by the existence problem relies on a mistaken logical assumption, namely that propositional attitude verbs are closed under logical consequence. Since it is widely known that they are not, there simply <i>is</i> no problem (and there never were) for Russell's analysis of definite or indefinite descriptions. (R1)
Response 2	The putative problem is not the result of an inadequate semantic analysis of definite or indefinite determiners but rather a result of complexities which it is well known that propositional attitude verbs introduce. (R2)
Response 3	The existence problem appears to pose a problem for Russell's analysis only because the full range of technical resources at our disposal — technical resources necessary for dealing with various intensional constructions — have not yet been utilized. If propositional attitude verbs are treated as quantifiers over worlds, there are good reasons to assume that such worlds must be ordered (according to some relevant parameter) and contextually restricted. Once these assumptions are technically implemented, we are in a position to dispel the puzzle. (R3)

Elbourne (2010) has already convincingly argued that the existence problem does not rely on the incorrect assumption that propositional attitude verbs are closed under classical consequence. But since this type of response is frequently voiced, I reiterate some of Elbourne's observations here and add a couple of novel ones. This, I hope, will conclusively disarm responses of type (R1).¹⁹

The response in (R1) starts with the following observation: the sentence in (110) logically entails that there are men, but this logical entailment fails when (110) is embedded in the scope of an attitude verb, cf. (111).

(110) There are honest men. (Kaplan 2005)
 \models There are men.

(111) Diogenes wondered whether there are honest men. (Kaplan 2005)
 $\not\models$ Diogenes wondered whether there are men.

Examples such as these are supposed to show that propositional attitude verbs are not closed under a classical consequence relation and from this

¹⁹ This response is tentatively endorsed by Kaplan (2005: 985) and emphatically stressed by Neale (2005: 846, 2007: 89-91). It has also often been raised in Q&As when I have presented this material.

it is concluded that from a sentence such as (94) — a simple Russellian paraphrase of (84) — one cannot infer (112).

(84) Hans wants the ghost in his attic to be quiet tonight.

(94) Hans wants there to be a unique ghost in his attic and for it to be quiet.

(112) Hans wants there to be a unique ghost in his attic.

The problem with this response is that it simply does not address the right objection.²⁰ The problem for Russell's analysis raised here is not a problem about entailments, it is a problem about the predicted truth conditions and the predicted asserted content. To illustrate, suppose a speaker utters the sentence in (113) but continues her speech by uttering either (113a) or (113b).²¹

(113) I am not sure there are any ghosts in my attic.

a. I want there to be exactly one ghost in my attic and for it to be quiet tonight.

b. # I want the ghost in my attic to be quiet tonight.

As Elbourne correctly points out, the continuation in (113a) is consistently judged felicitous by native speakers of English whereas discourse the continuation in (113b) is reported to sound contradictory and hence infelicitous. The problem here is that if (113a) is nothing but a paraphrase of (113b), as the proponents of the Russellian analysis maintain, both of these discourse continuations should be felicitous — yet they are not. This means that if the type of response given in (R1) is to be relevant to the objection raised here, it must be reformulated, i.e. it should be,

[...] the superficially similar but ultimately distinct claim that the utterance of [(113a)] is consistent with [me] being unsure that there is a ghost in [my] attic, whereas the utterance of [(113b)] is not. But [(113a)] is just a Russellian paraphrase of [(113b)]. So the Russellian paraphrase cannot be correct. (Elbourne 2010: Sec. 5)

²⁰ Moreover, whether the assumption that propositional attitudes are not closed under a classical consequence relation suffices to show that inferences from (84) to (112) are illicit is not completely obvious. The mere fact that classical entailments are not invariably preserved in propositional attitude contexts in no way mandates the conclusion that such inferences are *never* licensed. Ultimately, that will depend on what the "correct" closure conditions for attitude verbs are. And without any explicit contextual clues, if some speaker were to literally assert the sentence in (94), it would not seem totally unreasonable for her interlocutors to infer (112).

²¹ This argument is essentially the argument given in Elbourne (2010: Sec. 5).

In other words, the meaning of the complement clause of (113a) cannot be equivalent to the meaning of the complement clause of (113b). And so, that propositional attitude verbs are not closed under classical consequence is orthogonal to the existence problem. Even if we assume that they are not, this cannot explain why the *truth conditions* predicted by Russell's analysis for sentences such as (84) are intuitively incorrect.

The response in (R2) is a more general version of (R1) and the primary problem with this response is that the existence problem is not restricted to cases where a determiner phrase is embedded in the scope of a propositional attitude verb. As has already been emphasized several times, the exact same type of problem arises with conditionals.

Moreover, if the existence problem is simply blamed on the complexities of propositional attitude ascriptions, it would effectively be impossible to construct an adequate semantics for various attitude verbs — at least if the semantics of definite descriptions is assumed to be the standard Russellian semantics. First, remember that the Russellian analysis is incapable of distinguishing the meanings, viz. the truth conditions, of (91) and (92) (repeated below).

(91) The ghost is quiet.

(92) There is a unique ghost and it is quiet.

Suppose, as is standard, that attitude verbs take propositions as their arguments. If so, it is entirely unclear how *any* semantics of attitude verbs is supposed to avoid the existence problem: If an attitude verb Σ takes as its argument a proposition p , we need $\Sigma(p)$ to output one set of truth conditions when p is the proposition expressed by (91) and to output *a distinct* set of truth conditions when p is the proposition expressed by (92). But since (91) and (92) have identical truth conditions on the Russellian analysis, and thus express the same proposition, it seems that we need Σ to take identical arguments and yet *systematically output distinct truth conditions*. It is rather difficult to envision how this is supposed to be possible. And again, this simply highlights the importance of adopting a semantic analysis which can distinguish the meaning of (91) from the meaning of (92).

While the response in (R3) seems to be a commonplace reaction to the existence problem, it remains somewhat elusive to me what the concrete response is supposed to be. But here is a crude outline of what I suspect proponents of (R3) have in mind: Suppose we assume a Kratzerian (1977, 1981) semantics for modals where modals are quantifiers over possible worlds and modal statements depend on conversational backgrounds, namely *modal bases* and *ordering sources*. Moreover, let's also assume that a sentence such as $\text{DES}_a(\phi)$ is true only if at the worlds determined by the modal base, the best

ranked worlds are ϕ -worlds (where the rank of these worlds is determined by the ordering source). Following Heim (1992), we can assume that the modal base for DES depends on the subject of the attitude verb and that in (84) the modal base is the doxastic state of the subject, i.e. a set of worlds compatible with the beliefs of the subject. This set of worlds is then ordered according to the ordering source which is the subject's desires.

This is supposed to yield the following truth conditions: (84) is true if and only if the best ranked worlds (determined by the modal base and the ordering source) are such that the formula embedded inside the scope of the attitude verb in (85) is true at those worlds.

$$(85) \quad \text{DES}_h[\exists x[\text{ghost-in-attic}(x) \wedge \forall y[\text{ghost-in-attic}(y) \rightarrow x = y] \\ \wedge \text{quiet}(x)]]$$

And so, if the modal base is determined on the basis of Hans' beliefs, every world in the modal base will be a world where there is exactly one ghost in Hans' attic. And the best ranked worlds will be worlds where that ghost is quiet. These appear to be the correct truth conditions.

Unfortunately, this proposed analysis leaves one crucial question unanswered: In virtue of what is the set of relevant worlds restricted to include only worlds where there is precisely one ghost in Hans' attic? If one simply assumes that the modal base is Hans' doxastic states, then the question is this: Where in (84) do we locate the information that Hans *believes* that there is a ghost in his attic? It is difficult to see how one could plausibly answer this question without appealing to some presuppositional requirement triggered by the sentence itself.

And this now raises a different question, namely what the role of the existential quantifier in (85) is supposed to be. Since the formula in (85) expresses exactly the predicted truth conditions of (94) (repeated below), how are we supposed to predict that the truth conditions of (94) are distinct from the truth conditions of (84)?

- (84) Hans wants the ghost in his attic to be quiet tonight.
- (94) Hans wants there to be a unique ghost in his attic and for it to be quiet.

Again, the problem is the inability of the Russellian analysis to distinguish existential-there sentences from sentences with a definite description in subject position. Yet, if one then were to make the natural move and adopt a presuppositional analysis of the definite description, we are immediately back to the problem of adequately characterizing the asserted contents of these descriptions. And it is exactly *this* problem, viz. characterizing the

asserted content, that I have argued cannot be solved by either a standard Russellian nor the Frege/Strawson analysis.

This concludes my discussion of the existence problem in relation to the Russellian analysis and the Frege/Strawson analysis. I should emphasize that the existence problem is not a problem only for proponents of these analyses.²² It is also a problem for strictly speaking non-Russellian analyses, e.g. Szabó (2000), Ludlow and Segal (2004), and Hawthorne and Manley (forthcoming). Any analysis which maintains that a definite or an indefinite description uniformly *asserts* existence of a (unique) individual is subject to incorrect predictions in both non-doxastic attitude contexts and conditionals.

This concludes the predominantly negative part of this chapter. In the following, I outline the requirements that a solution to the existence problem must meet and I sketch an analysis which succeeds in doing exactly this.

3 Towards a Solution to the Existence Problem

In the previous section, I outlined several conditions for a solution to the existence problem. After summarizing these conditions, I sketch a semantic analysis of definite and indefinite descriptions that succeeds in satisfying these conditions. However, since this proposed analysis is embedded in a dynamic semantic framework, I must digress with a short introduction to this semantic framework. After this introduction, I explain how the proposed analysis provides a solution to the existence problem as it arises in conditionals. I then move to a discussion of attitude verbs and I demonstrate that my analysis is the way forward if the existence problem is to be avoided, but I also show that a full-fledged solution to this problem requires the development of a general dynamic semantics for attitude verbs and a systematic analysis of anaphora in intensional contexts. However, such a full-fledged solution is unfortunately beyond the scope of this chapter.

3.1 Three Desiderata

Here are three minimal conditions which in light of the discussion above it seems an analysis intended to solve the existence problem must meet.

²² For the former, cf. e.g. Neale (1990, 2005, 2007), Ludlow and Neale (1991). For the latter, cf. e.g. Heim (1991), Heim and Kratzer (1998), Elbourne (2005, 2010).

Condition 1	The analysis must distinguish, or minimally be capable of distinguishing, the meaning, i.e. the truth conditions, of an existential-there sentences such as (91) from the meaning of a sentence with a definite or an indefinite description in subject position such as (92). (C1)
Condition 2	The asserted content of definite and indefinite descriptions cannot be characterized using quantificational expressions or referential terms (at least when the descriptions are embedded in the scope of non-doxastic attitude verbs or in conditionals.) (C2)
Condition 3	If an analysis of indefinite descriptions (and indefinite determiners more generally) is to be overall adequate, these determiners must be analyzed as having quantificational force, i.e. as ranging over multiple individuals. (C3)

It is *prima facie* difficult to see how *any* analysis of descriptions could plausibly meet all of these conditions. And yet, the challenge is not only to devise an analysis which meets each of these conditions (even though that certainly *is* a major challenge), the analysis must also be *generally* adequate.

The most natural way of satisfying (C1) is to adopt a presuppositional analysis of not only definite descriptions but also indefinite descriptions. So, let's tentatively assume that definite descriptions *invariably* trigger existence presuppositions and that indefinite descriptions are *capable* of triggering existence presuppositions. This assumption will provide us with the required resources to distinguish between the meanings of (91) and (92), but also, potentially, the meanings of (114) and (115).²³

(114) There are some apples in the pantry and they are ripe.

(115) Some apples in the pantry are ripe.

However, if we assume that the indefinite description in (115) is capable of triggering an existence presupposition, the crucial question is how we should characterize its asserted content. The second condition above, (C2),

²³ There is ample data, for example the sentences considered with respect to (CII), suggesting that the sentences in (114) and (115) are not, contrary to popular belief, invariably equivalent in meaning. While it is not typically assumed that indefinite descriptions trigger existence presuppositions, the claim that they sometimes do has been defended by several people, e.g. Strawson (1952, 1964), Milsark (1977), Partee (1989), Diesing (1992), and von Stechow (1998), see also Heim (2010). Understanding when and why indefinites trigger presuppositions is a complex issue that I for reasons of space am unable to explore here. For now, I simply adopt this assumption with the justification that it seems intuitively correct for the cases in question. I refer the suspicious readers to the papers cited above and Appendix A where this assumption is discussed in detail.

says that the asserted contents of definite and indefinite descriptions cannot be existentially bound variables or referential terms (i.e. constants or iota-terms). We saw in the previous section that these characterizations lead to incorrect predictions. This leaves very few possibilities, but here is one: What is contributed to a sentence by a definite or a (presuppositional) indefinite description is simply an unbound variable. If we tentatively make this assumption, the challenge is to explain how this unbound variable is supposed to be assigned a semantic value.²⁴

So, let's suppose that a sentence of the form 'the F is G ' triggers the existential presupposition \mathcal{P} and asserts \mathcal{A} (cf. FIGURE I below).

FIGURE I

SENTENCE	<i>the F is G</i>	\mathcal{S}
PRESUPPOSITION	$\exists x[F(x) \wedge \forall y[F(y) \rightarrow x = y]]$	\mathcal{P}
ASSERTION	$G(x)$	\mathcal{A}

Again, presuppositions are effectively constraints on interpretability. This means that \mathcal{S} can be felicitously uttered only if its presupposition is satisfied. Hence, any context in which \mathcal{S} is felicitous — and any context in which \mathcal{S} is true — must also be a context in which the presupposition, the existential formula in \mathcal{P} , is true. Hence, if we could somehow have the seemingly unbound variable in \mathcal{A} be bound by the existential formula in \mathcal{P} , we would essentially be able to mimic the effects of wide-scoping the descriptions without making any illicit assumptions about movement of syntactic constituents (since the presuppositions would not be treated as proper syntactic constituents). This then provides at least the outline of an analysis which straightforwardly and elegantly satisfies (C1)-(C3). And if this idea could be adequately explicated, this should pave the way for a general solution the existence problem.

Luckily, there are several well-developed semantic frameworks which give us almost exactly what we need, namely so-called dynamic semantic

²⁴ The choice of unbound variables (as what is contributed by descriptions to asserted contents) is supposed to be motivated by the observation that neither quantificational expressions nor referential terms will work. However, I do acknowledge that this does not rule that there is a better alternative (even if I am not sure what that alternative would be). I thank an anonymous referee for making this clear.

systems.²⁵ In the next few subsections, I provide a short introduction to one specific dynamic framework. This framework will serve as background for the remaining discussion. I then introduce a dynamic analysis of definite and indefinite descriptions which, as I demonstrate, satisfy the three conditions described above.

3.2 *Meaning as Update Potentials*

Following Stalnaker (1970, 1974, 1998, 2002), let's assume that a discourse context is a set of possible worlds, *the context set*, where this set of worlds represents the mutually accepted presuppositions of the discourse participants, what Stalnaker calls *the common ground*: If the discourse participants mutually presuppose that ϕ , the context set will contain only ϕ -worlds, whereas if the discourse participants are agnostic or disagree about ϕ , the context set will contain both ϕ -worlds and $\neg\phi$ -worlds. Hence, if ϕ is asserted in a discourse and accepted by the discourse participants, every $\neg\phi$ -world is eliminated from the context set.

The principal difference between dynamic semantics and standard static semantics is that on the dynamic view the meaning of a sentence is not given by its truth conditions but instead by its *update potential*, viz. its potential effect on a discourse. For example, in Heim's (1982, 1983) influential implementation, sentence meaning is explicated in terms of *context change potentials* (CCPs) which informally speaking are instructions to update the current discourse context to a new revised discourse context, and formally speaking a set theoretic operation. I.e. updating a discourse context c with an atomic sentence S is formally to intersect the context set c with the set of worlds denoted by the proposition expressed by S . This set theoretic operation eliminates from the context set every world where the proposition expressed by S is not true.

The update operation for atomic sentences is however constrained by the condition that if S triggers a presupposition χ , an update of c with S is defined only if c entails χ — if the presuppositions of S are not entailed by c , the update is undefined.

For complex sentences ϕ , their CCPs are compositionally derived from the CCPs of their constituents. Furthermore, when updating a context with a

²⁵ I use the term 'dynamic semantics' as an umbrella term for various types of semantic systems which share certain non-standard features, most prominently the permissibility of binding across clause-boundaries, i.e. Discourse Representation Theory (Kamp 1981, Kamp and Reyle 1993), File Change Semantics (Heim 1982), and Dynamic Predicate Logic (Groenendijk and Stokhof 1991).

complex sentence ϕ , the update proceeds in increments, one constituent at a time, and requires that each constituent of ϕ is *admitted* by its corresponding *local* context. This means that updating with a conjunction $\phi \wedge \psi$ requires that c updated with ϕ , $c[\phi]$, is defined and that the resulting context, c' , updated with ψ is also defined, $c'[\psi]$. The local context for ϕ is thus c , whereas the local context for ψ is $c[\phi]$. In contrast, the *global* context for ψ is simply c . The figure below demonstrates how the CCPs of some additional complex sentences can be defined.

FIGURE II

ATOMIC	$c[p]$	$c \cap \{w \mid w \in p\}$
NEGATION	$c[\neg\phi]$	$c / c[\phi]$
CONJUNCTION	$c[\phi \wedge \psi]$	$c[\phi][\psi]$
IMPLICATION	$c[\phi \rightarrow \psi]$	$c / (c[\phi] / c[\phi][\psi])$

A central feature of Heim's (1982) system is that indefinites are not treated as quantificational expressions but instead assumed to introduce new so-called *discourse referents* into a discourse. Heim describes the discourse context as a collection of "file cards" that represent the introduced *discourse referents* and these file cards contain information related to the discourse referents. The context is thus updated with a sentence containing an indefinite description by adding a new file card to the context and writing the index (i.e. the number) of the indefinite on the card. The file card now represents this particular discourse referent and future occurrences of the index will be anaphorically linked to this discourse referent. Indefinite descriptions are thus formally analyzed as variables whose values depend on a model.

In contrast, if a sentence contains a definite NP (i.e. a pronoun), the information associated with the definite must be added to an already existing file card, namely the card whose index corresponds to the index of the definite. Uses of definite NPs require that the NP, in Heim's terminology, is *familiar* (that it has a file card associated with it) whereas uses of indefinite NPs require the opposite. This essentially means that definite NPs (NPs with a *+def* feature) are analyzed as anaphors which are *bound* by a previously introduced discourse referent.

3.2.1 Dynamic Semantics

In formal terms, Heim's file card metaphor is explicated using variable-assignments. So, instead of treating contexts as simple sets of worlds, we

define contexts as sets of pairs of worlds and variable assignments.²⁶ But first some basic structure.

- (A1) A model $\mathcal{M} = \langle \mathcal{W}, \mathcal{D}, \mathcal{I} \rangle$: where \mathcal{W} is a set of worlds, \mathcal{D} is a set of individuals, and \mathcal{I} is an interpretation function from basic expressions to functions from worlds to extensions.
- (A2) Let \mathbb{N} be the set of natural numbers and \mathcal{G} be a set of assignment functions g such that $g: \mathbb{N} \mapsto \mathcal{D}$.
- (A3) Let **Dom** be the set of discourse referents introduced into a discourse with $\mathbf{Dom} \subseteq \mathbb{N}$ and let $\mathbf{Dom}(g)$ denote the domain of an assignment function g .

We can now define a context c (relative to our model \mathcal{M}) as a set of ordered pairs $\langle w, g \rangle$.

- (A4) $c \subseteq \mathcal{W} \times \mathcal{G}$ such that $c = \{ \langle w, g \rangle : \forall i \in \mathbf{Dom}(g), g(i) \text{ is an individual which in } w \text{ verifies all the information shared by the interlocutors about } i. \}$

To illustrate how contexts, or information states, are updated, suppose we want to update a context c with the sentence in (116) where its logical form is assumed to be given by (116a).

- (116) A cat arrived.
 - a. $\text{cat}(x_i^{-\text{DEF}}) \wedge \text{arrived}(x_i^{+\text{DEF}})$

Indefinite NPs are assumed to introduce new discourse referents into the discourse, so updating with the first conjunct of (116a) requires that the *domain* of discourse, **Dom**, is expanded to include i . This can be captured as follows.

- (117) $\llbracket F(x_i^{-\text{DEF}}) \rrbracket = \lambda c: c \subseteq \{ \langle w, g \rangle : i \notin \mathbf{Dom}(g) \} .$
 $\{ \langle w', g' \rangle : \exists \langle w, g \rangle \in c \text{ where } g' \text{ is just like } g \text{ except that } i \in \mathbf{Dom}(g') \text{ and } g'(i) \in F(w') \text{ and } w' = w \}$

In natural language, (117) is a partial function from a context c into a new context c' which is defined only if the numerical index on the indefinite is not already included in the domain of the variable assignments. If defined, the function maps the current context c into a new context c' where the domain of discourse is extended to include i , and g maps i to an individual

²⁶ Here I follow loosely the exposition of Heim's system as it is presented in Roberts (2003: 309-310).

who is a member of the set F at all worlds of the context. This captures that indefinite NPs introduce new discourse referents, i.e. that a new file card must be opened. In formal terms, this simply means that the domain of the variable assignments in the context is extended to include the relevant index, i.e. the domain of discourse is *expanded*. Updating c with the first conjunct of (116) thus yields a set of worlds w where there is at least one cat at each w . This is thus equivalent to adding the information, or the proposition, to the context that there is a cat.

In order to interpret the second conjunct, we need an update rule for variables with the *+def* feature, i.e.

$$(118) \quad \llbracket F(x_i^{+DEF}) \rrbracket = \lambda c: c \subseteq \{ \langle w, g \rangle : i \in \mathbf{Dom}(g) \} \cdot \{ \langle w, g \rangle : g(i) \in F(w) \}$$

Since definite NPs require that their referents are *familiar*, viz. antecedently introduced into the discourse, (118) is treated as a partial function that is defined only if for every world-variable assignment pair in the context, the domain of the variable assignment includes the numerical index i . If defined, the update instruction is simply to eliminate every world/variable assignment pair from the context that fails to verify the relevant conditions on $g(i)$. I.e. every world w where $g(i)$ is not in F at w must be eliminated from c . Updating c with the second conjunct of (116) thus yields a set of worlds where the individual i , who is a member of the set of cats, is also a member of the set of individuals who arrived.²⁷

3.2.2 Presuppositional Definites

Now, in Heim's system, updating a context with a sentence of the form 'the F is G ' does not require that the referent of the F is the *unique* F , i.e. that

²⁷ As the CCP is defined in (118), it imposes the same constraint on every definite NP. But, it is clear that different definite NPs impose different further requirements on a discourse context; pronouns impose the semantic constraint that the intended referent satisfies the number, gender, and person features of the pronoun (the so-called phi-features) and definite descriptions impose the semantic constraint that their referents satisfy the descriptive content. And there are arguably more important differences between pronouns and definite descriptions, for example how salient their intended referents must be. This is an issue largely ignored in this chapter, but for illuminating discussion, see Roberts (2003). It does however raise the following question: how should these additional features be formally captured? For example, must such features be assumed to be given at logical form (similar to the *+def* feature)? These are difficult questions and to avoid an excessive and possibly exegetical discussion about logical forms, I henceforth make the simplifying assumption that occurrences of e.g. pronouns and definite descriptions can be distinguished as different inputs for semantic interpretation, and hence can be assigned different CCPs.

there are no other individuals who are F . However, I propose, for reasons to be explained, to treat ‘the F ’ as imposing the additional requirement that no other individual in the relevant context is F .²⁸

We are now in a position to define a CCP for sentences of the form ‘the F is G ’, a CCP that captures the constraints outlined in conditions (C1)-(C3) above.²⁹

$$(119) \quad \llbracket [\text{The } F]_i \text{ is } G \rrbracket = \lambda c: c \subseteq \{ \langle w, g \rangle: |F(w)| = 1 \wedge g(i) \in F(w) \} . \\ c \cap \{ \langle w, g \rangle: g(i) \in G(w) \}$$

The CCP in (119) is a partial function that is defined only if the set F has exactly one member at every world w , and g maps i to that individual at w . If the function is defined, the context is updated by eliminating every world/variable-assignment pair where the individual $g(i)$ is not a member of the set G at w .

The predictions of the CCP proposed in (119) are in standard extensional contexts more or less equivalent to the predictions of a Frege/Strawson analysis. The CCP in (119) predicts that an assertion of ‘the F is G ’ is felicitous only when it is antecedently established that there is exactly one individual who is F . This prediction mimics the predictions of the Frege/Strawson analysis where ‘the F is G ’ fails to express a proposition when the presupposition is not satisfied. When the presupposition is satisfied, the Frege/Strawson analysis predicts that ‘the F is G ’ is true only if the referent of ‘the F ’ is G . On the dynamic analysis, if the update is successful, that is if the update leaves a non-empty set of world/variable-assignment pairs, then for each world w in c , the unique individual who is F at w is also G , viz. the same truth conditions.

²⁸ Here I am essentially proposing a proper uniqueness constraint on definite descriptions (rather than Heim’s less strict familiarity-constraint). Whether definite descriptions trigger uniqueness presuppositions is a much debated issue, but nothing crucial to my main points in this chapter depends on this. For arguments in favor of a uniqueness assumption, cf. Roberts (2003). The purpose of the uniqueness constraint here is to maintain a clear difference between definite descriptions and presuppositional indefinite descriptions.

²⁹ This CCP is defined at the sentence level and I ignore issues related to its composition. This is but a convenient short cut, and not something which I believe is in any way problematic. One could, I suppose, convert (119) into a lexical entry for just the definite determiner by abstracting over both predicates and this would make it quite similar to the standard semantics for quantificational determiners (only it would output a CCP rather than a truth value). But, I remain agnostic as to what is the most optimal strategy here.

3.2.3 Presuppositional Indefinites

We have already seen how indefinite descriptions are generally analyzed in the Heimian dynamic framework sketched above, namely as expanding the domain of the variable assignments in c , viz. the domain of discourse. This would also be the analysis we would employ for sentences with existential-there sentences, i.e. ‘there is a/some F ’. However, now the question is how we should treat the special, and perhaps less frequent, occurrences of presuppositional indefinite descriptions? Here I propose a treatment similar to the treatment of definite descriptions but with one exception: presuppositional uses of indefinite descriptions do not require uniqueness.

Thus, we define the CCP for *presuppositional* uses of ‘an F is G ’ or ‘some F is G ’ as follows.³⁰

$$(120) \quad \llbracket [A^{Ps}/Some^{Ps} F]_i \text{ is } G \rrbracket = \lambda c: c \subseteq \{ \langle w, g \rangle: |F(w)| \geq 1 \wedge g(i) \in F(w) \} . \\ c \cap \{ \langle w, g \rangle: g(i) \in G(w) \}$$

As before, (120) is a partial function that is defined if only if for every $\langle w, g \rangle$ in the context c , F has *at least one* member at w and g maps i to an individual who is F at w . Again, if the function is defined, updating c means eliminating every world/variable-assignment pair where the individual $g(i)$ at world w is not a member of the set G at w . This makes presuppositional indefinite descriptions look very much like definite descriptions, but this is not surprising because they are the same in at least one sense — both trigger presuppositions. However, they are also importantly different, because uses of ‘some F is G ’ do not presuppose that there is *exactly one* individual who is F . This analysis captures this fact.

However, the most important feature of my proposed analysis is that on the presuppositional interpretation of indefinite descriptions, these do not *assert* existence (and neither, of course, do definite descriptions). A presuppositional indefinite description is treated as an anaphor and not as introducing a new discourse referent. This should now be sufficient to avoid the problematic assertions of existence without assuming that an indefinite description *refers* to a particular individual and yet retaining the quantificational force of the indefinite: Even though an indefinite description such as ‘an F ’ will be anaphorically linked to a particular discourse referent, this discourse referent can be mapped to different individuals across the worlds of the context as long as there are multiple individuals who are F .

³⁰ Again, I remain agnostic with respect to the question about *when* indefinites trigger presuppositions and also from what more general theoretical principles these presuppositional uses can be predicted.

3.2.4 Solving the Existence Problem for Conditionals

It should now be fairly straightforward to see that the analysis sketched above solves the existence problem as it arises in the cases involving conditionals. That is, using the CCP for ‘the F is G ’ given in (119) and the CCP for ‘if ϕ , ψ ’ given in section 3.2, the existence problem simply fails to arise. Consider the abbreviated version of (95) below.

(95) If [the ghost in my attic]_i is quiet tonight, then ψ .

In order to update a context c with a sentence of the form ‘if ϕ , ψ ’, $c[\phi][\psi]$ must be defined and this means that $c[\phi]$ must also be defined. In this case, since ϕ is ‘the ghost in my attic is quiet tonight’, $c[\phi]$ is defined only if the following holds:

$$c \subseteq \{\langle w, g \rangle : |\text{ghost-in-attic}(w)| = 1 \wedge g(i) \in \text{ghost-in-attic}(w)\}$$

That is, if the context fails to satisfy this constraint, the computation crashes. And if satisfied, the update then proceeds by intersecting c with the following set:

$$\{\langle w, g \rangle : g(i) \in \text{quiet-tonight}(w)\}$$

This yields a set of world/variable-assignment pairs where the individual $g(i)$ at world w is a member of the set of ghosts and the set of individuals who are quiet tonight. Next we update $c[\phi]$ with ψ and then we take the relevant complements (cf. the CCP for ‘ \rightarrow ’ in FIGURE 1). This now leaves us with a set of worlds where either the antecedent is false or the consequent is true.³¹

We also avoid the prediction that (95) is equivalent in meaning to (96).

(96) If there is a unique ghost in my attic and it is quiet tonight, then ψ .

The reason is that updating c with (96) will not require that the context satisfies the constraint described above, because this sentence triggers no existential presupposition. Once again, this highlights that the lexical entry in (119) is more or less explanatorily equivalent to the Frege/Strawson analysis of definite descriptions.

³¹ This lexical entry for the conditional is thus equivalent to material implication and I do not mean to suggest that this is an even remotely plausible semantics for indicative conditionals. However, it suffices for making the point that the lexical entry in (119) resolves the existence problem, since the mechanism employed to avoid the existence problem here could also be employed with a more sophisticated analysis of conditionals, e.g. the Kratzerian analysis of conditionals proposed in Heim (1992).

The real strength of my proposed analysis is that it works in precisely the same way with presuppositional uses of indefinite descriptions. The difference lies only in the presuppositional constraint imposed by *definite* descriptions on the ingoing context, viz. that the cardinality of the restrictor set is exactly one. We are therefore in a position to avoid the existence problem even as it arises with *indefinite* descriptions and we thus have a uniform solution to the existence problem as it arises for conditionals.

3.3 Propositional Attitudes

The next question is whether my proposed analysis suffices to make correct predictions in the problematic cases involving propositional attitude ascriptions, e.g. (106) (repeated below).

(106) Bertrand wants a murderer to be convicted tonight.

I have suggested that the indefinite description embedded inside the scope of the desire verb in (106) should be treated as triggering an existence presupposition and given my proposed analysis, this means that it should be analyzed as a variable with a *+def* feature, i.e. it should be interpreted semantically along the lines suggested in (120).

However, I have also argued in favor of the assumption that presuppositions triggered in the scope of an attitude verb project, in general, to belief contexts. So, it would not be correct to assume that a sentence such as (106) requires that a discourse referent x_i who is member of the set of murderers at w has already been introduced into the context. What is presupposed by (106) is not that there is a murderer, but that Bertrand *believes* that there is. So, we might conclude that the logical form of (106) should look something like (121) and that the context must verify (122) in order for (121) to be interpretable.

(121) $\text{DES}_b[\text{convicted-tonight}(x_i^{+\text{DEF}})]$

(122) $\text{BEL}_b[\text{murderer}(x_i^{-\text{DEF}})]$

The *intuitive* idea is that if a context verifies (122), it would then be a context where there is at least one murderer, x_i , at each of the worlds that are compatible with Bertrand's beliefs. If this context is then updated with (121), we should then predict that the update is successful only if the worlds where x_i is convicted are more desirable to Bertrand than the worlds where the set of convicted individuals is empty. Since the indefinite description is treated as a *+def*-variable, we avoid a problematic assertion of existence and since x_i is treated as a discourse referent, this discourse referent can be mapped to different individuals across Bertrand's belief worlds. And given that the

individual convicted across Bertrand's belief worlds differs, we capture that Bertrand's desire is *general*, not specific.³²

While this analysis seems both simple and intuitively correct, it faces two immediate problems which are both general and complex. These are therefore not problems that I am in a position to properly address in this chapter. However, I want to emphatically stress that there is no *prima facie* reason to believe that these problems *cannot* in principle be solved. And ultimately, I claim that (a) when these problems are solved, my suggested analysis of definite and (presuppositional) indefinite descriptions provides a full-fledged solution to the existence problem, and (b) none of the currently available analyses have any hopes of providing a *unified* solution to the existence problem, i.e. a solution that captures the conceptual similarity between definites and indefinites (even if the problems to be described below were in fact solved).³³

The first problem is the technical difficulty of extending the dynamic system above to intensional constructions such as propositional attitude ascriptions. While Heim (1992) does provide an analysis of certain attitude verbs, her analysis builds on Hintikka's (1969) semantics for attitudes which is founded on propositional modal logic. But to deal with sentences where variables are embedded in the scope of an attitude verb, we need a more expressive system. We need a system where we can characterize an agent's *existential beliefs*, i.e. beliefs such as 'there is a murderer'. And while we would normally capture such existential statements using world/variable-assignment pairs, it is not immediately clear to me how to extend this to e.g. beliefs. That is, it is not clear to me that an agent's beliefs can be sensibly described using world/variable-assignment pairs.³⁴ Solving this

³² This semantics for desire would then, in its essentials, follow Heim's (1992) proposal. However, since the aim here is only to provide an intuitive sketch of how my proposed analysis can be integrated into an already existing semantic analysis of desire, I do not mean to suggest that we *must* use the Heimian analysis of desire-verbs.

³³ To clarify, it would obviously be possible for a proponent of the Frege/Strawson analysis of definite descriptions to propose a novel analysis of e.g. indefinites where, let's suppose, these were not always treated as quantifiers. And, it is also possible that the existence problem could then be generally avoided. However, one would then not be able to capture the strong intuition that the meaning of definites, indefinites, and presuppositional uses of indefinites is closely related. The formal analyses of these determiners would likely be completely disjointed. In contrast, on my proposed analysis, these determiners share an essential component of meaning; they either introduce discourse referents or anaphorically linked to such referents.

³⁴ Interestingly, there is, it seems, a somewhat limited amount of research on the interaction of variables and propositional attitude verbs in Heim's File Change Semantics. However,

problem would require not only technical ingenuity but also a rather lengthy discussion of propositional attitudes and the more intricate details of the dynamic framework.

The second problem is, I think, more complex, but again a thoroughly general problem not related specifically to my proposed analysis. The problem is that in order for the semantic system to make correct predictions *in general*, it must be determined which constraints modals, i.e. propositional attitude verbs, impose on various anaphoric relations. Consider again (84) (repeated below).

(84) Hans wants the ghost in his attic to be quiet tonight.

For reasons already mentioned, the system should predict that a context can be updated with (84) only if the context is such that Hans *believes* that there is a ghost in his attic. But it should also predict that while (84) *does* trigger a presupposition about Hans' beliefs, (123) does not.

(123) Hans believes that there is a ghost in his attic and he wants the ghost in his attic to be quiet tonight.

The reason is that the first conjunct of (123) *asserts* what is presupposed by the second conjunct. But since I propose to analyze these descriptions as variables, this means that the system should predict that when a discourse referent is introduced under a belief verb — as 'a ghost in his attic' is in (123) — it can then function as an anaphoric anchor for a variable embedded under another attitude verb, e.g. 'the ghost in his attic' embedded under 'want' in (123). In other words, a discourse referent embedded under one attitude can *sometimes* function as a *binder* for a variable embedded under another different attitude verb. This should already be fairly obvious. However, the problem is determining *when* it can have this function. To illustrate, consider the sentences below.

(124) Hans hopes that there is [a ghost in his attic]_i and he wants [the ghost]_i to be quiet tonight.

(125) Jan expected to get [a puppy]. She intended to keep [the puppy]_i in her back yard.

(126) John wants to catch [a fish]_i. He plans to eat [the fish]_i for supper.

in the related dynamic system of Discourse Representation Theory, it seems that some steps towards an analysis of anaphoric expressions in attitude ascriptions has been taken, cf. e.g. Cumming (2007), Geurts (1998), Kamp *et al.* (2011: 326-387), see also van Rooy (2006).

- (127) Alice fears that there is [a squirrel in her kitchen cabinets]_i. She hopes to trap [the squirrel]_i alive.

As in (123), we here have discourse referents introduced under one attitude which appears to neutralize the presupposition triggered by the description embedded under another attitude. In other words, the discourse referents function as anaphoric anchors and the system should thus predict that these anaphoric relations are licensed, i.e. that there is an anaphoric relation in (124) from the discourse referent introduced under 'hope' to the variable under 'want' and in (126) from the discourse referent introduced under 'want' to the variable embedded under 'plan' etc. However, the system should also predict that anaphoric relations are not immediately licensed in (128)-(130).

- (128) # Hans wonders whether there is [a ghost in his attic]_i. [The ghost]_i is noisy.
(129) # Hans hopes there is [a ghost in his attic]_i and he believes [the ghost]_i is quite noisy.
(130) # Jan expected to get [a puppy]_i and she managed to housebreak [the puppy]_i quickly.

In order to provide a completely general solution to the existence problem, a systematic explanation of why anaphoric relations are licensed in e.g. (124)-(127) but not licensed in (128)-(130) is needed. What is *very important* to emphasize here is that this is simply a manifestation of a completely general problem already widely familiar from the study of pronouns. As an attentive reader will have noticed, one could replace the descriptions in (124)-(130) with suitable pronouns and generate the exact same puzzles.³⁵

- (124') Hans hopes there is [a ghost]_i in his attic and he wants it_i to be quiet tonight.
(125') Jan expected to get [a puppy]_i. She intended to keep it_i in her back yard.
(126') John wants to catch [a fish]_i. He plans to eat it_i for supper.
(127') Alice fears that there is [a squirrel in her kitchen cabinets]_i. She hopes to trap it_i alive.
(128') # Hans wonders whether there is [a ghost in his attic]_i. It_i is noisy.
(129') # Hans hopes there is [a ghost in his attic]_i and he believes it_i is quite noisy.

³⁵ In fact, several of the cases above are simply adapted from Roberts' (1996) paper on anaphora in intensional contexts.

(130') # Jan expected to get [a puppy]_i and she managed to housebreak it_i quickly.

This issue, which concerns the analysis of anaphora in intensional contexts, is a problem that any adequate semantic theory must ultimately address.³⁶ And, were it solved for the cases involving pronouns, the problem would then automatically be solved for definite and indefinite descriptions too (if these are analyzed as variables).

4 Conclusion

This might seem a rather disappointing conclusion of the chapter. After all, the existence problem remains unsolved — at least as it arises when descriptions are embedded in the scope of a propositional attitude verb. However, my primary aim in this chapter was to demonstrate that the existence problem is a *general* problem that neither a Russellian nor a Frege/Strawson analysis is in a position to solve. Another aim was to shed light on the debate between proponents of the Russellian analysis and proponents of the Frege/Strawson analysis. I have presented arguments against both these analyses and proposed a step towards an analysis combining elements of both Russell's analysis and the Frege/Strawson analysis. On my analysis, definite descriptions presuppose, rather than assert, existence. But it also retains in spirit the quantificational nature of Russell's analysis.

Second, I wanted to establish that a solution to the existence problem requires quite radical changes to our semantic analysis of e.g. definite and indefinite descriptions. In order to establish this, I have attempted to outline the factors that cause the existence problem and to then explicate three crucial conditions that an analysis of definite and indefinite descriptions must satisfy in order to solve the problem, viz. the conditions outlined in (C1)-(C3).

Third, I wanted to show that there is a semantic analysis which satisfies the relevant conditions, namely my proposed dynamic analysis, and that this analysis provides a uniform solution to the existence problem as it arises in conditionals, viz. a solution for both definite and indefinite descriptions.

Finally, I have attempted to show that a full-fledged solution to the existence problem relies on solutions to other complex problems. Since the existence problem arises when definite and indefinite NPs interact with

³⁶ This problem is an instance of the problem of *modal subordination* first discussed by Craige Roberts in her (1987) dissertation. An overview of this and related problems is provided in Roberts (1996).

propositional attitude verbs, this is not really surprising. But I hope to have convinced you that even though my solution is incomplete, this is not due to a problem with the proposed analysis, but rather the lack of solutions to these other general problems. I hope it is intuitively clear that were these other problems to be resolved, my proposed analysis would then provide a full-fledged solution.

I also want to emphasize that while I have relied in my exposition on a Heimian style dynamic framework, nothing really essential to my proposed analysis depends on adopting this *particular* dynamic framework. What I have argued could also be argued using e.g. Discourse Representation Theory and it is certainly possible that DRT would be better suited to (a) deal with the interaction of propositional attitude verbs and variables (i.e. pronouns) and (b) to outline the proper restrictions on anaphora in intensional contexts, see e.g. Kamp *et al.* (2011).

However, I do want to acknowledge that there are several important issues which I have not addressed in this chapter. For example, I have focused my discussion on definite and indefinite descriptions while acknowledging that the problem also arises for other determiner phrases. There is thus a question whether my proposed analysis could be generalized to these other determiners. While I see no reason to think that it could not be so extended, I recognize that this is largely an open question.

Moreover, the issue concerning what assumptions one must make about logical forms in order to properly distinguish definite descriptions, presuppositional indefinite descriptions, and pronouns from each other must be addressed. And there are additional questions concerning how to capture the apparently different salience requirements imposed by such expressions, cf. Ariel (2001) and Roberts (2003). While I am convinced that any plausible answer to these questions would complicate the theory that I have here advocated, I am also optimistic that plausible answers could be given.

In conclusion, with the analysis proposed here, there is the prospect of a full-fledged and uniform solution to the existence puzzle to emerge, and this is not the case for either the standard Russellian analysis or the Frege/Strawson analysis.

CHAPTER IV

ATTRIBUTIVE/REFERENTIAL: AGAINST THE ARGUMENT FROM REGULARITY

1 Attributive and Referential Descriptions

A very influential objection to Russell's analysis was put forward by Keith Donnellan in his seminal paper 'Reference and Definite Descriptions' (1966). Here, Donnellan argued that Russell's analysis explains only what Donnellan labeled *attributive* uses of definite descriptions but that it fails to account for so-called *referential* uses. According to Donnellan, a definite description is used attributively when a speaker uses it to predicate a property, say *G*, of an individual *x* but intends *G* to apply to *x* only in virtue of *x* satisfying the literal content of the description used. In other words, when a speaker uses a description attributively, the speaker needs not have any idea who *x* is nor be acquainted with *x* since what is predicated of *x* applies regardless of who *x* is. When 'the *F* is *G*' is used attributively, it is thus essentially equivalent to 'the *F*, *whoever that is*, is *G*'. To illustrate, consider the following example.

Context I Imagine that Smith has been murdered and that a detective discovers Smith's disfigured and mutilated body. After inspecting the body, the detective in chief concludes that whoever murdered Smith is clinically insane. He informs the other detectives by asserting (131).
(CI)

(131) Smith's murderer is insane.

Donnellan (1966)

Since the detective intends to attribute insanity to *whoever* murdered Smith, the detective's use of the description in (131) is *attributive*. The detective has no specific individual in mind, and so attributes the property of insanity to the individual who murdered Smith only *in virtue of* that individual murdering Smith.

The attributive uses contrast the so-called referential uses where the speaker entertains a singular thought, viz. an *object-dependent* thought, about a specific individual and communicates this thought using a definite description. Here the function of the description is only to help the interlocutors identify the relevant individual—the individual that the speaker has in mind. In other words, when 'the *F* is *G*' is used referentially, the speaker *refers* to a particular individual *a*, and if the intended referent is *G*, the speaker succeeds in asserting something true of *a*. To illustrate, here is a variation on one of Donnellan's famous examples.

Context II Suppose that Jones is on trial for the murder of Smith and as the trial progresses, it becomes blatantly obvious that Jones is suffering from a severe mental disorder. During his testimony, as Jones is making various clearly absurd claims, the criminal investigator, Jackson, who investigated the murder of Smith, turns to the person sitting next to him and whispers (131). (CII)

(131) Smith's murderer is insane.

Here Jackson uses the description referentially, i.e. Jackson has *referred* to Jones and asserted something *about* Jones which is true or false regardless of whether Jones in fact murdered Smith. So, even if it turned out that Jones didn't murder Smith and that Jones therefore fails to satisfy the description 'Smith's murderer', the intuition is that Jackson has referred to Jones and asserted something true *about* Jones. Here is another example.

Context III Suppose Sue is at a party and that she notices an interesting-looking man (call him 'Bob') who appears to be drinking a martini. Sue is curious about the identity of this man and therefore utters (132). (CIII)

(132) Who is the man drinking a martini?

Had it turned out that the intended referent, Bob, was in fact drinking water, it seems that Sue would nevertheless have asked a question in particular

about *Bob*. Similarly, had Sue predicated a property of the man drinking a martini using e.g. a declarative, Donnellan maintains that the truth of the declarative depends only on whether Bob has the relevant property — not on whether Bob satisfies the descriptive content of the definite description. So, if Bob is a philosopher, then Sue could, in the scenario described, have used (133) to assert something true of Bob.

(133) The man drinking a martini is a philosopher.

Donnellan observes that both Russell's (1905, 1957) analysis and the Frege/-Strawson analysis¹ of definite descriptions are incapable of explaining the referential uses. On Russell's analysis, 'the *F* is *G*' is a complex quantificational construction which asserts that there is a unique individual *x* who is *F* and that every *x* is *G*. Russell's analysis predicts that the declarative sentences, as asserted in (CII) and (CIII), are false. In other words, Russell's analysis is unable to capture the reported intuitions in (CII) and (CIII), namely that whether the assertions are true depends only on the properties of the intended referent—not on whether there exists a unique *F* who is *G*.

The Frege/Strawson analysis differs in important respects from Russell's but Donnellan argues that it also fails to explain referential uses. On the Frege/Strawson analysis, 'the *F*' is a referential term which *presupposes* the existence of a unique *F*. But when this presupposition is unsatisfied, the Frege/Strawson analysis predicts that no proposition is expressed, i.e. Strawson (1950) famously maintained that when there is no unique *F*, the speaker is making a linguistic mistake rather than asserting something false. The Frege/Strawson analysis therefore predicts that no propositions are expressed in (CII) or (CIII). Again, Donnellan argues, this does not fit with our judgments about these cases.

Nowadays, people generally agree that speakers often use definite descriptions referentially. But what conclusions we should draw from Donnellan's cases remains controversial. Donnellan is himself somewhat unclear on this matter even though his arguments strongly suggest that there are important semantic differences between attributive and referential uses. In misdescription cases, such as (CII) and (CIII), a speaker often succeeds, according to Donnellan, in asserting something *true* despite having misdescribed the intended referent. This seems to suggest that referential uses of descriptions are semantically distinct from attributive uses and thus that the definite article in English is at least two-way ambiguous.

¹ The so-called 'Frege/Strawson analysis' traces back to Frege's seminal (1892) paper. However, it was Strawson (1950, 1952, 1964) who in his responses to Russell made this analysis famous.

2 Kripke's Pragmatic Response

Responding to Donnellan, Kripke (1977) contended that to show that a certain phenomenon in natural language is a counter-example to a proposed semantic analysis, it must be shown that were the proposed analysis correct (i.e. stipulated), the problematic phenomenon would fail to arise. In less abstract terms, Kripke argues that referential uses of descriptions are inconsistent with Russell's quantificational analysis only if it can be shown that a language that worked according to Russell's analysis would not permit referential uses. In conclusion, if it's assumed that Russell's analysis of descriptions is correct, the existence of referential descriptions in English is a problem for Russell's analysis only if Russell's analysis is incompatible with referential uses of descriptions.

To show that Russell's analysis is compatible with referential uses, Kripke imagines what he labels *weak*, *intermediate*, and *strong* Russell languages—I focus only on the *strong* Russell language here. In the strong Russell language, definite descriptions are banned from the language and what one would normally express in the weak Russell language using the sentence 'the *F* is *G*', one must, in the strong Russell language, express by saying 'There is unique individual *x* who's *F* and every *x* is *G*'. Assuming that speakers of the strong Russell language are as fallible as any normal English speaker, these speakers would be equally prone to making occasional subtle mistakes. For example, suppose the context is as described in (CIII) and that Sue is a speaker of the strong Russell language. Let's again assume that Sue is entertaining a singular thought about Bob and intends to communicate that singular thought. Since definite descriptions are banned from Sue's language, she cannot assert (133) but must instead assert (134).

- (134) There is exactly one man in the corner drinking a martini and every man in the corner drinking a martini is a philosopher.

Kripke poses the following question: did Sue, despite having misdescribed the intended referent, succeed in asserting something *true* about Bob? Did Sue succeed in communicating a singular proposition about Bob? Immediate judgments suggest that Sue succeeded in both. Sue used a bona fide quantificational determiner but succeeded in expressing a thought *about* Bob and thus in communicating a singular proposition. It is therefore consistent with Russell's analysis of definite descriptions that speakers of English nevertheless use such descriptions to express singular thoughts and to convey singular propositions. Kripke concludes,

Since the phenomenon Donnellan cites would arise in all the Russell languages, if they were spoken, the fact that they do arise in English,

as actually spoken, can be no argument that English is not a Russell language. (Kripke 1977: 266)

Kripke's argument applies *mutatis mutandis* to a Strawsonian analysis of definite descriptions. For example, imagine a language where sentences of the form 'the *F* is *G*' are stipulated as having a Frege/Strawson semantics, viz. presuppose the existence of a unique *F* such that for whoever is *F*, the description 'the *F*' refers to this individual. Assuming that speakers of this language are not infallible, these speakers could surely succeed in communicating a singular thought using sentences containing definite descriptions which strictly speaking fail to refer. That is, these speakers could surely succeed in communicating singular thoughts using sentences that fail to express propositions. In other words, if the problem of referential readings is dismissed for the Russellian analysis for the reasons suggested by Kripke, it should be dismissed for the same reasons for the Frege/Strawson analysis.

2.1 *Referential Uses of Proper Names and Quantifiers*

In addition to Kripke's argument about Russell languages, Kripke points out that Donnellan's cases also appear to arise with proper names. For example, suppose that two speakers (*a* and *b*) are observing an individual from a distance whom they believe is Jones. As it happens, the individual in the distance is not Jones but Smith. Now suppose the following dialogue takes place.

- (135) *a.* What is Jones doing?
b. He's ranking the leaves.

Although the name 'Jones' refers invariably to Jones, there is a clear sense in which *a* and *b* succeed in referring to Smith and that *b* succeeds in asserting something (possibly) true about Smith. Kripke points out that this is hardly evidence in favor of the view that proper names are ambiguous, i.e. have an essentially *indexical* use on which they refer to whoever the speaker *intended*. But if not, it's not clear that the existence of referential uses of descriptions constitutes an argument in favor of there being a semantically distinct category of referentially used descriptions which have an indexical use.

A related observation is made by Neale (1990) who emphasizes that referential uses are possible with a whole range of natural language determiners.²

² Neale attributes this observation to Sainsbury (1979).

Suppose it is common knowledge that Smith is the only person taking Jones' seminar. One evening, Jones throws a party and Smith is the only person who turns up. A despondent Jones, when asked the next morning whether his party was well attended, says,

Well, everyone taking my seminar turned up.

fully intending to inform me that only Smith turned up. The possibility of such a scenario would not lead us to complicate the semantics of 'every' with an ambiguity [...] (Neale 1990: 87-88)

In order to explain the existence of referential uses and the fact that such uses have no semantic ramifications, Kripke invokes a rather subtle distinction between speaker's reference and semantic reference, i.e. what *a speaker* can refer to when using certain words and what the meaning of these words *semantically determine* as their denotation. I leave a discussion of Kripke's positive pragmatic proposal for another occasion. I turn now to a discussion of a quite recent defense of Donnellan's semantic distinction.

2.2 A Defense of Referential Descriptions

Donnellan's and Kripke's highly influential contributions prompted numerous papers on the referential/attributive distinction. Not everyone was swayed by Kripke's arguments and today there is a quite extensive literature on the topic; while some researchers remain sympathetic to Donnellan's distinction, many are highly skeptical.^{3,4} It is not possible to survey this voluminous literature here and I therefore restrict my attention to a recent argument due to Reimer (1998) and Devitt (1997, 2004, 2007) in favor of referential descriptions with semantic import. This argument, labelled 'the argument from convention' by Neale, has gained currency among many of those sympathetic to Donnellan's distinction. For example, Pupa (2008) states that the argument "undermines the pragmatic approach" (2008: 108) and "provides a very strong case for adopting" an ambiguity view (2008: 109). Similarly, Neale, one of the most vigilant defenders of the unitary Russellian analysis, proclaims that Reimer and Devitt have "presented an intuitive and powerful argument for an ambiguity in definite descriptions"

³ The proponents include, but are not restricted to, e.g. Barwise and Perry (1983), Devitt (1981b), Devitt and Sterelny (1997), Wettstein (1981, 1983), Wilson (1978), Schiffer (1998, 2005).

⁴ The opponents include e.g. Sainsbury (1979), Searle (1979), Bach (1981, 1987, 2004, 2007), Evans (1982), Neale (1990), Ludlow and Neale (1991) Salmon (1982, 1991), Soames (1986, 1994).

(2004: 173). Finally, in light of Reimer's and Devitt's argument, Abbott (2010) states that "the weight of evidence appears to be in favor of considering the referential use of definite descriptions to be semantically encoded" (2010: 152).

In the next sections, I intend to demonstrate that the intuitive appeal of Reimer's and Devitt's argument is both elusive and misleading. Rather, this is an argument based on a number of questionable empirical assumptions which makes a number of unfortunate predictions about the communication of singular thoughts. I will also show that when Donnellan's distinction is motivated in terms of Reimer's and Devitt's argument, one can no longer explain the cases that originally motivated the distinction, namely cases such as (CII) and (CIII). Moreover, as Kent Bach (2004: 227-228) has already emphasized, it seems that when one examines Reimer's and Devitt's argument in detail, it becomes clear that it severely overgenerates, i.e. it is easily extended to many cases where an ambiguity view is extremely implausible. A substantial part of this chapter is therefore devoted to (a) fleshing out Bach's claim in detail and (b) demonstrating that Reimer's and Devitt's efforts to show the contrary fail.⁵ But, first things first, let's consider Reimer's and Devitt's argument.

3 The Argument from Regularity

3.1 Regular Usage

The fact that sentences of the form 'the *F* is *G*' *can* be used to communicate singular propositions (or object-dependent propositions), is not a problem for the Russellian analysis according to Reimer and Devitt. What *is* a problem

⁵ However, this is not to be construed as a defense of the standard Russellian semantics for definite descriptions. I have already argued in Chapter II and Chapter III that the orthodox Russellian analysis (and various Neo-Russellian variants) suffer from other significant problems. The aim here is only to demonstrate that the semantic ambiguity hypothesis, as it is defended by Reimer and Devitt, is untenable. Moreover, it's often assumed that a defense of the unitary Russellian analysis of descriptions (or some other unitary analysis of descriptions, e.g. the Frege/Strawson analysis that is standard in linguistics or the analysis I have argued in favor of in Chapter III) *must* explain referential readings in terms of either *generalized* or *particularized* conversational implicatures. I am *not* attempting to bolster such a line of defense because I do not believe that such a Gricean explanation has any hope of succeeding. Towards the end of this Chapter I will instead propose an alternative explanation of referential uses based on the semantic analysis of definite descriptions proposed in Chapter III.

is the *regularity* (i.e. the *statistical frequency*) of the referential uses. This use is, Reimer and Devitt maintain, the *standard* use.

Puzzle: given that Donnellan, Chastain, and others have made *the frequency of referential uses of descriptions apparent*, why are so many philosophers so committed to the view that such uses do not exemplify a convention, or at least not a semantic convention? (Devitt 2004: 305, my italics)

[...] the fact that the referential use of definite descriptions is a *standard* (i.e. *statistically common*) use of such expressions, *does* pose a problem for Russell's Theory. (Reimer 1998: 89, italics in original)

Reimer's and Devitt's central contention is that definite descriptions are regularly used referentially, viz. to communicate singular thoughts, and that given this high frequency of referential uses, it is implausible to maintain that the referential meaning is somehow mediated by—or pragmatically derived from—the attributive meaning. Rather, if a standard use of 'the *F* is *G*' is referential, one should instead conclude that on such uses these sentences *literally express* singular propositions. That is, if the standard use of definite descriptions is referential, we should conclude that this is a semantic convention rather than attempt to explain these uses in terms of pragmatic, e.g. Gricean, inferences, as argued by e.g. Neale (1990: 83-91). This is, in its essentials, the argument.

Following Bach (2004: 227), I label this argument *the argument from regularity* (rather than convention) and throughout this chapter, I will refer to the thesis that referentially used descriptions *literally express* (viz. semantically express) singular propositions as the RDT (the referential descriptions thesis). I now discuss the putative advantages of motivating the RDT using the argument from regularity.

First, when the RDT is motivated in terms of the putative empirical observation concerning frequency of use, namely that speakers regularly use definite descriptions to convey singular thoughts, this allows one to avoid the problem raised by Neale (cf. above). The reasoning is that while it is plausible to analyze definite descriptions as genuinely ambiguous between attributive and referential meanings, this is not plausible for other determiner phrases, e.g. 'every *F*', 'both *Fs*', 'seventeen *Fs*', 'many *Fs*' etc. Even though these determiner phrases can be used — in appropriately loaded contexts — to communicate object-dependent propositions (as Neale demonstrated, cf. above), they are not standardly so used. It therefore seems implausible to think that such object-dependent uses constitute a *semantic* convention rather than a phenomenon which has a pragmatic explanation.

But Reimer and Devitt do acknowledge that when an expression E is standardly used (i.e. with a high statistical frequency) to mean E^* , this does not guarantee that E *literally* means E^* .

However, as a little reflection will show, *standard* use is no guarantee of *literal* use. In order for a (simple or complex) expression e to be capable of *literally* expressing p it must be the case that p is appropriately constrained by the *linguistic* meaning. (Reimer 1998: 95)

In other words, the literal meaning of any expression E is constrained by its *linguistic* meaning. In other words, E could literally express E^* only if E^* is consistent with the linguistic meaning of E . As an illustration, Reimer says,

If (e.g.) I say ‘She is tall’, and my intended referent does not satisfy the indexical’s linguistic meaning (i.e., is not a female), then it seems plausible to suppose that, while I may well have *communicated* a singular proposition, no proposition was *literally expressed* (on account of reference failure). (Reimer 1998: 93)

While the semantic import of ϕ -features such as *gender* is a notoriously difficult topic, cf. e.g. Heim (2008) and Sauerland (2008), Reimer’s general idea, as it applies to e.g. definite descriptions, appears to be the following. Suppose that for some expression ϕ , ϕ takes linguistic constituents $e_1...e_n$ as its arguments. The literal meaning of $\phi(e_1...e_n)$ is then (at a minimum) constrained by the linguistic meaning of the constituents $e_1...e_n$ and the product of their composition. This way, the constraint that ‘the F is G ’ (used referentially) can be true only if the intended referent is F , is built into the analysis, since F is a linguistic constituent of ‘the F ’.⁶ Let’s refer to this as the *linguistic meaning constraint* (LMC).

Now, the (LMC) helps avoid several of the problems that Donnellan’s analysis faces. First, Donnellan’s analysis of referential descriptions has the consequence that ‘the F is G ’ can be true even if no individual is F . What is required to make ‘the F is G ’ true on its referential use is simply that the intended referent is G . But assuming the (LMC), this consequence of Donnellan’s view (which is generally considered very counter-intuitive and thus problematic) is avoided. Even if ‘the F is G ’ is used referentially, it can only be true if the intended referent is F .

Another fairly clear problem with Donnellan’s analysis is that the restrictor predicate, F , appears to make no semantic contribution in ‘the F ’ when

⁶ Reimer’s and Devitt’s remarks on this issue are unfortunately somewhat vague, but Devitt states “I [have] assumed that ‘ F ’ also contributes to the meaning of ‘the F ’ pointing out the prima-facie implausibility of claiming otherwise” (2004: 291).

this description is used referentially. Given that it is possible for ‘the F is G ’ to be true even when no individual is F , the restrictor is in a sense semantically vacuous. This consequence is also avoided if the (LMC) is adopted, because it ensures that the restrictor predicate plays an important *semantic* role — it constrains what the speaker can semantically refer to using ‘the F ’ and thus what the speaker can *literally* assert. Hence,

The proposition expressed [by a referential use of ‘the F is G ’] will be of the form: x is G , where the identity of x will vary with the context of utterance. In such cases, the object or individual referred to (x) contributes directly to the proposition literally expressed, **provided it satisfies the (univocal) linguistic meaning of the definite description: provided it is the (contextually) unique F .** Thus, the contribution of the linguistic meaning in such cases is merely *indirect*. (Reimer 1998: 93) — [*my emphasis*]

Finally, the (LMC) has one additional advantage, namely that when the RDT is motivated by observations about regularity of use (that is, *frequency* of use), the argument threatens to justify the conclusion that for some non-literal meaning χ pragmatically derived from another expression $\phi(e_1...e_n)$, if the meaning of χ is *regularly* conveyed by uses of $\phi(e_1...e_n)$ then it should be plausible to assume that the literal meaning of χ is identical to the meaning of $\phi(e_1...e_n)$. But this conclusion would be too strong. However, if one, following the (LMC), assumes that the literal meaning of an expression $\phi(e_1...e_n)$ is constrained by the linguistic meaning of the constituents of ϕ , this problematic consequence can plausibly be avoided. For example,

It would also seem plausible to suppose that the linguistic meaning of every constituent of the sentence uttered must contribute in *some* way (directly or indirectly) to the proposition literally expressed. To illustrate. While sentences of the form *Could you do x ?* are standardly used to mean *Do x* , **such is not their literal meaning**, which concerns a query as to the hearer’s ability to do x . (Reimer 1998: 95) — [*my emphasis*]

While it’s not perfectly clear how Reimer intends to get from the general claim about the contribution of linguistic constituents to the specific claim about the meaning of modals like ‘could’, I assume here that this point should be explicated in terms of a constraint like the (LMC). I.e. since the linguistic meaning of the constituents ‘you’ and ‘could’ in (136a) appear to make no direct contribution to the meaning of (136b), the (LMC) is supposed to prevent us from concluding that (136a) *literally* means what is expressed by (136b).

(136) a. Could you pass the salt?

- b. Pass the salt.

To further illustrate, adopting the (LMC) provides a straightforward explanation for the intuitive implausibility of assuming that the *literal* meaning of (137a) is identical to the literal meaning of (137b) even though speakers *regularly* convey the literal meaning of (137b) using utterance tokens of (137a).

- (137) a. John kicked the bucket.
b. John died.

Moreover, in support of this claim, consider what one would normally convey by an utterance of (137c).

- (137) c. John *literally* kicked the bucket.

So, to avoid the conclusion that (137b) is the literal meaning of (137a), we should adopt the independently plausible assumption that the linguistic constituents of (137a) must make *some* semantic contribution. And if (137a) had the same *literal* meaning as (137b), the constituents ‘kick’, ‘the’, and ‘bucket’ would make no such contribution. We therefore conclude that (137b) cannot be the literal meaning of (137a).

3.2 Grasping Referential Uses

As shown above, Reimer’s and Devitt’s argument relies importantly on the putative observation that the referential use is standard. This, they maintain, is a compelling reason to assume that there is a *semantic* convention and thus that referential definite descriptions behave semantically as descriptively constrained indexicals. And when Reimer and Devitt state that it is implausible that the meaning of a referentially used description is derived pragmatically from a semantically attributive meaning, they also emphasize that interlocutors in general grasp this referential meaning, i.e. the singular thought conveyed, both quickly and effortlessly.

When a person has a thought with a particular *F* object in mind, there is a regularity of her using ‘The *F*’ to express that thought. *And there need be no special stage setting enabling her to conversationally imply what she has not literally said*, nor any sign that her audience needs to use a Gricean derivation to understand what she means. (Devitt 2004: 283 – my italics)

[Speakers] grasp the meaning [of referentially used definite descriptions] *immediately and directly* because that is the meaning it conventionally has. (Devitt 2004: 285 – my italics)

Devitt's claim here is not only that definite descriptions are *regularly* used to refer, but also that referential uses appear to incur no extra costs with respect to processing. A speaker is not required to indicate that a referentially used description should be interpreted non-literally for communication to succeed. Moreover, it doesn't seem that the interlocutor must perform a computational task as cognitively demanding as deriving a Gricean implicature in order to grasp the singular thought conveyed.

[...] in a linguistic community (such as our own) where [the referential] use was standard, *it is plausible to suppose that the intended meaning would be grasped immediately*: that is without the mediation of any Gricean-style inferences (Reimer 1998: 99 — my italics)

That the intended 'singular' reading of a sentence of the form in question *could be grasped without first grasping the 'general' reading of such an utterance*, suggests (even if it does not establish) that definite descriptions can be used literally in utterances of the form 'the *F* is *G*', to express singular propositions. (Reimer 1998: 99 — my italics)

Here Reimer emphasizes the quickness with which referential uses are grasped. The reasoning is that since the computational costs associated with the processing of a referentially used description is not obviously higher than the computational costs of processing an attributively used description, this adds further support to the thesis that the referential meaning is conventional. If the referential meaning had to be derived pragmatically from an attributive meaning, one would expect the referential meaning to be more computationally laborious.

We will have occasion to examine these claims in detail later, but let's first see how Reimer and Devitt respond to Kripke's objections to Donnellan.

3.3 *Kripke's Objections Revisited*

In response to Kripke's argument concerning Russell-English, Reimer and Devitt argue that in a community of Russell-English speakers (where in Russell-English, definite descriptions are stipulated to have a Russellian semantics, viz. be complex quantificational constructions) referential uses of definite descriptions would *not* be standard. If the referential uses *were* standard, then it should be possible to analyze these descriptions as literally expressing singular propositions — but by stipulation, it is not.

The phenomena differ because there is no convention in Russell-English of using definites to express singular thoughts. (Devitt 2004: 287)

The problem is, in essence, that Kripke's argument fails to factor in the diachronic aspect of language evolution; a once non-literal meaning χ ,

pragmatically derived from an expression ϕ , can (over time) become the *literal* meaning of ϕ .⁷

Reimer and Devitt also reject the standard Russellian explanation that analyzing the referential use in terms of the attributive use and relevant pragmatic principles obviates the need for stipulating ambiguity. Against this explanation, Reimer offers the following consideration: the verb ‘incense’ once literal meant ‘to make fragrant with incense’ and was used only metaphorically to mean ‘to make angry’. But the once metaphorical meaning has now been conventionalized and so what used to be a metaphor is now dead. In other words, the verb ‘incense’ is now ambiguous in that it has two truth conditionally distinct *literal* meanings.

The problem is that using the Kripke-style reasoning, we could show that the verb ‘incense’ is *not* ambiguous. Just imagine a different language, English’, where ‘incense’ is stipulated to be unambiguous, i.e. its literal meaning is simply ‘to make fragrant with incense’. Could speakers of English’ use this verb in the metaphorical sense, i.e. to implicate ‘to make angry’? Almost surely. But this means that the metaphorical uses could occur in English’ even when in that language the verb is stipulated to have only *one* literal meaning. As a result, the existence of such metaphorical uses is not an argument against an analysis which analyzes ‘incense’ as having only one literal meaning. And if we adopt the Gricean dictum that “senses are not to be multiplied beyond necessity” (1989: 47) and consequently wield Grice’s modified Occam’s Razor, we can then conclude that the English’ analysis of ‘incense’ is superior since it avoids positing a lexical ambiguity.

This argument, Reimer and Devitt emphasize, must be unsound, because if it *were* sound, it could be used to show that there are no dead metaphors in any language. The problem, Reimer argues, is that Kripke’s argument disregards the important relation between *standard* use and *literal* meaning. If meaning supervenes on use, then if an expression ϕ is *standardly* used to express ψ then it is likely that the ψ -meaning becomes lexicalized so that ϕ (at some point) comes to *literally* mean ψ .

⁷ A problem that Kripke (1977) himself mentions in the very last paragraph of his paper: “In particular, I find it plausible that a diachronic account of the evolution of language is likely to suggest that what was originally a mere speaker’s reference may, if it becomes habitual in a community, evolve into a semantic reference. And this consideration may be one of the factors needed to clear up some puzzles in the theory of reference.” (Kripke 1977: 271)

3.4 Analyzing *The Argument from Regularity*

The argument from regularity relies importantly on both empirical assumptions and various theoretical considerations, so in order to assess its strengths and weaknesses, it will be useful to have a clear outline of its structure. Here is thus a structured outline of the argument and its premises.

THE ARGUMENT FROM REGULARITY (RA)

Premise 1	Sentences of the form ‘the F is G ’ are <i>standardly</i> (or <i>regularly</i> , <i>frequently</i>) used to communicate singular propositions (whereas sentences such as ‘every F is G ’, ‘most F s are G s’, ‘both F s are G s’ etc. are not). (P1)
Premise 2	The meaning of ‘the F is G ’ is grasped quickly and effortlessly. (P2)
Premise 3	The literal meaning of an expression ϕ (taking arguments $e_1 \dots e_n$) is constrained by the linguistic meaning of $e_1 \dots e_n$: ϕ can <i>literally</i> mean ψ only if the linguistic meaning of each constituent of ϕ , viz. $e_1 \dots e_n$, makes <i>some</i> semantic contribution to the meaning of ψ . (P3)
Premise 4	When a is the unique F , it is consistent with the linguistic meaning of ‘the F is G ’ to assume that its literal meaning is ‘ a is G ’. (P4)
Premise 5	For any expression ϕ , if ϕ is standardly (regularly, frequently) used to communicate ψ and ψ is grasped quickly and effortlessly on the basis of uses of ϕ , then if ψ is consistent with the linguistic meaning of ϕ , it’s plausible to assume that the literal meaning of ϕ is ψ . [For example, assume ϕ = ‘the F ’ and ψ = ‘ a ’ where a is the unique individual who is F .] (P5)
Since ‘the F is G ’ is standardly (regularly, frequently) used to communicate a singular proposition, namely the proposition that a is G (when a is the unique F), and this singular meaning is grasped quickly and effortlessly, then since it’s consistent with the meaning of ‘the F is G ’ that it literally means ‘ a is G ’...	
Conclusion	... it’s plausible to assume that the literal meaning of ‘the F is G ’ is ‘ a is G ’. (∴)

To summarize the motivation for the above premises: (P1) and (P2) are supposed to be empirical facts, although I should mention that even if these premises might seem intuitively plausible, Reimer and Devitt provide no empirical evidence in their support.

(P3) is motivated by the independently plausible assumption that linguistic constituents always make some semantically significant contribution.

Thus, if an expression ϕ is composed of the constituents $e_1...e_n$, then the literal meaning of ϕ is constrained by the linguistic meaning of $e_1...e_n$.

(P4) is the observation that when it is assumed that the literal meaning of 'the F is G ' is ' a is G ', this is consistent with premise (P3) as long as a is the unique F . I.e. in that case, the nominal F is playing a semantic role, namely that of constraining what 'the F ' can refer to and this is consistent with it referring to a because a is the unique F .

(P5) is supposed to be independently motivated, i.e. it's roughly a combination of the Wittgensteinian slogan that meaning supervenes on use and the restriction that constituent linguistic meanings make semantic contributions and thereby constrain *literal* meaning.⁸ Reimer says,

Surely the fact that an expression is *standardly* used to mean such-and-such suggests (even if it does not establish) that it can be used – literally – to mean such-and-such. (Reimer 1998: 98)

In sum, referential uses of 'the F is G ' literally express singular propositions not merely because these uses *can* facilitate communication of singular propositions, but rather because they *consistently do* facilitate such communication and this communication appears to proceed quickly and effortlessly.

This concludes my exposition of Reimer's and Devitt's argument from regularity. I now turn to a detailed discussion the argument and I argue that it faces several significant problems.

4 Evaluating the Argument from Convention

There are, I think, several serious worries with the (RA), but here I focus on four particular issues. First, I show that Reimer's and Devitt's view runs into trouble when it comes to explaining certain intuitions about the use of referential descriptions and the conditions for communicative success. Second, I demonstrate that, contrary to Reimer's and Devitt's repeated claims, there is no reason to think that the (RA) *cannot* be extended to other determiners of English, and hence that the argument generalizes in a way that both Reimer and Devitt quite explicitly intend to avoid. Third, I explicate in detail an objection to Reimer's and Devitt's view, which to my knowledge was first stated by Bach (2004: 227-228), that the (RA) proves too much, i.e. that their argumentative strategy can be applied to numerous other cases for which an ambiguity view is at best implausible and highly controversial.

⁸ cf. Wittgenstein (1953: 25e).

4.1 *Communicative Success and Referential Descriptions*

One of Donnellan's most important insights was that when, for example, the object of a singular thought is perceptually accessible to a speaker and audience, and the speaker wishes to communicate a thought about that object, what matters for successful communication is only that the speaker and audience come to entertain a thought about the *same* individual. Donnellan demonstrated that this is possible even in cases where the vehicle with which the speaker expresses the singular thought, say a definite description, is strictly speaking inadequate, i.e. not satisfied by the intended referent. That is, if a speaker intends to communicate a singular thought, it effectively makes no difference whether the chosen vehicle of language is adequate as long as the speaker can reasonably assume that using that vehicle of language will enable her audience to entertain the same singular thought.⁹ Thus, as regards successful communication, it is often irrelevant from both the speaker's and the audience's point of view whether the intended referent satisfies the description.

This observation led Donnellan to conclude that definite descriptions have referential uses where their *semantic* referent is simply the intended referent, but this conclusion also leads to the quite counter-intuitive consequence that 'the *F* is *G*' can be true even when there are no individuals who are *F*. In order to avoid this counter-intuitive consequence, Reimer and Devitt maintain that definite descriptions are semantically referential only when the intended referent satisfies the descriptive content. So, when a speaker entertains a singular thought, intends and succeeds in communicating that singular thought, but misdescribes the intended referent, she fails to *refer* and hence she fails to literally express a singular proposition.

Now, it is important to recognize that the assumption that 'the *F* is *G*' can be true only if the intended referent is *F* in fact undermines the motivation that Donnellan originally put forward in favor of his distinction, namely misdescription cases. When this assumption is adopted (which perhaps it should be), misdescription cases can no longer be explained in terms of Donnellan's distinction, and instead we must resort to a pragmatic explanation. Reimer and Devitt are thus committed to the view that the explanation of why communication succeeds in normal cases is importantly different from the explanation of why communication succeeds in misdescription cases; the first is explained simply in terms of the proposition *literally* expressed

⁹ Although, if speakers are obeying Grice's (1989) cooperative principle, it makes sense for them to choose the vehicle which is the least cognitively taxing for the interlocutors, since one might think that this improves the chances of communicative success.

whereas the latter can only be explained by appeal to pragmatics. But since the phenomenology of these cases can be quite similar (which essentially motivated Donnellan's view), one might be inclined to believe that there should be a general and uniform explanation in both types of cases. And so, if one wants to avoid the consequence that 'the *F* is *G*' can be true when no individual is *F*, it would seem quite reasonable to accept Kripke's general pragmatic explanation.

A related worry is that Reimer's and Devitt's view fails to explain why — when interlocutors fail to grasp what a speaker has *literally* said — communication nevertheless appears to proceed smoothly. Now, whether an utterance of 'the *F* is *G*' succeeds in communicating a singular proposition depends on various factors that a speaker has no control over. Suppose I'm entertaining the singular thought about my sister that she just completed a marathon and suppose further that I want to communicate this thought to my interlocutors. In this context, it would be natural for me to assert (138).

(138) My sister just completed a marathon.

However if my interlocutors have never met my sister and in fact have no idea who my sister is, I can only succeed in communicating a *general* proposition.¹⁰ My assertion will enable my interlocutors to entertain only the *general* thought that my sister, *whoever she is*, just completed a marathon. Consequently, if my interlocutors are not acquainted with my sister (in any relevant sense), it is simply not possible for me to successfully communicate the *singular* thought that *she* just completed a marathon — even if I so desired. The communication of the singular thought thus depends on both the mental states of the speaker *and* the interlocutors.

Given Reimer's and Devitt's view, it seems plausible that there are situations where a speaker *literally* expresses a singular proposition but where the interlocutors are in a position to grasp only a *general* proposition. For example, given the scenario described with respect to (138), I would effectively have asserted the singular proposition about my sister that she just completed a marathon, but my interlocutors might only be capable of recovering a general proposition. In that case, Reimer's and Devitt's view entails that my interlocutors fail to grasp what I have *literally* said. And even though this might seem *prima facie* unproblematic, since this could also happen with e.g. proper names (when the names are not antecedently introduced) and deictic uses of pronouns (where the required demonstration

¹⁰ I.e. suppose that there is no relevant perceptual link between my sister and my interlocutors, as Devitt (2004: 290) argues is required for singular thought.

fails), the effect is manifestly clear — communication breaks down. For example, if I assert (139) but fail to demonstrate anything, i.e. suppose I just point into the void, my interlocutors would be forced to bring the conversation to a halt and ask me to elaborate or explain myself.

- (139) HE_f just completed a marathon.
[pointing into the void]

In contrast, there is no breakdown in communication when I assert (138) and my interlocutors are not acquainted with my sister. So even though my interlocutors apparently failed to grasp what I *literally* said, communication appears to be completely unaffected. The possibility of systematic miscommunication without any effects on the discourse seems like something that needs to be explained.

4.2 Generalizing the (RA) to other Determiners

The objection to Donnellan's view that a speaker can use almost any determiner phrase of English to communicate singular propositions (object-dependent propositions) and that the view leads to the implausible conclusion that all determiners of English are lexically ambiguous, is a familiar one.¹¹ Reimer's and Devitt's view is supposed to avoid this problem because the RDT is motivated in terms of regularity of use.¹² But this motivation is based on a bona fide empirical claim about what is standard usage and what is not. The question then is whether this empirical claim holds up under closer scrutiny.

Let's grant that a sentence such as 'the *F* is *G*' is standardly (regularly, frequently) used to convey object-dependent propositions and that in the general case, there are no misdescriptions. Consider now standard uses of *plural* definite descriptions; if speakers can use singular definite descriptions to express object-dependent propositions, namely to *literally* express singular propositions, speakers can presumably also use plural definite descriptions for the same purpose.¹³ But to be sure, consider the example below.

¹¹ Cf. the quote from Neale above, but see also Kripke (1977), Sainsbury (1979), Neale (1990), Ludlow and Neale (1991)

¹² See in particular Reimer (1998: 96) and Devitt (2004: 283)

¹³ I.e. it seems to me that if the definite article is analyzed as ambiguous between referential and attributive readings, this should apply regardless of the morphology of the NP-argument.

Context IV Suppose Sue is at a party and that she notices two interesting-looking men standing in the corner (call them ‘Bob’ and ‘Bill’). Suppose further that Bob and Bill both appear to be drinking martinis. Now imagine that Sue either asks the question in (140) or asserts the declarative in (141). (CIV)

(140) Who are the men in the corner drinking martinis?

(141) The men in the corner drinking martinis are quite handsome.

That Sue is entertaining an object-dependent thought, namely a thought whose content depends on Bob and Bill, seems a reasonable assumption. It also seems reasonable to assume that (141) expresses an object-dependent proposition, namely a proposition *about* Bob and Bill. And given that the mark of the referential use is Sue’s intention to refer, viz. Sue’s intention to express a proposition about Bob and Bill, we should, following Reimer and Devitt, assume that Sue has effectively asserted a proposition which *literally* expresses an object-dependent proposition. If so, Sue has effectively used the description *referentially*.¹⁴

Plural determiner phrases do not pose any *prima facie* obstacles with respect to referentiality. What is required for succeeding in referring is simply that the speaker has a particular number of individuals in mind and that she intends to assert something *about* those individuals. And again, for the communication of an object-dependent proposition to succeed, the interlocutor must also have the relevant individuals in mind. However, now the question is whether there is any convincing reason to maintain that *quantificational* determiner phrases are not *regularly* so used.

Context V Suppose Sue is at a party and that she notices four interesting-looking men standing in the corner (call them Bob, Bill, Jack, and John). Suppose further that Bob, Bill, Jack, and John are wearing Armani suits. Finally, suppose that Sue has antecedent reasons to believe that most men at this party are married. Now imagine that Sue either asks the question in (142) or that Mary, when noticing Sue’s interest in Bob, Bill, Jack, and John, asserts the declarative in (143). (CV)

¹⁴ The fact that the proposition expressed by Sue depends on *two* rather than one object should make no difference to whether the expression is genuinely referential. For example, if Sue had pointed at the two men in the corner in (CIV) and asserted “they are quite handsome”, the plural pronoun would clearly *refer* to Bob and Bill. Sue must therefore be interpreted as asserting something *about* Bob and Bill.

- (142) Is every man in the corner wearing an Armani suit married?
 (143) Every man in the corner wearing an Armani will stay with his wife until he dies.

Is there any reason to think that Sue in asking the question in (142) is not asking a question in particular *about* Bob, Bill, Jack and John? For example, just imagine that Sue wants to ask out either Bob, Bill, Jack, or John (because of, say, their beautiful brown eyes) and as a result she wants to know whether Bob, Bill, Jack, and John are married.¹⁵ If so, Sue clearly isn't intending to ask a question about any individual satisfying the description 'in the corner wearing an Armani suit'. Rather, Sue is intending to ask a question about Bob, Bill, Jack, and John. Hence, it seems to me that insofar as Bob, Bill, Jack and John are the only individuals in the relevant context who satisfy the description 'in the corner wearing an Armani suit', Sue clearly succeeds in communicating an object-dependent proposition about Bob, Bill, Jack, and John. And a similar argument could easily be devised for Mary's utterance of (143).

Now we might ask what the relevant difference between (CV) and a case involving a *plural* definite description is supposed to be? Reimer and Devitt must maintain that the difference is that the determiner phrases in (CV) are not *standardly* used to refer but that plural definite descriptions are.¹⁶ Now, even if we were inclined to accept the claim that singular and plural definite descriptions are *regularly* used to communicate object-dependent propositions at face value (i.e. without empirical justification), why should we accept the claim that a quantificational determiner such as 'every *F*' is not? After all, it seems that in (CV), the definite determiner and a quantificational determiner like 'every' can do precisely the same work. For example, consider the following questions.

- (144) a. Do you mean every man over there drinking martinis?
 b. Do you mean the men over there drinking martinis?
 c. Do you mean all of the men over there drinking martinis?

¹⁵ One could make a case that the definites in (140-141) are more natural here. But 'every' is the natural choice in (CV) if for example the speaker wants signal her presumption that most men are married. However, even so, the speaker might very well be interpreted as intending to communicate an object-dependent proposition, a proposition about the individuals in question.

¹⁶ And for reasons that are unclear to me, Bach (2004: 226) concedes that "this is an important difference" between definite descriptions and other determiner phrases in English despite the fact that the goal of Bach's paper is to rebut Devitt's arguments for referential descriptions.

If we assume that Bob, Bill, Jack, and John are drinking martinis, then we can easily imagine that Sue would answer each of these questions in the affirmative. If so, Sue intended to express a thought *about* Bob, Bill, Jack, and John because if Sue had merely intended to talk about *whoever* had the property of being ‘in the corner wearing an Armani suit’ she would not answer ‘yes’ to the above questions. It looks like ‘every *F*’, ‘all of the *Fs*’, and ‘the *Fs*’ serve the exact same purpose; the speaker entertains an object-dependent thought about some individuals, intends to communicate that thought and, if the relevant individuals are perceptually accessible, she chooses a vehicle of language which will enable her interlocutors to come to entertain the same object-dependent thought. Plural definite descriptions and determiner phrases such as ‘every *F*’ will often be equally well suited for this purpose.

Given the lack of empirical evidence to the contrary, I do not think there are any clearly compelling reasons to accept the claim that the use of the determiner phrase in (CV) is either nonstandard or even infrequent. Nor do I think that there are any immediate reasons to believe that the object-dependent proposition is somehow cognitively more taxing to compute than the general proposition, i.e. one would be hard pressed to argue, with no evidence from psycholinguistic research, that grasping the referential use of the determiner phrase in e.g. (143) is slower, or requires more effort, than grasping the referential use of the plural definite description in (141).

In conclusion, without at least *some* empirical evidence, we should reject (P1). And whether there is any empirical support for this premise is questionable. Yet if we accept that various other determiner phrases of English are frequently used to communicate object-dependent propositions, the (RA) now becomes potentially an argument for an across-the-board semantic ambiguity — not just for definite descriptions. Reimer and Devitt could choose to bite the bullet, but since it was emphasized as a virtue of the (RA) that it avoided this consequence, I assume that even they find this problematic. But, there are even more disconcerting applications of the (RA).

4.3 *Linguistic Meaning vs. Literal Meaning*

The role of the constraint in (P3) is ensure that the restrictor, *F*, in ‘the *F* is *G*’ makes a semantic contribution to the sentence in which it occurs. According to Reimer, this contribution is a constraint on possible semantic referents.

DET (D)	RESTRICTOR (NP)	NUCLEAR SCOPE (VP)
<i>The</i>	<i>F</i>	<i>G</i>
The	man in the corner wearing an Armani suit	is drinking a dry martini.

The nature of the constraint is explicated in terms of linguistic meaning, i.e. ‘the *F* is *G*’ can be true only when the intended referent satisfies the property determined by the linguistic meaning of the restrictor. Given the linguistic meaning of the restrictor’s constituents, and their semantic composition, the sentence above can semantically refer only to individuals who satisfy the derived property, namely individuals who are male, in the corner, and wearing an Armani suit. The semantic referent cannot be e.g. a woman, a man in the middle of the room, or a man in the corner wearing a Hugo Boss suit. Each lexical constituent of *F* has a linguistic meaning which contributes compositionally to determine the property which is linguistically expressed. The sentence which embeds this description can therefore be true only if the intended referent has that property.

Now, as explained earlier, this constraint on literal meaning (referred to earlier as the LMC) is also supposed to help prevent the (RA) from overgenerating. The problem is that it doesn’t. I will now describe two distinct cases where it seems quite clear that the (RA) overgenerates even when the independently plausible constraint on literal meaning, (P3), is accepted.

4.3.1 Overgeneration I: Scalar Implicatures

It’s commonly agreed that a speaker who asserts a sentence such as (145) often *conversationally implicates* what is literally expressed by (146). These implicatures are called quantity implicatures, or *scalar* implicatures, owing to the fact that they are usually explained in terms of Grice’s (1989) maxim of quantity and so-called *Horn scales* (Horn 1972, 1989).

(145) Some students passed the exam.

(146) Some students, but not all, passed the exam.

In other words, it is commonly agreed that interlocutors quite generally infer from utterances of sentences such as (145) that the speaker meant (146). It is also assumed that (146) is a *pragmatically derived*, and hence *non-literal*, meaning of (145). But using an argument in the style of the (RA), we can now demonstrate that this standard assumption is incorrect and that it is more plausible to assume that (146) is not a pragmatically derived implicature but rather the *literal* meaning of (145). Let’s consider such an argument.

THE SCALAR ARGUMENT FROM REGULARITY (SRA)

Premise 1	Sentences of the form ‘some <i>Fs</i> are <i>Gs</i> ’ are <i>standardly</i> (regularly, frequently) used to communicate that some <i>Fs</i> , but not all, are <i>Gs</i> . (P1)
Premise 2	The scalar meaning of ‘some <i>Fs</i> are <i>Gs</i> ’, namely ‘some <i>Fs</i> , but not all, are <i>Gs</i> ’, is grasped quickly and effortlessly. (P2)
Premise 3	The literal meaning of an expression ϕ (taking arguments $e_1...e_n$) is constrained by the linguistic meaning of $e_1...e_n$: ϕ can <i>literally</i> mean ψ only if the linguistic meaning of each constituent of ϕ , viz. $e_1...e_n$, makes <i>some</i> semantic contribution to the meaning of ψ . (P3)
Premise 4	The assumption that the <i>literal</i> meaning of ‘some <i>Fs</i> are <i>Gs</i> ’ is ‘some <i>Fs</i> , but not all, are <i>Gs</i> ’ is consistent with the <i>linguistic</i> meaning of ‘some <i>Fs</i> are <i>Gs</i> ’. (P4)
Premise 5	For any expression ϕ , if ϕ is standardly (regularly, frequently) used to communicate ψ and ψ is grasped quickly and effortlessly on the basis of uses of ϕ , then if ψ is consistent with the linguistic meaning of ϕ , it’s plausible to assume that the literal meaning of ϕ is ψ . [For example, assume ϕ = ‘some <i>Fs</i> ’ and ψ = ‘some <i>Fs</i> , but not all’.] (P5)
<p>Since sentences of the form ‘some <i>Fs</i> are <i>Gs</i>’ are <i>standardly</i> used to convey that some <i>Fs</i>, but not all, are <i>Gs</i> and this is grasped quickly and effortlessly, then since it’s consistent with the <i>linguistic</i> meaning of ‘some <i>Fs</i> are <i>Gs</i>’ that it <i>literally</i> means ‘some <i>Fs</i>, but not all, are <i>Gs</i>’...</p>	
Conclusion	...it’s plausible to assume that the literal meaning of ‘some <i>Fs</i> are <i>Gs</i> ’ is ‘some <i>Fs</i> , but not all, are <i>Gs</i> ’. (∴)

I assume that the conclusion of the (SRA) is untenable and hence that it should be resisted. But if the (SRA) is somehow unsound, the crucial question is whether this argument can be shown to be unsound without thereby also showing that the (RA) is unsound.¹⁷ Since (P3) and (P5) in the

¹⁷ I should mention that there is an on-going debate about the scalar meanings of e.g. ‘some’. For example, Chierchia (2004), Spector (2007), Chierchia *et al.* (2010) argue that sentences such as ‘some *Fs* are *Gs*’ indeed are ambiguous between the ‘some, but not all’ and the ‘some, and possibly all’-readings (but not because of a lexical ambiguity in the word ‘some’). However, the arguments provided in favor of this view rely on entirely different considerations than those suggested by the above argument and I’m fairly sure that these researchers would not want to adopt an ambiguity view on the basis of the (SRA). Moreover, what’s more important for my purposes is that I strongly suspect that

(SRA) are unchanged from the (RA), there are only three options: Either (P1), (P2), or (P4) in the (SRA) must be false. Let's consider each of these options in turn.

Now, if one was inclined, with regards to the (RA), to accept (P1) without any empirical justification, viz. to accept that referential uses of definite descriptions are *standard*, there is no obvious reason to be less inclined to accept that the scalar use of 'some *F*s' is also standard. I.e. it seems equally plausible that the frequency of the uses of 'some *F*s' on which the speaker intends to convey 'some *F*s but not all' is higher than the frequency of the uses where the speaker only intends to convey 'some *F*s and possibly all'. In other words, if one was to conclude that the above argument fails because the first premise is false, one would need to explain why one shouldn't conclude that the (RA) fails for the exact same reason. One would effectively need to provide empirical evidence that the scalar uses of 'some *F*s' are statistically *infrequent* whereas the referential uses of 'the *F*' are statistically frequent but it is *prima facie* difficult to believe that the data would confirm such a hypothesis.

Next, one might question whether scalar implicatures are grasped as quickly and effortlessly as the literal 'some *F*s, and possibly all'-meaning of (145). If there were reasons to believe that it takes more cognitive effort to derive the meaning in (146) than understanding the literal meaning of (145), one might conclude that (146) could only be pragmatically implicated rather than *literally* asserted by (145). This would be grounds for rejecting (P2). Several psycholinguistic experiments on this question have been conducted. Some of these experiments suggest that scalar implicatures take longer to process and thus are computationally more laborious, but other experiments appear to show the opposite.¹⁸ More than anything, these experiments demonstrate that claims about putative computational processing costs are incredibly hard to assess and even more difficult to empirically justify. The evidential weight of such claims (when no empirical justification is provided) is therefore highly questionable. Moreover, it is not clear that processing speed — i.e. how quickly and effortlessly certain meanings are grasped — reliably indicates whether a standardly conveyed meaning is the *literal* meaning. For example, many psycholinguistic studies have shown that idiomatic meaning is grasped more rapidly, or at least as rapid, as literal meanings, see e.g. Swinney and Cutley (1979), and Gibbs (1985). Also,

neither Reimer nor Devitt are inclined to accept the conclusion of the (SRA) and thus that they would want to resist it.

¹⁸ For the former cf. Bott and Noveck (2004). For the latter, cf. Grodner *et al.* (2010).

Gibbs (1986) has shown that when speakers are confronted with an utterance of (136), they often don't compute the literal meaning of this sentence and instead immediately interpret it as conveying the meaning in (136b).

- (136) a. Could you pass the salt?
b. Pass the salt.

In other words, it is far from clear that if some content standardly conveyed is not the *literal* content, this is clearly manifested by the speed or ease with which that content is grasped. And, let's not forget, that even if processing speeds were relevant, Reimer and Devitt supply no evidence beyond raw intuitive judgments in support of these claims. So, if one concluded that the above argument fails because scalar implicatures are not "grasped immediately" and thus that the second premise is false, one would be hard pressed to nevertheless maintain that the second premise of the the (RA) is true (or even that it is more plausible that it is true).

Summing up, it is not feasible to reject either (P1) or (P2) of the (SRA) without thereby also rejecting (P1) or (P2) of the (RA). The remaining option is to reject (P4), namely the premise that it is consistent with the linguistic meaning of 'some *F*s are *G*s' that it literally means 'some *F*s, but not all, are *G*s'. Now, remember that according to (P3), it is consistent to assume that an expression ψ is the literal meaning of an expression ϕ (taking arguments $e_1...e_n$) only if the linguistic meaning of the constituents of ϕ make some semantic contribution to the literal meaning, viz. to ψ . So, for example, the literal meaning of 'the *F*' is constrained by the linguistic meaning of its restrictor, *F*. And given the linguistic meaning of the restrictor, *F*, 'the *F*' can only semantically refer to individuals who satisfy this constraint. Accepting this prima facie plausible assumption, it seems reasonable to assume that the literal meaning of 'some *F*s' must be constrained in a similar way. The question is whether accepting this constraint then rules out that 'some *F*s' could literally mean 'some *F*s, but not all'. The answer is clearly 'no'. Let me illustrate.

FIGURE II

DET (D)	RESTRICTOR (NP)	NUCLEAR SCOPE (VP)
<i>Some</i>	<i>F</i> s	<i>G</i> s
Some	students	passed the exam

As stated in (P3), we assume that the linguistic meaning of any ingoing constituent must contribute to the truth conditions of the sentence and

thereby constrain what could be literally expressed. With regards to (145) (repeated below), this means, more concretely, that the nominal ‘students’ must be assumed to constrain what ‘some students’ could literally mean.

(145) Some students passed the exam.

In particular, we must conclude that the relevant domain of quantification for the indefinite determiner is the domain of students (and nothing else). And so, we are licensed in concluding that (145) could not *literally* mean what is expressed by e.g. (147)-(149).

(147) Some professors passed the exam.

(148) Some mathematicians, but not all, passed the exam.

(149) Some students and philosophers passed the exam.

In other words, the nominal restrictor constrains the domain of individuals that the sentence could be interpreted as *literally* being about. However, the restrictor only constrains what domain of quantification the indefinite determiner ranges over, nothing more and nothing less. And hence, such a constraint is satisfied by sentences such as (150) and (151).

(150) Some students in the third grade passed the exam.

(151) Some students, who had flunked every previous exam, passed the exam.

Now we can of course straightforwardly rule out that (145) literally means what is expressed by either (150) or (151). We can rule this out simply by observing that the meanings of (150) and (151) are not regularly conveyed by utterances of (145).¹⁹ But, since (150) and (151) are consistent with the constraint on literal meaning described in the third premise of the (RA) and the (SRA), we must also accept that (146) is consistent with this constraint. And as regards (146), we cannot rule out on the basis of the previously mentioned independent grounds that it is not literally expressed by (145), because (145) *is* regularly and standardly used to convey what is expressed by (146).

(146) Some students, but not all, passed the exam

¹⁹ Of course, things are not quite this simple. It’s well known that natural language determiners are subject to implicit domain restrictions and depending on how such domain restrictions are formally cashed out, one might think that in certain contexts (150) and (151) are *literally* expressed by utterances of (145). I’ll generally ignore this issue here since it’s orthogonal to my main point. For discussions of the vexing issue of quantifier domain restrictions, cf. von Stechow (1994) or Stanley and Szabo (2000).

In conclusion, it seems that we are forced to either (a) accept the conclusion of the (SRA) or (b) reject (P1), (P2), or (P4). The former option saddles us with an implausible view of scalar meanings, while the latter renders the conclusion of the (RA) false. The most plausible conclusion seems to be that the (RA) is somehow fallacious because even if one was inclined towards a semantic ambiguity view of scalar terms such as ‘some’, one would hardly think that an argument such as the (SRA) could in fact settle this issue.

4.3.2 Overgeneration II: Ability Modals

Let’s now consider Reimer’s own example, the apparent difference in meaning between (152) and (153), and the claim that the (RA) does not commit one to the view that (153) is *literally* expressed by (152).

- (152) Can you pass the salt?
 (153) Pass the salt.

First some preliminary ground clearing: (152) is an interrogative sentence consisting of two major components: a modal auxiliary (*can*) and its pre-jacent (*you pass the salt*). Let’s assume uncontroversially that on the most natural and *literal* interpretation of (152), the modal auxiliary is interpreted as an ability modal (a *bouletic* modal) and that (152) can be more or less accurately paraphrased by (152’).

- (152’) Are you capable of passing the salt?

FIGURE III

MODAL (AUX)	PREJACENT (S)
<i>modal</i>	ϕ
can	you pass the salt

Now, in contrast to (152), the sentence in (153) is an imperative (a non-interrogative clause-type) which is ordinarily used to make commands or requests. One could thus paraphrase more or less accurately the meaning of (153) (when used as a command) using a declarative sentence such as (153’).

- (153’) I command you to pass me the salt.

What is important for now is that there is an intuitive difference between what we take the *literal* meanings of (152) and (153) to be. Hence, the question is this: if one accepts the (RA), can one nevertheless maintain that

such an argument cannot be also be used to show that (152) *literally* means (153)?

THE MODAL ARGUMENT FROM REGULARITY (MRA)

Premise 1	Questions of the form ‘can you F ?’ and ‘could you F ?’ are <i>standardly</i> used as imperatives, viz. to issue commands or requests that the addressee F s. (P1)
Premise 2	The imperative meaning of ‘can you F ?’ and ‘could you F ?’ is grasped quickly and effortlessly. (P2)
Premise 3	The literal meaning of an expression ϕ (taking arguments $e_1 \dots e_n$) is constrained by the linguistic meaning of $e_1 \dots e_n$: ϕ can <i>literally</i> mean ψ only if the linguistic meaning of each constituent of ϕ , viz. $e_1 \dots e_n$, makes <i>some</i> semantic contribution to the meaning of ψ . (P3)
Premise 4	Assuming that the <i>literal</i> meaning of ‘can you F ?’ and ‘could you F ?’ is ‘ $F!$ ’ (where ‘!’ signals the imperative mood) is consistent with the linguistic meaning of ‘can you F ?’ and ‘could you F ?’ (P4)
Premise 5	For any expression ϕ , if ϕ is standardly (regularly, frequently) used to communicate ψ and ψ is grasped quickly and effortlessly on the basis of uses of ϕ , then if ψ is consistent with the linguistic meaning of ϕ , it’s plausible to assume that the literal meaning of ϕ is ψ . [Assume ϕ = ‘can you F ?’ and ψ = ‘ $F!$ ’.] (P5)
	Since questions of the form ‘can you F ?’ and ‘could you F ?’ are <i>standardly</i> used as imperatives to issue commands or requests and this meaning is grasped quickly and effortlessly, then since it’s consistent with the <i>linguistic</i> meaning of ‘can you F ?’ and ‘could you F ?’ that it <i>literally</i> means ‘ $F!$ ’...
Conclusion	...it’s plausible to assume that the <i>literal</i> meaning of ‘can you F ?’ and ‘could you F ?’ is ‘ $F!$ ’ (∴)

I assume that the conclusion of the (MRA) is false, or at least not adequately justified on the grounds of the (MRA) alone: It is not plausible to assume that a sentence such as (152) *literally* means the same as (153).²⁰ But if this argument is unsound, then this again means that either (P1), (P2), or (P4) must be false — since (P3) and (P5) are unchanged from the (RA).

²⁰ I assume that Reimer agrees, because — as we have seen — she states herself that her argument for referential descriptions does not commit her to this equivalence in *literal* meaning.

First, it seems almost impossible to seriously challenge (P1); surely it is true that (152) is standardly used as an imperative, i.e. to express a command or a request in the form of (153). If one was to reject this premise, it would be completely ad hoc to nevertheless maintain that (P1) in the (RA) is true.

As regards (P2), I have already mentioned the psycholinguistic studies which show that subjects compute the imperative meaning in (153) from an utterance of (152) even faster than they compute the literal meaning of (152). In fact, the literal meaning is often not even computed at all, cf. Gibbs (1986). As a result, it seems impossible to challenge (P2) in the (MRA) while maintaining that (P2) of the (RA) is true.

This leaves only (P4) to be questioned and judging from Reimer's albeit vague remarks, this premise is supposed to be false. The (P3) constraint is supposed to rule out the possibility of running an (RA)-type argument for sentences such as (152). That is, just as we should think of *F* as imposing a constraint on the meaning of 'the *F*', we should presumably also think of the prejacent as imposing a constraint on the meaning of modal+prejacent constructions, i.e. that the linguistic meaning of ϕ as it occurs in e.g. *can*(ϕ) imposes a constraint on the literal meaning of *can*(ϕ). And this then is supposed to avoid the conclusion that the literal meaning of (152) is (153).

The problem is that it doesn't. For each constituent of the prejacent of (152), it is easy to demonstrate that its linguistic meaning *does* constrain what e.g. command or request that sentence can *literally* express, but that these linguistic meanings nevertheless are consistent with the meaning of (153). We can demonstrate this by varying the meaning of each constituent and observing that the result is a request or command which (152) could *not* be used to express (given the way these expressions are currently used). So, for example, (152) could *not* be used to express the commands or requests in (154)-(157).

(152) Can you pass the salt?

- | | | |
|-------|--|----------------------------------|
| (154) | Pass the pepper. | [varying the meaning of 'salt'] |
| (155) | Pass some of the salt. | [varying the meaning of 'the'] |
| (156) | Throw the salt | [varying the meaning of 'pass']. |
| (157) | [addressing someone other than the addressee]:
Pass the salt. | [varying the meaning of 'you'] |

Each constituent of the prejacent, namely *you*, *pass*, *the*, and *salt*, clearly constrains what request (152) could be (and is) conventionally used to make. So, even though the word 'you' doesn't actually occur in (153), its linguistic meaning could nevertheless be considered as playing an important role in giving the meaning of (153), i.e. it makes a significant semantic contribution.

The same considerations apply in an obvious way to the constituents ‘pass’, ‘the’, and ‘pepper’. But what about the conventional meaning of the modal auxiliary ‘can’? This word does not occur in (153) either, so what semantic contribution is this word then making (since every word should play some semantic role)? And since the linguistic meaning of the modal concerns an *ability* (as Reimer phrased it), this aspect of the meaning of (152) is surely incompatible with the literal meaning of (153)? That’s correct, but we must bear in mind that the (MRA) is an argument about the *literal* meaning of certain *modals* precisely as the (RA) is an argument about the meaning of the definite determiner. Since the meaning of the definite determiner is under discussion, the constraint on the linguistic meaning of ‘the *F*’ in the (RA) comes from the restrictor, viz. *F*. The exact same could be said to apply in the case of the (MRA): the meaning of the modal auxiliary is under discussion, and so the linguistic constraint on what (152) can literally mean comes from its prejacent. In less abstract terms, Reimer and Devitt maintain that the semantic contribution of ‘the *F*’ on its referential use is an individual, but whether it semantically refers to the intended referent depends on whether the individual satisfies the property expressed by the restrictor — this is the linguistic constraint. The same considerations appear to apply here: the contribution of the modal auxiliary ‘can’ on its *imperative* use, or *request/command*-use, is simply an imperative meaning constrained by the linguistic meaning of the constituents in the prejacent. If the reasoning underlying the (RA) is sound, I see no obvious reason to think that this reasoning about the modal ‘can’ is somehow not.

It seems that Reimer and Devitt must accept that if the (RA)-type argument is a genuine argument for the (RDT), it is also a straightforward argument for the (less fortunate) conclusion that ‘can’ is ambiguous between a standard bouletic modal meaning and a quasi-imperative meaning. And hence that (152) and (153) can have the same *literal* content.

Summarizing the objections outlined above, we find the following:

- Reimer’s and Devitt’s view fails to shed a *general* light on successful communication of singular thoughts, since every misdescription case requires a pragmatic explanation. The cases that motivated Donnellan are not explained by invoking this distinction.
- Reimer and Devitt are also committed to the counter-intuitive conclusion that when interlocutors fail to grasp the singular thought that a speaker intended to convey using a sentence of the form ‘the *F* is *G*’, they have literally *failed to grasp* what the speaker said.

- Contrary to the claims of the proponents of the (RA), it can easily be generalized to other determiners of English and therefore it cannot plausibly be restricted to the definite determiner alone.
- The main argument for the view could just as well be used to show that scalar implicatures are in fact not implicatures and that sentences such as (152) *literally* expresses the order in (153). These are merely two detailed cases that demonstrate that the (RA) overgenerates, i.e. that an argument of this type is easily extended to various other, standardly assumed pragmatic, phenomena in natural language. However, as Bach (2004: 227-228) also observes, the (RA) could easily be extended to cover numerous other cases, e.g. attitude verbs such as 'believe' and 'suspect', and temporal and causal interpretations of natural language conjunction 'and' to mention just a few.

4.4 Explanatory Value

I want to conclude my discussion of Reimer's and Devitt's case for referential descriptions by raising a question: When the semantic distinction between attributive and referential descriptions is motivated as proposed by Reimer and Devitt, what are the explanatory advantages? The answer to this question is, I think, somewhat unclear. Now, on Reimer's and Devitt's view, the meaning of a referentially used description differs importantly across contexts from the meaning of an attributively used description. I.e. if 'the *F* is *G*' is used referentially, it expresses different propositions in different utterance contexts, whereas if it is used attributively, it expresses the same proposition in every utterance context. The thought is that this difference in meaning is captured when it is assumed that referentially used descriptions *literally* express singular propositions. In contrast, if definite descriptions are assumed to uniformly express only general propositions, this difference in meaning is not captured.

But even when we acknowledge this difference in meaning, the remaining question is whether it is somehow advantageous to posit an ambiguity in order to capture it. After all, in those cases where there appears to be a genuine ambiguity between these different meanings, it's not clear that we *need* an ambiguity analysis in order to explain it. As Heim (2010) points out.

But what difference in linguistic behavior and judgment, if any, corresponds to this technical difference [between indexical/referential uses and attributive uses]? If we just use truth-value judgment tasks, we cannot distinguish the two cases. In making a truth value judgment, a speaker contemplates whether a given sentence would be true if the world in which it was uttered had such and such properties. The

same imagined world here serves both as utterance world and world of evaluation, and therefore the difference between indexical [referential] and widest-scope non-indexical [attributive] meanings is systematically neutralized. (Heim 2010: 32)

Heim's point is that the difference in meaning between referential and attributive uses of descriptions (as it is proposed by Reimer and Devitt) has a detectable impact on our judgments only when these descriptions occur in intensional environments. But as Heim also points out, there are independent reasons for believing that various formal tools are required to deal with intensional constructions and that these tools will enable us to capture the relevant differences in meaning. In conclusion, it's just unclear whether there are any explanatory advantages to Reimer's and Devitt's view.

I tentatively conclude that as the (RA) stands, it fails to offer a convincing reason to favor the (RDT) and that when the (RDT) is motivated as Reimer and Devitt suggest, it remains unclear what explanatory advantages the view has.

5 Gödelian Completions

Let's suppose that an adequate analysis of the semantics and pragmatics of definite descriptions should be capable of explaining the following facts.

- Definite descriptions are used to convey both *general* and *singular* thoughts.
- The communication of singular thoughts do not always depend on the intended referent being adequately described, i.e. speakers will misdescribe an object/individual and nevertheless often succeed in communicating a singular thought about that individual/object.
- When the intended referent is misdescribed, we have an intuition that the speaker "did something right but also that [the speaker] did something wrong. After all, the description [the speaker] used failed to fit the person he wanted to 'talk about,' and to that extent the speech act was defective." (Neale 1990: 91).

If a uniform analysis of definite descriptions is to capture these facts, it must somehow be able to capture both the general and the singular aspect of the meaning of definite descriptions. One quite promising attempt towards that aim is a proposal put forward by Neale (2004).

Now, Neale's proposal uses a mechanism which it is often assumed is necessary for dealing with incomplete definite descriptions, i.e. descriptions

where the descriptive content is inadequate for determining a unique denotation. To account for these cases, it is argued that the content of these incomplete descriptions is enriched by context (or perhaps more precisely, speaker intentions). For example, when a speaker asserts the sentence ‘the table is covered with books’ in some context *c*, it is assumed that the speaker might in fact have asserted a descriptively more rich proposition, e.g. that the table in the corner of this room is covered with books. The content of the description is thus enriched in context and it is precisely facts about the context that are supposed to explain how interlocutors are capable of determining precisely which proposition was in fact expressed.²¹

How exactly these content enrichments, or supplementations, are to be explained, and whether the process in question is semantic or pragmatic, is a widely debated and controversial issue.²² But, assuming that such content supplementations are needed for independent reasons, Neale (2004, 2005) proposes to use these supplementations to explain the referential uses. The suggestion is to analyze the descriptive content of referentially used descriptions as supplemented with an identity statement where one relata of the identity relation is a singular term. Neale refers to these supplementations of content as Gödelian completions.²³ Hence, when the definite description in (133) (repeated below) is used referentially, its descriptive content is not (158) but rather (159).²⁴

(133) The man drinking a martini is a philosopher.

(158) [The *x*: drinking-a-martini(*x*)](philosopher(*x*))

(159) [The *x*: drinking-a-martini(*x*) \wedge *x* = *a*](philosopher(*x*))

It is assumed here that the constant *a* is a referential term referring to the individual that the speaker intended and if the intended referent is indeed drinking a martini, the speaker succeeded in literally expressing a true proposition *about* that particular individual. Conversely, if the intended referent is not drinking a martini, the speaker said something obviously false,

²¹ This is oversimplifying the issue a great deal, but since ultimately I won’t be relying on this mechanism, it is for illustrative purposes only.

²² This issue is discussed in several of the papers cited in the introduction, but see e.g. Devitt (1981a), Wettstein (1981), Salmon (1982), Recanati (1986), Soames (1986), and Neale (1990).

²³ cf. Neale (2001: 118-128, 2004: 171-175, 2005: 849).

²⁴ Since Neale uses a restricted quantifier notation (cf. the introduction), I adopt this practice in describing his view. The truth conditions are here assumed to be the standard Russellian truth conditions.

namely that the intended referent is identical to the unique individual (in the context) drinking a martini. Neale thus concludes.

On such an account, ‘the’ used referentially would (perhaps as a matter of implicit convention) amount to a Gödelian description $(\iota x)(\phi x \wedge x=a)$. If this is a plausible story, then the debate between the Russellian and the ambiguity theorist grinds to a compromise: when ‘the’ is used referentially, ‘ $C(\text{the } \phi)$ ’ is used to express a proposition that is both Russellian and object-dependent in the relevant way. (Neale 2005: 849)

While Neale’s proposal is deceptively simple, it explains why definite descriptions can be used to convey both singular and general thoughts. It also straightforwardly explains (a) why speakers can succeed in conveying singular thoughts even when the descriptions fail to denote and (b) why it seems that the speaker has done something right in such cases but also something wrong. Moreover, since it does not explain referential uses in terms of Gricean conversational implicatures, it is exempt from having to provide an additional explanation of how exactly these referential readings are supposed to be pragmatically derived. In sum, this seems like a very promising proposal.

However, in response to Neale’s proposal, Devitt poses the following question.

But is the Gödelian completion really Russellian? Note that it yields a proposition that is *equivalent* to a conjunction of singular propositions:

Fa.Ga

So it is hard to see how the Gödelian completion yields something that is in any *interesting* sense a general proposition and hence Russellian. We might say that it is syntactically Russellian but not semantically so; it is only “pseudo-Russellian”. (Devitt 2007: 29)

One might think that Devitt makes a valid point here which should concern diehard Russellians. But even if Gödelian completions do appear to be a significant departure from the standard Russellian paradigm that ‘the F is G ’ invariantly expresses a general proposition, there is also a sense in which this proposal maintains the core Russellian features, namely that definite descriptions are quantificational expressions akin to other determiners of English *and* that the definite article in English is not ambiguous.

In the remaining section, I argue that by utilizing the resources made available by a dynamic semantic framework, a proposal along the lines of Neale’s proposal can be maintained without making any assumptions about supplementation with identity statements. I also hope to show that when the proposal is simply a consequence of the semantic system in question, Devitt’s point is less of a worry.

5.1 Gödelian Completions and Dynamic Semantics

On standard referential uses of definite descriptions, both speaker and interlocutor have a particular individual *in mind*, and in the paradigm cases the intended referent is perceptually accessible to both. In such cases, if the speaker intends to refer to *a* and *a* is the unique *F*, then if the speaker utters the sentence ‘the *F* is *G*’, the interlocutor comes to entertain the singular thought that *a* is *G* as a result of the speaker’s utterance.

However, suppose the speaker or interlocutor are not (perceptually) acquainted with *a*. In such cases, a singular thought is not communicated; if the speaker is not acquainted with *a*, then the speaker could not be entertaining a singular thought about *a* and hence not be in a position to convey such a thought. In contrast, if the interlocutor is not acquainted with *a*, then the interlocutor is incapable of recovering a singular thought about *a*.

But now suppose that the speaker misdescribes the intended referent, *a*. If so, the interlocutor cannot come to entertain a singular thought about *a*, *unless* the interlocutor either (incorrectly) assumes that the intended referent, *a*, is the unique *F*, or instead simply assumes (perhaps for the sake of the conversation) that the speaker (incorrectly) believes that the intended referent, *a*, satisfies the descriptive content.

What is important to notice here is that these types of scenarios can be explained solely in terms of backgrounded presuppositions. For example, it seems reasonable to assume that when ‘the *F* is *G*’ is intended to express a thought about a particular individual, say *a*, the speaker is, in standard cases, *presupposing* that *a* is the unique *F*. I.e. suppose that the speaker is entertaining the singular thought that *a* is *G*. He then chooses to convey that thought using a sentence of the form ‘the *F* is *G*’ because he assumes that *a* is the unique *F* and he assumes that by describing *a* as ‘the *F*’ he is helping his interlocutor identify *a*. So here it seems that the description functions merely as a descriptive guide to the intended referent.

Now, if the interlocutor is *also* presupposing that *a* is the unique *F* (or at least willing to accommodate that presupposition, i.e. by in context turning his head and determining *who* the unique *F* is), then when confronted with the speaker’s utterance, the interlocutor will come to entertain the very thought that the speaker intended to convey, namely that *a* is *G*. This seems both a simple and intuitive explanation.

Let’s remind ourselves of the analysis of definite descriptions that was proposed in Chapter III.

$$(119) \quad \llbracket [\text{The } F]_i \text{ is } G \rrbracket = \lambda c: c \subseteq \{ \langle w, g \rangle: |F(w)| = 1 \wedge g(i) \in F(w) \} . \\ c \cap \{ \langle w, g \rangle: g(i) \in G(w) \}$$

Again, the context change potential of ‘the F is G ’ is formally a partial function that is defined only if the restrictor set F has exactly one member at every world w and g maps i to that individual at w . If the function is defined, the context is updated by eliminating every world/variable-assignment pair where the individual denoted by $g(i)$ is not a member of the set G at w .

So, suppose it is common ground that some particular individual a is the unique F . In terms of the dynamic semantic system, this means that the common ground is such that ‘the F ’ is mapped to the same individual at every world in the context.²⁵ Hence, when it is common ground that a is the unique F , an assertion of ‘the F is G ’ will invariably enable the interlocutors to entertain the singular thought about a that a is G . There is therefore a sense in which the speaker *has* asserted a singular proposition, because the denotation of the definite description does not vary across the worlds of the context and the effects of Gödelian completions (viz. enriching contents with identity statements) are thereby achieved simply in virtue of the dynamic system and relevant assumptions about the discourse participants’ beliefs. Yet as explained above, these assumptions are independently plausible, because in order for the speaker to communicate a singular thought, both speaker and interlocutor must be acquainted with a and believe that a is the unique F . In conclusion, for the standard cases, i.e. paradigm cases, of referential uses of descriptions, the dynamic semantic framework has plenty of resources to explain these uses.

But how should the ubiquity of referential uses be explained? Here is a plausible assumption. Speakers typically talk about individuals and objects in their immediate environments, viz. individuals and objects with which they are acquainted (or which they could become acquainted given their immediate physical environment). Given this, it will be common ground in most contexts that a is the unique F , b is the unique G , c is the unique H etc. That is, the majority of definite descriptions are going to be mapped to the same individual across every world of the context. It then follows that in the majority of cases, singular thoughts (in the sense described above) are unproblematically conveyed by sentences of the form ‘the F is G ’. In other words, the fact that ‘the F is G ’ is standardly used to communicate singular thoughts is neither surprising nor problematic on this analysis.

But what about cases where both speaker and hearer are not acquainted with the individual denoted by ‘the F ’? In these cases, if the speaker asserts ‘the F is G ’, the speaker could not be communicating a singular thought, and

²⁵ I.e. a is now to be thought of as a member of the domain; an individual which a variable assignment could map a natural number onto — and not as a constant of the language.

this is exactly what the dynamic analysis proposed here predicts; an assertion of ‘the F is G ’ requires that it is common ground that there is a unique F , and if this is common ground, yet neither speaker nor interlocutor have any idea who the unique F is, the semantic value of ‘the F ’ varies across the worlds of the context set. In other words, ‘the F ’ will be mapped to different individuals across the worlds of the context set and so the speaker will not be communicating a singular thought because the content communicated will not be about any *particular* individual.

The most complicated cases are those involving misdescriptions, but even here the dynamic analysis provides a fairly good explanatory strategy. Suppose that the speaker misdescribes the intended referent. What happens in such cases? Well, if the interlocutor is under the same misapprehension as the speaker, i.e. that the intended referent is the unique F , the information update proceeds without any problems and the speaker succeeds in communicating a singular proposition. Why? Because if both speaker and interlocutor (incorrectly) presuppose that b is the unique F , then at every world in the context set, the semantic content of ‘the F ’ is mapped to b . Hence the speaker would be successfully communicating something about b despite the fact that b is not, in actual fact, the unique F .

But what about cases where the speaker incorrectly believes (and hence presupposes) that, say, b is the unique F but the interlocutor does not? Well, in such cases, we are confronted with what Stalnaker (2002) calls a defective context — a context where the presuppositions of the discourse participants are not identical.

A nondefective context is a context in which the participants’ beliefs about the common ground are all correct. Equivalently, a nondefective context is one in which each of the parties to the conversation presuppose the same things. A defective context may go unnoticed; that is, it can happen that all participants believe that the context is nondefective even when it is defective. But it can also happen that one or more parties to the conversation are aware that the context is defective, and in this case some kind of corrective action is called for. (Stalnaker 2002: 717)

Let’s consider one context where corrective actions is called for, namely the example from CIII.

Context III Suppose Sue is at a party and that she notices an interesting-looking man (call him Bob) who appears to be drinking a martini. Sue is curious about the identity of this man and therefore utters (133).
(CIII)

(133) The man drinking a martini is handsome.

Here Sue has misdescribed Bob and Sue's interlocutor is aware of this (suppose Sue's interlocutor knows that Bob is teetotaler). If so, the interlocutor might nevertheless decide to disregard that Sue misdescribed Bob and for the sake of the conversation presuppose that Bob indeed is drinking a martini. Stalnaker says,

The case of the man drinking an alleged martini illustrated the kind of defective context that is most easily remediable: a case where one party not only recognized that the context was defective, but also recognized exactly how the presuppositions of the two parties diverged, and so knew how to repair or accommodate. (Stalnaker 2002: 718)

Thus the explanation of the communicative success here, viz. the fact that a singular proposition about Bob appears to have been communicated, is explained just as it was in the case where both speaker and interlocutor falsely believes that the intended referent is the unique *F*.

6 Concluding Remarks

The dynamic analysis and the associated framework thus appears to provide a very simple and intuitive explanation of the facts outlined in the beginning of Section 5. This explanation requires no assumptions about content completions as it follows directly from the way the semantic system is defined. Moreover, there is no sense in which it is "pseudo-Russellian" because it is not Russellian at all.

However, this is not to deny that there are more complicated cases, for example cases where the interlocutor refuses, or is simply unable to, accommodate the relevant presupposition. It is evident that explaining such cases requires a definition of context that permits diverging presuppositions among the discourse participants. How precisely Stalnaker's system is to be revised in order to achieve this is a question that I am at present unable to answer. However, I do not believe that there is any principled reason to maintain that this would be impossible.

Finally, since my proposed analysis has the consequence that when it is common ground that, say, *a* is the unique *F*, sentences of the form 'the *F* is *G*' then expresses a singular proposition, one might object that even when it is common ground that *a* is the unique *F*, one can nevertheless use 'the *F*' attributively. I.e. by asserting 'the *F* is *G*' one can communicate that the property of being *G* applies to *whoever* is *F* (incl. the actual *F*). One might think that an example of this would be an assertion of e.g. (160).

(160) The president is wise.

Asserting (160) would then be intended to express that *whoever* is the president is wise, even though, let's suppose, it is common ground that the president is Barack Obama.

My initial reaction to this objection is simply that it is unclear whether this is in fact possible. Setting generic uses aside as special case, e.g. reference to kinds, I would argue that when it *is* common ground that *a* is the unique *F*, then if one intends to use 'the *F*' in the strictly attributive sense (as above), one would normally resort to modals or adverbial quantifiers to make one's point. I.e. (161) or (162).

(161) The president is always wise.

(162) The president should be wise.

But since modals and adverbial quantifiers are intensional operators, other tools, possibly quite complex, are required for analyzing these cases. For a start, we must permit that when a definite description is embedded under a possibility modal, its semantic value is not necessarily determined by the worlds of the context, but rather by some alternative set of worlds where the denotation is not Barack Obama at every world. With these tools in hand, we should then be perfectly capable of accounting for both the "referential" reading of the president in (161) and the attributive reading.

My second response would be that if the notion of context is redefined so that a non-defective context can contain multiple discourse participants with inconsistent presuppositions, we should also be able to model temporary revisions of the context. I.e. one can imagine that if (160) can be used attributively in a context where 'the president' is known to denote Barack Obama, the context can be temporarily reset so that the value of 'the *F*' varies across worlds.

These are admittedly fairly vague responses and ultimately they rely essentially on various technical innovations. But again, even though I do not at present know how to extend the Stalnakerian framework in the way suggested, I have no reason to think that it is in principle impossible. And it is difficult to deny that the explanation of referential and attributive uses presented here, which essentially falls out of the dynamic analysis proposed in the previous chapter, seems a quite promising start.

APPENDIX A

INDEFINITES AND PRESUPPOSITIONS

The solution to the existence puzzle that was proposed in Chapter III relied on one important assumption, namely that the weak determiners in examples (106) and (107) should be interpreted as triggering existential presuppositions. However, in order to maintain the focus of that chapter, I decided not to discuss this assumption in detail (and instead just pointed to other places in the literature where this assumption is defended). But, since this is not an uncontroversial assumption, I suspect it might engender some skepticism about my proposed solution. In this appendix, I therefore intend to demonstrate that the assumption that weak determiners *can* and *do* trigger existence presuppositions has considerable evidence in its favor. The purpose of this appendix is thus largely to consider the available data, but in the last section I turn briefly to a discussion of one potential objection.

1 Preliminary Observations

Milsark (1977) observed that there is a syntactic two-way division of natural language determiners witnessed by the distributional effects of so-called existential-there sentences (or there-insertion sentences). In particular, insertions of a determiner in the postcopular NP position of 'there BE'-constructions yield grammatical sentences for only a restricted set of determiners.

(163) STRONG DETERMINERS

- a. * There is every book on the table.
- b. * There is the book on the table.
- c. * There are all three books on the table.

- d. * There are both books on the table.

(164) WEAK DETERMINERS

- a. There is some/a/one book on the table.
- b. There are many books on the table.
- c. There are three books on the table.
- d. There are no books on the table.

Why natural language exhibits this syntactic distribution is a matter of controversy and I introduce the distinction here mainly to single out a set of relevant determiners, namely determiners that are grammatical in there-insertions, viz. weak determiners.

The distributional effects demonstrated above are sometimes explained by positing a difference in semantic properties; *strong* determiners are presupposition triggers whereas *weak* determiners are not. Hence, weak determiners are treated as *asserting* existence, i.e. the asserted semantic content of a singular weak determiner phrase is some existentially quantified formula ranging over individuals while the asserted semantic content of a plural determiner phrase, for example numerical determiners, is either an existentially quantified formula ranging over groups (with some specified cardinality) or a range of successive existentially quantified formulas ranging over individuals.¹ The reasoning is that strong determiners, which presuppose existence, are ungrammatical in existential-there sentences because the explicit purpose of an existential-there sentence is to establish existence. But this purpose of existential-there sentences would be superfluous if it was combined with a strong determiner which already presupposes existence. Hence, these are incompatible.

However, Milsark also observed that weak determiners appear to have presuppositional uses too. In particular, there seems to be a contrast in the meaning of weak determiners when used in existential-there sentences and when used as simple sentence subjects. For example, in (165) below there is no immediate temptation to infer that the existence of apples in the pantry is taken for granted, but this is not clearly the case for (166) and (167). The latter sentences elicit a much stronger intuition that the speaker is taking this

¹ I'm ignoring various alternative approaches to the semantics of indefinites, most notably choice function approaches, see e.g. Winter (2001). Plural determiners in general introduce various complexities, but to the extent possible I simply ignore these. I assume throughout that a presuppositional treatment of strong determiners is already independently justified and thus focus on supplying arguments for the claim that weak determiners can trigger existence presuppositions. For arguments in favor a presuppositional treatment of strong determiners, see e.g. De Jong and Verkuyl (1985).

information for granted. This contrast is subtle, but one way to detect it is to consider which sentence would be more natural to assert if the existence of apples in the pantry was not already established in the discourse.

(165) There are some ripe apples in the pantry.

(166) Some apples in the pantry are ripe.

(167) Several apples in the pantry are not ripe.

However, as Reinhart (1995/2006), von Fintel (1998), Heim (2010) point out, even if it is acknowledged that (166) and (167) seem to give rise to presuppositional inferences, and (165) does not, this subtle contrast fails to conclusively establish that the weak determiners in (166) and (167) trigger existence presuppositions. The primary reason is that the availability of the putative presuppositional inferences is consistent with the assumption that these determiners both assert *and* presuppose existence. Since existential interpretations are always true in contexts where existence is presupposed, this makes it particularly difficult to conclusively establish that a weak determiner is a genuine presupposition trigger. Hence, as von Fintel points out,

The most important methodological point to remember is that it will not be enough to show that a particular sentence (with a supposedly presuppositional indefinite) can be uttered in a context where an existence presupposition is satisfied. What has to be shown is that such a sentence can *only* be uttered in such a context. A supposedly presuppositional indefinite needs to be shown to *require* a context satisfying its presupposition. (von Fintel 1998: 3)

In other words, the contrast between (165) and (166-167) is not sufficient for ruling out an analysis where the presuppositional inference is explained in purely pragmatic terms. In order to conclusively establish that some weak determiner is a genuine presupposition trigger, stronger evidence is needed, i.e. evidence of the kind suggested by von Fintel.

1.1 Entailments, Implications and Presuppositions: Projection

Logical entailments, implicatures, and presuppositions license inferences, but what distinguishes a presuppositional inference from a logical entailment and an implicature is its *projection behavior*.²

² This is today a thoroughly standard assumption, see e.g. Heim (1990), Beaver (1997, 2001) Geurts (1999). For more critical assessments of projection tests as tests for presuppositions, see Simons *et al.* (2009), Simons *et al.* (2010), Simons (2010).

An inference is said to *project* when it survives certain complex embeddings. For example, suppose some atomic sentence S licenses the inference I and that S is embedded under an *entailment-cancelling* operator, e.g. a negation or a possibility modal. Now, if $\neg S$ or $\diamond S$ continues to license I , I is said to *project*. Consequently, if we are aiming to determine whether some expression E is a genuine presupposition trigger, we must examine whether the inferences, that are licensed by sentences containing E , survive when they are embedded under the relevant operators. And so, to justify the assumption that the weak determiners in (166) and (167) trigger existence presuppositions, it should be shown that the inferences licensed by these sentences project.

Now relying on the clear exposition of data in von Stechow (1998), I now demonstrate that when indefinites are subjected to the right tests, we find that the inferences triggered by these determiners behave just as we would expect if these were presuppositional inferences.³

2 Projection Tests

2.1 Questions

It is well known that if a declarative sentence ϕ licenses a presuppositional inference, then if ϕ is converted into a polar question, viz. a yes/no question, a presuppositional inference licensed by ϕ survives this conversion. This behavior of presuppositional inferences contrast logical entailments and conversational implicatures since these tend to disappear when the declarative sentence which triggered them is converted into a polar question.⁴

³ The data in sections A1.2 and A1.3 is adapted from von Stechow's (1998) paper. I have added the data in sections A1.4 and A1.5. Now, this appendix is only intended to establish that presuppositions can be triggered by weak determiners. It is not an attempt to explain or derive these empirical observations from more general theoretical principles. To my knowledge, there is no real consensus on this question, but some of the explanations that have been suggested in the literature include: (a) that indefinites are lexically ambiguous between presuppositional and non-presuppositional interpretations, cf. Diesing (1992), (b) that the presuppositional effects arise when weak determiners are headed by individual-level predicates (and so should be explained in terms of the individual-level/stage-level distinction, cf. von Stechow (1998), Heim (2010), and (c) that these presuppositional effects come from treating the weak determiners phrases as topical, cf. Büring (1996). While I remain neutral on this question, I do provide some evidence against (b).

⁴ In the following examples, I indicate when inferences are intuitively licensed using ' \rightsquigarrow ' and ' \rightsquigarrow^* ', respectively.

(168) ASPECTUAL VERBS

- a. Bertrand stopped smoking.
- ~> Bertrand used to smoke. *(presuppositional inference)*
- ~> Someone stopped smoking. *(classical entailment)*

- b. Did Bertrand stop smoking?
- ~> Bertrand used to smoke. *(presuppositional inference)*
- ~> Someone stopped smoking. *(classical entailment)*

(169) DEFINITE DESCRIPTIONS

- a. The Governor of Texas wants to secede.
- ~> There is a Governor of Texas. *(presuppositional inference)*
- ~> Someone wants to secede. *(classical entailment)*

- b. Does the Governor of Texas want to secede?
- ~> There is a Governor of Texas. *(presuppositional inference)*
- ~> Someone wants to secede. *(classical entailment)*

The difference between the presuppositional inferences and the classical entailments is clear; the presuppositional inferences triggered by the declarative (a)-sentences are preserved when these are converted into polar questions, i.e. the interrogative (b)-sentences. In contrast, the classical entailments are not. Now, if weak determiners, i.e. indefinite descriptions, are uniformly analyzed as asserting existence, adding an overt assertion of existence (i.e. an existential there-construction) should not affect which inferences are licensed. In other words, we should expect the putative existential entailments licensed by weak determiners to behave as classical entailments.

- (170) a. There are some rotten apples in the pantry. *(declaratives)*
- b. There are many rotten apples in the pantry.
- c. There are several rotten apples in the pantry.
- ~> There are apples in the pantry.

- (171) a. Are there some rotten apples in the pantry? *(interrogatives)*
- b. Are there many rotten apples in the pantry?
- c. Are there several rotten apples in the pantry?
- ~> There are some/many/several apples in the pantry.

As previously indicated, it seems that when the declarative existential-there sentence is converted into a polar question, the inference that there are apples

in the pantry is no longer intuitively available. This strongly suggests that it is a classical entailment. However, now consider the examples in (172) and (173).

- (172) a. Some apples in the pantry are rotten. *(declaratives)*
 b. Three apples in the pantry are rotten.
 c. Several apples in the pantry are rotten.

↪ There are some/many/several apples in the pantry.

- (173) a. Are some apples in the pantry rotten? *(interrogatives)*
 b. Are three apples in the pantry rotten?
 c. Are several apples in the pantry rotten?

↪ There are some/many/several apples in the pantry.

It seems quite natural to infer from the examples in (173) that there are apples in the pantry, i.e. the inference licensed by the declaratives in (172) appears to survive when converted into a polar question. And this of course contrasts the result of the conversions in (170) and (171). What should be emphasized here is that if the weak determiners in (172) simply asserted existence as in (170) and the existence of a non-empty restrictor set was a classical entailment of (172), this inference should not survive the conversion. In contrast, if in these cases the weak determiners are analyzed as presupposing the existence of a non-empty restrictor set, the availability of the existential inference in (172) is straightforwardly explained; the existential presuppositions project.

2.2 Conditionals

Another standard method for detecting presuppositional expressions is considering their inferential potentials in conditionals. It is widely agreed that a sentence which contains a presuppositional expression licenses a presuppositional inference even when it is embedded in the antecedent of a conditional. Again, the presuppositional inferences thereby contrast classical entailments and implicatures since the latter tend to disappear when embedded in such syntactic environments.

- (174) a. Bertrand stopped smoking.
 ↪ Bertrand used to smoke. *(presuppositional inference)*
 ↪ Someone stopped smoking. *(classical entailment)*

- b. If Bertrand stopped smoking, his boss gave him a raise.
 \leadsto Bertrand used to smoke.
 \leadsto Someone stopped smoking.
- (175) a. The Governor of Alaska resigned
 \leadsto There is a Governor of Alaska. (*presuppositional inference*)
 \leadsto Someone resigned. (*classical entailment*)
- b. If the Governor of Alaska resigned, she's running for president.
 \leadsto There is a Governor of Alaska.
 \leadsto Someone resigned.

Now, to demonstrate that there is a difference between existential-there sentences and sentences with weak determiners in subject position, we carefully consider which inferences these sentences license in relevant complex syntactic environments. And since the difference is quite subtle, I use a method devised by von Stechow (1998) to make it easier to detect. The method is to preface each conditional with a statement indicating that the speaker is agnostic about the existence of *F*s (where *F* is the restrictor for the determiner). This should have the following effect: If the speaker is openly agnostic about existence of *F*s, the existence of *F*s cannot be common ground. Hence, it should then be infelicitous for a speaker to use an expression which requires that the existence of a non-empty set of *F*s is already common ground. Each sentence in (176a-176c) should therefore be read as immediately following the initial assertion in (176).

- (176) I'm not sure whether there are any apples in the pantry.
 - a. If there are some apples in the pantry and they're ripe, we should bake an apple pie.
 - b. If there are several apples in the pantry and they're ripe, we should bake an apple pie.
 - c. If there is an apple in the pantry and it's ripe, we should bake an apple pie.

The discourse continuations in (176a-176c) are felicitous even though the speaker has openly declared herself agnostic about the existence of apples in the pantry. But, now compare these to the continuations in (177a-177c).

- (177) I'm not sure whether there are any apples in the pantry.
 - a. # If some apples in the pantry are ripe, we should bake an apple pie.
 - b. # If several apples in the pantry are ripe, we should bake an apple pie.

- c. # If an apple in the pantry is ripe, we should bake an apple pie.

The discourse continuations in (177a-177c) should be equivalent to the continuations in (176a-176c). That is, there should be interpretations of (177a-177c) available where the weak determiners take scope inside the if-clauses.⁵ But these interpretations are simply unavailable as the continuations in (177a-177c) are clearly infelicitous. And given the standard existential analysis, this cannot be explained.

In contrast, if these determiners are assumed to trigger existence presuppositions, the infelicity is to be expected. If the weak determiners require that an existential presupposition is common ground, it is no surprise that a speaker cannot first declare herself agnostic about the status of some proposition *p* and then proceed to use a term that requires that *p* is common ground.⁶

⁵ In fact, as already pointed out in Chapter III, the interpretation on which the determiner phrase is embedded inside the scope of the if-clause should not only be *compatible* with these sentences, it should be *mandatory* since if-clauses are scope islands.

⁶ While certain discourse particles, e.g. ‘but’ or ‘however’, seem to improve felicity judgments about (177a-177c), these discourse particles are prototypical *contrast markers*. Their discursive function is to contrast information already given with information to be conveyed. To explicate precisely the semantic and pragmatic effects of discourse particles such as ‘but’ or ‘however’ is a complicated affair, but observe that these discourse particles also succeed in improving felicity judgments for sentences that contain standardly accepted presuppositional expressions, e.g. factives and aspectual verbs.

- (1) I’m not sure whether Peter is having a party.
 a. ...but if John knows that Peter is having a party, we should go!
 b. # ...and if John knows that Peter is having a party, we should go!
 c. # If John knows that Peter is having a party, we should go!
- (2) I’m not sure whether Sue used to smoke.
 a. ...but if she stopped, she’s probably drinking more coffee.
 b. # ...and if she stopped, she’s probably drinking more coffee.
 c. # If she stopped, she’s probably drinking more coffee.

The contrast is perfectly clear. The (b) and (c)-sentences are downright infelicitous whereas the (a)-sentences are not. This, I take it, is not evidence that factive or aspectual verbs fail to trigger presuppositions. Rather, it is evidence that these discourse particles facilitate coherent interpretations. Explaining this data is a complicated task that I am unable to undertake here, but my tentative (albeit rather vague) suggestion is that these contrast markers shift modal salience; the contrast markers indicate to the interlocutors that the continuation should be interpreted at a possible world where the presupposition is satisfied. I.e. one could imagine that these discourse particles license

2.3 Modals

The final projection test that I consider here is projection under modal embeddings. Similar to the effects observed with if-clauses, presuppositional inferences tend to survive when the inference-licensing sentence is embedded under a modal. In contrast, classical entailments and implicatures seem to disappear in such syntactic environments.

- (178) a. Bertrand stopped smoking.
 ↪ Bertrand used to smoke. (*presuppositional inference*)
 ↪ Someone stopped smoking. (*classical entailment*)
- (179) a. Bertrand might/ought-to have stopped smoking.
 ↪ Bertrand used to smoke.
 ↯ Someone has stopped smoking.
- (180) a. The Governor of Alaska is running for president.
 ↪ There is a Governor of Alaska. (*presuppositional inference*)
 ↪ Someone is running for president. (*classical entailment*)
- (181) a. It's possible that the Governor of Alaska is running for president.
 ↪ There is a Governor of Alaska.
 ↯ Someone is running for president.

Given this behavior of presuppositional inferences, we should not expect existential there-sentences embedded under modals to give rise to any existential inferences. In contrast, if weak determiners are capable of triggering existence presuppositions, we *should* expect these to license existential inferences. And this is precisely what we find.⁷

- (182) I'm not sure whether there are any apples in the pantry.

local accommodation (see ch. III). Another example that supports this explanation is the following.

- (3) I might go to Disney Land with you tomorrow, but I'll never go again.
 (4) # I might go to Disney Land with you tomorrow, and I'll never go again.

Here the relevant interpretation of (3) is the modally subordinated reading, viz. '... but if I go, I'll never go again'. This explanation also seems to largely conform to Roberts' (1996) account of modal subordination.

⁷ I use different modal expressions in order to ensure that the putative presuppositional effects are not caused by incidental features of particular modals.

- a. There might be several apples in the pantry that are ripe.
- b. There ought to be a couple of apples in the pantry that are ripe.
- c. It’s possible that there are three apples in the pantry that are ripe.
- d. Perhaps there is an apple in the pantry which is ripe.

As expected, when embedded under modals, these existential-there sentences are perfectly felicitous even though the speaker has declared herself openly agnostic about the existence of apples in the pantry. Compare this to the discourse continuations in (183).

- (183) I’m not sure whether there are any apples in the pantry.
- a. # Several apples in the pantry might be ripe.
 - b. # A couple of apples in the pantry ought to be ripe.
 - c. # It’s possible that three apples in the pantry are ripe.
 - d. # Perhaps an apple in the pantry is ripe.

These discourse continuations in (183a-183d) are clearly infelicitous, which by now is fairly unsurprising. Again, if it is assumed that the existential inference triggered by a non-complex sentence such as (166) is explained simply by the presence of an existential quantifier, these inferences should disappear when the quantifier is embedded in the prejacent clause, i.e. in the scope of the modal — and yet they do not.⁸

3 The ‘Hey, Wait a Minute’-Test

The last data point in favor of the claim that weak determiners can trigger presuppositions considered here is the so-called ‘Hey, wait a minute!’-test (HWMT).⁹

The idea behind the (HWMT) is to exploit the fact that non-asserted components of content can be questioned in ways that asserted components of content cannot. If for example a speaker S_1 asserts the sentence (184), it is felicitous for an interlocutor S_2 to give a response prefaced by the locution ‘hey, wait a minute!’ only if S_2 is not intending to question what

⁸ Several people have now pointed out to me that the projection behavior for determiners embedded in conditionals and under modals is strictly speaking consistent with the view that existence is both presupposed and asserted (in the semantic sense). However, this analysis is inconsistent with the data in section A1.1 (conversions into yes/no questions) and A1.5, cf. below. Also, it remains unclear to me what would motivate such a view.

⁹ This type of test was proposed by Shanon (1976), but I here use a variation due to von Stechow (2004).

S₁ has outright asserted. Rather, S₂ must be questioning something else, for example something that S₁ appears to be presupposing.

- (184) The president of Buganda is bald.
- a. Hey, wait a minute. I didn't know that Buganda has a president!
 - b. ? Hey, wait a minute. I didn't know that the president of Buganda is bald!

The (HWMT) is thus a useful heuristic for distinguishing assertive from presuppositional components.

To apply the (HWMT) to the case at hand, let's consider how existential-there sentences compare to sentences with weak determiners in subject position.

- (185) There are three apples in the pantry which are ripe.
- a. ? Hey, wait a minute. I didn't know that there are apples in the pantry!
- (186) There are several apples in the pantry which are ripe.
- a. ? Hey, wait a minute. I didn't know that there are apples in the pantry!

The 'Hey, wait a minute'-responses sound somewhat awkward here and the most natural explanation is that since the content being questioned is genuinely asserted, it is inappropriate for the interlocutor to preface his responses with the 'hey, wait a minute'-phrase. But now consider the following sentences with weak determiners in subject position.

- (187) Three apples in the pantry are ripe.
- a. Hey, wait a minute. I didn't know that there are apples in the pantry!
- (188) Several apples in the pantry are ripe.
- a. Hey, wait a minute. I didn't know that there are apples in the pantry!

These seem significantly more natural and this is straightforwardly explained if it is assumed that the questioned content is taken for granted by the speaker, i.e. if it is presupposed. While the (HWMT) is not exactly a foolproof test and there are reasons to be skeptical of this data, it is nevertheless suggestive.

4 Analysis

Given the above results of applying the projection tests and the (HWMT) test to weak determiners, what should we conclude from this data? At the very least, it seems safe to conclude that the data reveals the impropriety of certain conversational moves. For example, we can conclude that it is inappropriate for a speaker to raise to salience the possibility of the non-existence of x s and subsequently proceed to predicate a property of x . Such impropriety would for example explain why contrast markers, e.g. ‘but’ or ‘however’, can sometimes be used to restore coherence.

More importantly, it also seems safe to conclude that for the orthodox existential analysis of weak determiners, e.g. the Russellian, there is an explanatory problem. The data reveals that there are contexts where certain discourse continuations are appropriate *only* if the speaker uses an existential-there sentence, namely contexts where the speaker asserts a hypothetical (using e.g. a conditional or a possibility modal) about x but where the actual existence of x is not already common ground. In other words, there is an important contrast between existential-there sentences and sentences containing weak determiners in subject position.

Since this strongly indicates that weak determiners can trigger existence presuppositions, I conclude that it is no obstacle to the solution to the existence problem (the solution presented in Chapter III.) that it relies on the assumption that some weak determiners trigger presuppositions.

4.1 Weak Determiners and Presupposition Accommodation

Büring (1996) has argued that any presuppositional analysis of weak determiners is prohibited from appealing to the standard view of presupposition, viz. the view according to which a presuppositional expression imposes the contextual requirement that its presuppositions must be common ground. Büring argues as follows.

It should be noted, however, that [the putative existential presupposition triggered by a weak determiner] cannot be a presupposition in the usual sense, e.g. knowledge assumed to be shared by speaker and hearer, since examples like [(189)] clearly do not involve any existential commitments on the part of the person asking. (Büring 1996: sec. 4)

- (189) a. Are there any cookies in this room?
 b. [SOME]_T cookies are [in the cupboard]_F

Since the purpose of the interlocutor’s assertion in (189a) is to determine whether the domain of cookies is empty, she cannot be interpreted as

presupposing that the domain of cookies is non-empty. But if the weak determiner in (189b) is interpreted as triggering an existence presupposition, it requires that it is common ground that the domain of cookies is non-empty. But given that it makes sense to even ask the question, it cannot be.

Now, in response to Büring, one could just assume that (189b) is not an instance of a presuppositional weak determiner, but even if we did assume this, Büring's argument is overlooking a very important factor, namely the phenomenon of *presupposition accommodation*. Consider the sentences below.

- (189) a. Are there any cookies in this room?
- (190) a. The cookies are in the cupboard. *definite*
 b. The [chocolate CHIP]_F cookies are in the pantry. *definite*
 c. All of the cookies are under the sink. *partitive*
 d. Bertrand puts them in the fridge. *pronoun*

If, as standardly assumed, *definites*, *partitives*, and *pronouns* trigger existence presuppositions, then it should be mysterious, according to Büring's view, that speakers can felicitously respond using (190a-190d). But this is not mysterious at all. Even if these expressions trigger existence presuppositions that are not antecedently established in the common ground, then assuming that the speaker is cooperative, the interlocutor is simply prompted to revise or *repair* the common ground so that it admits a sentence that presupposes a non-empty domain of cookies. In a word, the presuppositions are *accommodated*.¹⁰

While providing an exact characterization of the phenomenon of presupposition accommodation is exceedingly complicated, it is nevertheless practically indisputable that such a mechanism is employed by speakers of natural language. At least, if one believes that the notion of presuppositions is a sensible one, one needs a mechanism like accommodation in order to explain why uses of presuppositional expressions are sometimes informative.¹¹

Now, one particular reason why it is hard to provide an exact characterization of accommodation is that the procedure relies extensively on extra-linguistic facts. For example, while it is quite probable that the

¹⁰ The notion of accommodation is standardly attributed to Lewis (1979). I provide only a cursory introduction to the topic, but for detailed discussions of accommodation, see the excellent Beaver and Zeevat (2007) and von Stechow (2008).

¹¹ Some researchers have also used precisely this point to argue against the standard views of presuppositions, see e.g. Gauker (2008).

presupposition triggered in (191a) is accommodated in most contexts, the presupposition triggered in (191b) is not.

- (191) You were supposed to be here at 9.15. It's 10.30 now!
- a. Yes, I know. I'm sorry, my car broke down.
 - b. ?? Yes, I know. I'm sorry, my spaceship broke down.

This difference must almost certainly be explained in terms of the extra-linguistic fact that it's normal to own a car, whereas it's quite unusual to own a spaceship.

Now, it's widely agreed that discourse initial uses of weak determiner phrases are perfectly acceptable whereas uses of e.g. definite descriptions are often not, and this contrast is standardly assumed to be explained in terms of presuppositions; (singular) definite descriptions are subject to a clear conversational constraint, namely that it is common ground that the restrictor of the definite description has precisely one member.

But it is important to observe that this is a strong presuppositional requirement. Definites presuppose not only existence of a non-empty domain, but also existence of a *unique* individual in the domain satisfying the restrictive component. Since, in many contexts, it will not be automatically clear that the definite uniquely denotes, initiating a discourse using a definite description will in such contexts often be infelicitous. But weak determiners, in particular indefinites, are the obvious contrast. When these determiners trigger presuppositions, they are subject to a comparatively weak presuppositional requirement, namely *just* that their domain is non-empty, viz. the requirement that the interlocutors agree that *some* individual in the domain satisfies the restrictive component. This requirement is typically satisfied because it is, in the vast majority of cases, completely uncontroversial to assume that *some* individual has the relevant property. And it is therefore no surprise that discourse initial uses of weak determiners are felicitous – their presuppositions are in general antecedently established in the common ground or, if not, simply accommodated.

As an analogy, consider the determiner phrase 'every *F*'. Uses of this determiner phrase are often assumed to be licensed only if the domain of *F*s is non-empty, but discourse-initial uses of 'every' are nevertheless generally unproblematic. The most natural explanation is that in most cases, the presupposition triggered by 'every *F*' is antecedently established in the common ground and, if it is not, it is accommodated. In other words, it is the concomitant uniqueness presupposition triggered by definite descriptions that makes them hard to accommodate, not their existential presupposition.

In conclusion, assuming that weak determiners can trigger presuppositions is not inconsistent with the standard notion of presuppositions, because

(true) discourse-initial uses of e.g. indefinites — uses in contexts where it is *not* already common ground that the set of *F*s is non-empty — are (a) rare and (b) can plausibly be explained in terms of accommodation.

REFERENCES

- Abbott, Barbara 2010. *Reference*. Oxford Surveys in Semantics and Pragmatics, Oxford: Oxford University Press.
- Abusch, Dorit 1994. 'The Scope of Indefinites'. *Natural Language Semantics*, 2, 2: 83–135.
- Ariel, Mira 2001. 'Accessibility Theory: An Overview'. In Sanders, Ted, Schilperoord, Joost and Spooren, Wilbert (eds.) *Text Representation*, chap. 2. John Benjamin Publishing Company, pp. 29–89.
- Bach, Kent 1981. 'Referential/Attributive'. In Davis, S. (ed.) *Pragmatics: A Reader*. Oxford: Oxford University Press, pp. 17–32.
- 1987. *Thought and Reference*. Oxford: Clarendon Press.
- 2004. 'Descriptions: Points of Reference'. In Bezuidenhout and Reimer (2004): 189–229.
- 2007. 'Referentially Used Descriptions: A Reply to Devitt'. *European Journal of Analytic Philosophy*, 3, 2: 33–48.
- Barwise, Jon and Cooper, Robin 1981. 'Generalized Quantifiers and Natural Language'. *Linguistics and Philosophy*, 4: 159–219.
- Barwise, Jon and Perry, John 1983. *Situations and Attitudes*. Cambridge, Massachusetts: MIT Press.
- Bäuerle, Rainer 1983. 'Pragmatischsemantische Aspekte der NP-Interpretation'. In Faust, M., Harweg, R., Lehfeldt, W. and Wienold, G. (eds.) *Allgemeine Sprachwissenschaft, Sprachtypologie und Textlinguistik*. Tübingen: Narr.
- Beaver, David I. 1997. 'Presupposition'. In van Benthem, Johan and ter Meulen, Alice (eds.) *Handbook of Logic and Language*. Cambridge, Massachusetts: MIT Press.
- 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.
- Beaver, David I. and Clark, Brady Z. 2008. *Sense and Sensitivity: How Focus Determines Meaning*. Wiley Blackwell.

- Beaver, David I. and Zeevat, Henk 2007. 'Accommodation'. In Ramchand, Gillian and Reiss, Charles (eds.) *Oxford Handbook of Linguistic Interfaces*, chap. 16. Oxford: Oxford University Press, pp. 503–539.
- Bezuidenhout, Anne and Reimer, Marga (eds.) 2004. *Descriptions and Beyond*. Oxford: Oxford University Press.
- Bott, Lewis and Noveck, Ira A. 2004. 'Some utterances are underinformative: The onset and time course of scalar inferences'. *Journal of Memory and Language*, 51, 3: 437–457.
- Büring, Daniel 1996. 'A Weak Theory of Strong Readings'. *SALT*, VI.
- Chierchia, Gennaro 2004. 'Scalar Implicatures, Polarity Phenomena and the Syntax/Semantics Interface'. In Belletti, A. (ed.) *Structures and Beyond*. Oxford: Oxford University Press, pp. 39–104.
- Chierchia, Gennaro, Fox, Danny and Spector, Benjamin 2010. 'Scalar Implicature as a Grammatical Phenomenon'. In Portner, P., Maienborn, C. and von Stechow, K. (eds.) *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter (forthcoming).
- Cumming, Samuel J. 2007. 'Proper Nouns'. PhD thesis (Rutgers University).
- De Jong, Franciska and Verkuyl, Henk 1985. 'Generalized Quantifiers: The Properness of Their Strength'. In van Benthem, Johan (ed.) *Generalized Quantifiers in Natural Language*. Dordrecht: Foris Publications.
- Dever, Josh 2005. 'Semantic Value'. In *The Encyclopedia of Language and Linguistics*. Elsevier.
- Devitt, Michael 1981a. *Designation*. Columbia University Press.
- 1981b. 'Donnellan's Distinction'. *Midwest Studies in Philosophy*, 6: 511–524.
- 1997. 'Meanings and Psychology: A Response to Mark Richard'. *Noûs*, 31: 115–131.
- 2004. 'The Case for Referential Descriptions'. In Bezuidenhout and Reimer (2004), chap. 7. Oxford University Press, pp. 280–306.
- 2007. 'Referential Descriptions and Conversational Implicatures'. *European Journal of Analytic Philosophy*, 3, 2: 7–32.
- Devitt, Michael and Sterelny, Kim 1997. *Language and Reality*. Cambridge, Massachusetts: MIT Press.
- 1999. *Language and Reality: An Introduction to Philosophy of Language*. Cambridge, Massachusetts: MIT Press.
- Diesing, Molly 1992. *Indefinites*. Cambridge, Massachusetts: MIT Press.
- Donnellan, Keith 1966. 'Reference and Definite Descriptions'. *Philosophical Review*, 77: 281–304.
- Elbourne, Paul 2005. *Situations and Individuals*. Cambridge, Massachusetts: MIT Press.

- 2010. 'The Existence Entailments of Definite Descriptions.' *Linguistics and Philosophy*, Online First.
- Evans, Gareth 1982. *Varieties of Reference*. Oxford University Press.
- von Fintel, Kai 1994. 'Restrictions on Quantifier Domains.' PhD Dissertation, University of Massachusetts, Amherst.
- 1998. 'Evidence for Presuppositional Indefinites.' Unpublished ms.
- 2004. 'Would You Believe It? The King of France is Back! (Presuppositions and Truth-Value Intuitions)'. In Bezuidenhout and Reimer (2004): 315–342.
- 2008. 'What is Presupposition Accommodation, Again?' *Philosophical Perspectives*, 22, 1: 137–170.
- von Fintel, Kai and Heim, Irene 2007. 'Intensional Semantics.' Unpublished ms.
- Fodor, Janet D. 1970. 'The Linguistic Description of Opaque Contexts.' Ph.D. thesis, Massachusetts Institute of Technology. Published in 1976 by Indiana University Linguistics Club and in 1979 in the Garland Series *Outstanding Dissertations in Linguistics*.
- Fodor, Janet D. and Sag, Ivan 1982. 'Referential and Quantificational Indefinites.' *Linguistics and Philosophy*, 5: 355–398.
- Frege, Gottlob 1892. 'On Sinn and Bedeutung'. In Beaney, Michael (ed.) *The Frege Reader* (1997). Blackwell Publishing, pp. 151–172.
- 1918. 'The Thought'. In Beaney, Michael (ed.) *The Frege Reader* (1997). Blackwell Publishing, pp. 325–346.
- Gauker, Christopher 2008. 'Against Accommodation: Heim, van der Sandt, and the Presupposition Projection Problem.' *Philosophical Perspectives*, 22: 171–205.
- Geurts, Bart 1997. 'Good News about the Description Theory of Names.' *Journal of Semantics*, 14: 319–348.
- 1998. 'Presuppositions and Anaphors in Attitude Contexts.' *Linguistics and Philosophy*, 21: 545–601.
- 1999. *Presuppositions and Pronouns*. Oxford: Elsevier.
- 2007. 'Existential Import'. In Comorowski, I. and von Heisinger, K. (eds.) *Existence: Syntax and Semantics*. Springer, Dordrecht, pp. 253–271.
- Gibbs, Raymond W. Jr. 1985. 'On the Process of Understanding Idioms.' *Journal of Psycholinguistic Research*, 14, 5: 465–472.
- 1986. 'Skating on Thin Ice: Literal Meaning and Understanding Idioms in Conversation.' *Discourse Processes*, 9, 1: 17–30.
- Glanzberg, Michael 2005. 'Presuppositions, Truth Values, and Expressing Propositions'. In Preyer, G. and Peters, G. (eds.) *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.

- 2006. 'Quantifiers'. In Lepore, Ernest and Smith, Barry C. (eds.) *The Oxford Handbook of Philosophy of Language*, chap. 31. Oxford: Oxford University Press, pp. 794–822.
- 2007. 'Definite Descriptions and Quantifier Scope: Some Mates Cases Reconsidered'. *European Journal of Analytic Philosophy*, 3: 133–158.
- Grice, Paul 1989. 'Logic and Conversation'. In *Studies in The Way of Words*. Cambridge, Massachusetts: Harvard University Press, pp. 22–41.
- Grodner, Daniel J., Klein, Natalie M., Carbary, Kathleen M. and Tanenhaus, Michael K. 2010. "Some," and possibly all, scalar inferences are not delayed: Evidence for immediate pragmatic enrichment'. *Cognition*, 116: 42–55.
- Groenendijk, Jeroen and Stokhof, Martin 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.
- 1991. 'Dynamic Predicate Logic'. *Linguistics and Philosophy*, 14: 39–100.
- Hamblin, Charles L. 1973. 'Questions in Montague English'. *Foundations of Language*, 10: 41–53.
- Hausser, Roland 1983. *The Syntax and Semantics of English Mood. Questions and Answers*. Dordrecht: Reidel.
- Hawthorne, John and Manley, David forthcoming. 'The Reference Book (tentative title, early draft, 2008)'.
- Heim, Irene 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- 1983. 'On the Projection Problem for Presuppositions'. In Barlow, M., Flickinger, D. and Wescoat, M. (eds.) *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*. Reprinted in Partee and Portner (2002). Citations from that edition.
- 1990. 'Presupposition'. In van der Sandt, R. (ed.) *Presupposition, Lexical Meaning and Discourse Processes: Workshop Reader*. University of Nijmegen.
- 1991. 'Artikel und Definitheit'. In von Stechow, Arnim and Wunderlich, Dieter (eds.) *Semantik: ein internationale Handbuch der zeitgenössischen Forschung*. Berlin: Walter de Gruyter, pp. 487–535.
- 1992. 'Presupposition Projection and the Semantics of Attitude Verbs'. *Journal of Semantics*, 9: 183–221.
- 2008. 'Features on Bound Pronouns'. In Harbour, Daniel, Adger, David and Béjar, Susana (eds.) *Phi Theory, Phi-Features across Modules and Interfaces*, chap. 2. Oxford: Oxford University Press, pp. 35–57.
- 2010. 'Definiteness and Indefiniteness'. In Maienborn, Claudia, von Heusinger, Klaus and Portner, Paul (eds.) *Semantics: An International*

- Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter (forthcoming).
- Heim, Irene and Kratzer, Angelika 1998. *Semantics in Generative Grammar*. Blackwell Publishing.
- Higginbotham, James and May, Robert 1981. 'Questions, Quantifiers and Crossing'. *Linguistics Review*, 1: 41–79.
- Hintikka, Jaako 1969. 'Semantics for Propositional Attitudes'. In Davis, J. W., Hockney, D. J. and Wilson, W. K. (eds.) *Philosophical Logic*. Dordrecht: Reidel, pp. 21–45.
- Horn, Laurence R. 1972. *On the Semantic Properties of Logical Operators in English*. Ph.D. thesis, University of California, Los Angeles.
- 1989. *A Natural History of Negation*. Chicago: Chicago University Press.
- Kadmon, Nirit 2001. *Formal Pragmatics*. Blackwell Publishing.
- 2009. 'Contrastive Topics and The Focal Structure of Questions'. Unpublished ms.
- Kamp, Hans 1981. 'A Theory of Truth and Semantic Representation'. In Groenendijk, Jeroen, Janssen, Theo M.V. and Stokhof, Martin (eds.) *Formal Methods in the Study of Language, Part 1*. Mathematical Centre Tracts, pp. 277–322.
- Kamp, Hans and Reyle, Uwe 1993. *From Discourse to Logic*. Dordrecht: Kluwer.
- Kamp, Hans, van Genabith, Josef and Reyle, Uwe 2011. 'Discourse Representation Theory'. In *Handbook of Philosophical Logic (Second Edition)*, vol. 15, chap. 3. Springer, Dordrecht, pp. 125–387.
- Kaplan, David 2005. 'Reading 'On Denoting' on its Centenary'. *Mind*, 114: 933–1003.
- Karttunen, Lauri 1974. 'Presuppositions and Linguistic Context'. *Theoretical Linguistics*, 1: 181–194.
- Keenan, Edward L. and Stavi, Jonathan 1986. 'A Semantics Characterization of Natural Language Determiners'. *Linguistics and Philosophy*, 9: 253–326.
- Keshet, Ezra 2008. *Good Intentions: Paving Two Roads to a Theory of the De re/De dicto Distinction*. Ph.D. thesis, Massachusetts Institute of Technology.
- Kratzer, Angelika 1977. 'What 'must' and 'can' must and can mean'. *Linguistics and Philosophy*, 1: 337–355.
- 1981. 'The Notional Category of Modality'. In Eikmeyer, H.-J. and Rieser, H. (eds.) *Words, Worlds, and Contexts, New Approaches to Word Semantics*. Berlin: Walter de Gruyter, pp. 38–74.

- 1998. 'Scope or Pseudoscope? Are there wide scope indefinites?' In Rothstein, S. (ed.) *Events and Grammar*. Dordrecht: Kluwer Academic Publishers, pp. 163–196.
- Krifka, Manfred 1991. 'A Compositional Semantics for Multiple Focus Constructions'. In *Proceedings of SALT 1*. Cornell University Working Papers in Linguistics.
- 2001. 'For a Structured Meaning Account of Questions and Answers'. In *Audiatur Vox Sapientia. A Festschrift for Arnim von Stechow*. Akademie Verlag, pp. 287–320.
- 2004. 'The Semantics of Questions and the Focussation of Answers'. In *Topic and Focus: A Cross-Linguistic Perspective*. Dordrecht: Kluwer Academic Publishers, pp. 139–151.
- Kripke, Saul 1977. 'Speaker's Reference and Semantic Reference'. *Midwest Studies in Philosophy*, 2, 1: 255–276.
- 1980. *Naming and Necessity*. Cambridge, Massachusetts: Harvard University Press.
- 2005. 'Russell's Notion of Scope'. *Mind*, 114: 1005–1037.
- Lambert, Karel 1983. *Meinong and The Principle of Independence*. Cambridge: Cambridge University Press.
- Lasnik, Howard 1993. 'Existence Presuppositions and Background Knowledge'. *Journal of Semantics*, 10: 113–122.
- Lewis, David 1979. 'Scorekeeping in a Language Game'. *Journal of Philosophical Logic*, 8: 339–359.
- Ludlow, Peter 1995. 'The Logical Form of Determiners'. *Journal of Philosophical Logic*, 24, 1: 47–69.
- 2007. 'Stanford Encyclopedia of Philosophy: Descriptions'.
- Ludlow, Peter and Neale, Stephen 1991. 'Indefinite Descriptions: In Defense of Russell'. *Linguistics and Philosophy*, 14: 171–202.
- Ludlow, Peter and Segal, Gabriel 2004. 'On a Unitary Semantical Analysis for Definite and Indefinite Descriptions'. In *Bezuidenhout and Reimer (2004)*: 420–437.
- Maier, Emar 2006. *Belief in Context: Towards a Unified Semantics of De Re and De Se Attitude Reports*. Ph.D. thesis, Radboud Universiteit Nijmegen, Nijmegen.
- 2009. 'Presupposing acquaintance: a unified semantics for *de dicto*, *de re*, and *de se* belief reports'. *Linguistics and Philosophy*, 32: 429–474.
- May, Robert 1985. *Logical Form: Its Structure and Derivation*. MIT Press.
- Meinong, Alexius 1904. 'Über Gegenstandstheorie'. In Meinong, A. (ed.) *Untersuchungen zur Gegenstandstheorie und Psychologie*. Leipzig: Johann Ambrosius Barth.

- Milsark, Gary 1977. 'Towards an Explanation of Certain Peculiarities of the Existential Construction in English.' *Linguistic Analysis*, 3, 1: 1–29.
- Mostowski, Andrzej 1957. 'On a Generalization of Quantifiers.' *Fundamenta Mathematicae*, 44: 12–36.
- Musan, Renate 1995. *In the Temporal Interpretation of Noun Phrases*. Ph.D. thesis, MIT, Cambridge, Massachusetts.
- Neale, Stephen 1990. *Descriptions*. Cambridge, Massachusetts: MIT Press.
- 2001. *Facing Facts*. Oxford University Press.
- 2004. 'This, That, and the Other.' In Bezuidenhout and Reimer (2004): 68–182.
- 2005. 'A Century Later.' *Mind*, 114: 809–871.
- 2007. 'Heavy Hands, Magic, and Scene-Reading Traps.' *European Journal of Analytic Philosophy*, 3, 2: 77–132.
- Parsons, Terence 1980. *Nonexistent Objects*. New Haven: Yale University Press.
- Partee, Barbara H. 1989. 'Many Quantifiers.' In Powers, J. and de Jong, K. (eds.) *Proceedings of the 5th Eastern Conference on Linguistics*. Columbus, Ohio. Reprinted in Partee (2004).
- 2004. *Compositionality in Formal Semantics*. Oxford: Blackwell Publishing.
- Partee, Barbara H. and Portner, Paul (eds.) 2002. *Formal Semantics - The Essential Readings*. Oxford: Blackwell Publishing.
- Percus, Orin 2000. 'Constraints on Some Other Variables in Syntax.' *Natural Language Semantics*, 8, 3: 173–229.
- Pupa, Francesco 2008. *Ambiguous Articles – An Essay on the Theory of Descriptions*. Ph.D. thesis, The City University of New York, New York.
- Recanati, François 1986. 'Contextual Dependence and Definite Descriptions.' *Proceedings of the Aristotelian Society*, 87: 57–73.
- Reimer, Marga 1998. 'Donnellan's Distinction/Kripke's Test.' *Analysis*, 58, 2: 89–100.
- Reinhart, Tanya 1995/2006. *Interface Strategies: Optimal and Costly Computations*. Cambridge, Massachusetts: MIT Press.
- 1997. 'Quantifier Scope: How labor is divided between QR and choice functions.' *Linguistics and Philosophy*, 20: 335–397.
- Roberts, Craige 1987. *Modal Subordination, Anaphora, and Distributivity*. Ph.D. thesis, University of Massachusetts, Amherst, Massachusetts.
- 1996. 'Anaphora in Intensional Contexts.' In Lappin, Shalom (ed.) *Handbook of Contemporary Semantic Theory*, chap. 8. Blackwell Publishing, pp. 215–247.
- 1998. 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics.' In Yoon, Jae-Hak and Kathol, Andreas

- (eds.) *OSU Working Papers in Linguistics: Vol. 49: Papers in Semantics*, (1998 revision) edn. Dept. of Linguistics, Ohio State University.
- 2003. 'Uniqueness in Definite Noun Phrases.' *Linguistics and Philosophy*, 26: 287–350.
- 2004. 'Context in Dynamic Interpretation.' In Horn, Laurence R. and Ward, Gregory (eds.) *Handbook of Pragmatics*. Blackwell Publishing, pp. 197–221.
- 2009. 'Topic.' In Maienborn, Claudia, von Stechow, Klaus and Portner, Paul (eds.) *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter (draft – forthcoming).
- Romoli, Jacopo and Sudo, Yasutada 2009. 'De Dicto/De Re Ambiguity and Presupposition Projection.' In Riester, Arndt and Solstad, Torggrim (eds.) *Proceedings of Sinn und Bedeutung*, vol. 13. University of Stuttgart.
- Rooth, Mats 1985. *Association with Focus*. Ph.D. thesis, University of Massachusetts, Amherst.
- 1992. 'A Theory of Focus Interpretation.' *Natural Language Semantics*, 1.1: 75–116.
- 1996. 'Focus.' In Lappin, Shalom (ed.) *Handbook of Contemporary Semantic Theory*. Blackwell, pp. 271–297.
- Rothschild, Daniel 2006. 'Non-Monotonic NPI-Licensing, Definite Descriptions, and Grammaticalized Implicatures.' In Tancredi, Christopher, Kanazawa, Makoto, Imani, Ikumi and Kusumoto, Kiyomi (eds.) *Proceedings of SALT 16*. CLC Publications.
- 2007. 'Presuppositions and Scope.' *Journal of Philosophy*, 2: 71–106.
- Russell, Bertrand 1903. *Principles of Mathematics*. London: George Allen and Unwin.
- 1905. 'On Denoting.' *Mind*, 14: 479–493.
- 1917. 'Knowledge by Acquaintance and Knowledge by Description.' In *Mysticism and Logic*. London: George Allen and Unwin, pp. 152–167.
- 1957. 'Mr. Strawson on Referring.' *Mind*, 66: 385–389.
- Ruys, Eddy G. 1992. *The Scope of Indefinites*. Ph.D. thesis, Utrecht University, Utrecht.
- Sainsbury, R.M. 1979. *Russell*. London: Routledge and Kegan Paul.
- Salmon, Nathan 1982. 'Assertion and Incomplete Definite Descriptions.' *Philosophical Studies*, 42: 37–45. — reprinted in Salmon (2007).
- 1991. 'The Pragmatic Fallacy.' *Philosophical Studies*, 63, 83–97. — reprinted in Salmon (2007).
- 2007. *Content, Cognition, and Communication: Philosophical Papers II*. Oxford: Oxford University Press.
- Sauerland, Uli 2008. 'On the Semantic Markedness of Phi-Features.' In Harbour, Daniel, Adger, David and Béjar, Susana (eds.) *Phi Theory*,

- Phi-Features across Modules and Interfaces*, chap. 3. Oxford: Oxford University Press, pp. 57–82.
- Scha, Remko 1983. *Logical Foundations for Question Answering*. Ph.D. thesis, Groningen University.
- Schiffer, Stephen 1998. 'Descriptions, Indexicals, and Belief Reports (But Not the Ones You Expect)'. In Ostertag, Gary (ed.) *Definite Descriptions: A Reader*. Cambridge, Massachusetts: MIT Press, pp. 369–393.
- 2005. 'Russell's Theory of Definite Descriptions'. *Mind*, 114, 1135–1183.
- Schoubye, Anders J. 2009. 'Descriptions, Truth Value Intuitions, and Questions'. *Linguistics and Philosophy*, 32, 6: 583–617.
- Schwarz, Florian 2009. *Two Types of Definites in Natural Language*. Amherst, Massachusetts: Graduate Linguistics Students' Association, Dept. of Linguistics, University of Massachusetts, Amherst.
- Schwarzschild, Roger 2002. 'Singleton Indefinites'. *Journal of Semantics*, 19, 3.
- Searle, John R. 1979. *Expression and Meaning*. Cambridge: Cambridge University Press.
- Shanon, Benny 1976. 'On the Two Kinds of Presuppositions in Natural Language'. *Foundations of Language*, 13: 247–249.
- Simons, Mandy 2010. 'A Note on Projection and Local Implication'. Unpublished ms.
- Simons, Mandy, Roberts, Craige, Tonhauser, Judith and Beaver, David I. 2009. 'Presuppositions, Conventional Implicature, and Beyond: A unified account of projection'. In Klinedinst, Nathan and Rothschild, Daniel (eds.) *Proceedings of Workshop on New Directions in the Theory of Presuppositions*. ESSLI 2009.
- 2010. 'What Projects and Why'. In *Proceedings of SALT 20*. pp. 309–327.
- Soames, Scott 1986. 'Incomplete Definite Descriptions'. *Notre Dame Journal of Formal Logic* — reprinted in Soames (2009), 27: 349–375.
- 1994. 'Donnellan's Referential/Attributive Distinction'. *Philosophical Studies* — reprinted in Soames (2009), 73: 149–168.
- 2009. *Philosophical Essays: Natural Language, What It Means and How We Use It*. Princeton University Press.
- Spector, Benjamin 2007. 'Scalar Implicatures: Exhaustivity & Gricean Reasoning'. In Aloni, M. and Dekker, P. (eds.) *Questions in Dynamic Semantics, Current Issues in the Semantics/Pragmatics Interface*, vol. 17. Oxford: Elsevier, pp. 225–250.
- Stalnaker, Robert C. 1970. 'Pragmatics'. *Synthese*, 22, 1/2: 272–289.

- 1974. 'Pragmatic Presuppositions.' In Munitz, Milton and Unger, Peter (eds.) *Semantics and Philosophy: Essays*. New York: University Press.
- 1978. 'Assertion.' In Cole, P. (ed.) *Syntax and Semantics 9*. New York: New York Academic Press, pp. 315–32.
- 1998. 'On The Representation of Context.' *Journal of Logic, Language, and Information*, 7: 3–19.
- 2002. 'Common Ground.' *Linguistics and Philosophy*, 25, 701–721.
- Stanley, Jason 2002. 'Nominal Restriction.' In Preyer, Gerhard and Peter, Georg (eds.) *Logical Form and Language*. Oxford: Clarendon Press, pp. 365–388.
- Stanley, Jason and Szabó, Zoltan Gendler 2000. 'On Quantifier Domain Restriction.' *Mind and Language*, 15: 219–261.
- von Stechow, Arnim 1989. 'Focusing and Background Operators.' In *Arbeitspapier Nr. 6*. Fachgruppe Wissenschaft, Universität Konstanz.
- 1991. 'Current Issues in the Theory of Focus.' In von Stechow and Wunderlich, D. (eds.) *Semantik/Semantics: An International Handbook of Contemporary Research*. de Gruyter, pp. 804–825.
- Strawson, Peter F. 1950. 'On Referring.' *Mind*, 59: 320–344.
- 1952. *Introduction to Logical Theory*. London: Methuen.
- 1964. 'Identifying Reference and Truth-Values.' *Theoria*, 30: 96–118 (Reprinted in P.F. Strawson *Logico-Linguistic Papers*. 1971. Ashgate. pp. 57–73).
- Swinney, D. and Cutley, A. 1979. 'The access and processing of idiomatic expressions.' *Journal of Verbal Learning and Verbal Behavior*, 18: 523–534.
- Szabó, Zoltan Gendler 2000. 'Descriptions and Uniqueness.' *Philosophical Studies*, 101: 29–57.
- Szabolcsi, Anna 2001. 'The Syntax of Scope.' In Baltin, Mark and Collins, Chris (eds.) *The Handbook of Contemporary Syntactic Theory*, chap. 19. Oxford: Blackwell, pp. 607–633.
- van Rooy, Robert 2006. *Attitudes and Changing Contexts*. Springer.
- van der Sandt, Rob 1992. 'Presupposition Projection as Anaphora Resolution.' *Journal of Semantics*, 9: 333–377.
- Wettstein, Howard 1981. 'Demonstrative Reference and Definite Descriptions.' *Philosophical Studies*, 2: 241–257.
- 1983. 'The Semantic Significance of the Referential-Attributive Distinction.' *Philosophical Studies*, 44: 187–196.
- Wiggins, David 1980. *"Most" and "All": Some comments on a Familiar Programme, and on the Logical Form of Quantified Sentences*. Routledge and Kegan Paul, pp. 318–346.

- Wilson, G. 1978. 'On Definite and Indefinite Descriptions.' *The Philosophical Review*, 87: 48–76.
- Winter, Yoad 2001. *Flexibility Principles in Boolean Semantics*. Cambridge, Massachusetts: MIT press.
- Wittgenstein, Ludwig 1953. *Philosophical Investigations*. Oxford: Wiley Blackwell.
- Yablo, Stephen 2006. 'Non-Catastrophic Presupposition Failure.' In Thomson, J. J. and Byrne, Alex (eds.) *Content and Modality: Themes from the Philosophy of Robert C. Stalnaker*. Oxford University Press.
- 2009. 'Must Existence-Questions Have Answers?' In Chalmers, David J., Manley, David and Wasserman, Ryan (eds.) *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press, pp. 507–525.